Powers of Perceptual Control:

An Inquiry into Language, Culture, Power, and Politics

M. Martin Taylor
Martin Taylor Consulting
mmt@mmtaylor.net

We thought we understood everything
but then we got more data
and saw how naïve we were.
Malcolm Collins (quoted in Science 9 Nov. 2018, p626)

Every sentence I utter must be understood
not as an affirmation, but as a question.
Niels Bohr

The first principle is that you must not fool yourself.
And you are the easiest person to fool.
Richard Feynman

Volume 2

Creativity, Consciousness, Communication

We build with what we have.
Others have different tools.
We talk, we walk,
We teach, we learn.

We invent upon inventions.
Others learn different things.
And invent different tools.
Do inventions lead us astray?

This is a work in progress, and the Table of Contents may not be up-to-date.

Table of Contents

**Overview of Volume II

Volume II has four conceptual parts, starting in Chapters II.1 and II.2 with concepts of autocatalytic and homeostatic loops, which though apparently leading toward stability nevertheless provide creativity that can sometimes generate what I call “Constructive Revolutions”. In the course of this section we use the concepts to suggest a plausible evolutionary approach to the constructive revolution that perhaps produced the first life, following ideas by Stuart Kauffman.

Chapters II.3 through II.7 form the second part of Volume II. They consist of the development of the measure called “rattling” and the development and refinement of conscious perceptual control analogous to the crumpling of paper in a process that we call, big surprise, “crumpling”. I argue that using the crumpling analogy leads directly to PCT-based insights, for which I use as examples the first “words” of a baby communicating with its mother, which we follow in detail in the fourth part of Volume 2, potential avalanches of reorganization following unexpected shocks, and the sometimes unexpected behaviour of teenagers.

Crumpling appears to be closely related to category perception, verbal language, and conscious control leading to reorganization, whereas rattling measures the likelihood that collectives will reorganize and how, statistically, the reorganization will be manifest. The measure itself is of the same nature as “variance” or “uncertainty”, and is widely applicable to the development of structure in general. The structures of most interest to us are of the interactions among control units in the Powers hierarchy and among language elements used in protocols as well as internal thoughts.

The third part of Volume 2, Chapters 22 through 24, is less theoretical and more practical. The reader might not lose too much in a quick scan of the potential “Powers of Perceptual Control” by skipping the first two parts of Volume 2. The third part, Chapters 22 through 24, deals with the initiation and development of language within a real human family and then in a family of“synthetic people” that starts with a small formal language programmed into an interactive robot, which develops into a language with all the flexibility of the languages real humans use.

The final part of Volume 2, Chapters 25 and 26, develops the concept of the “protocol” motif and the control of belief that was introduced in Volume 1. Control of belief is seem as the “syntax” of all protocols. This leads to discussions of various uses of protocols for cooperative and deceitful communications, and the end of Volume 2.

Part 5: Networks of Creative Stability

Chapter xII.1 concerns the use of energy in control. Energy is supplied to a control loop by the disturbance, and it takes energy supplied from outside at low entropy to dispose of the entropy that would be introduced by an unopposed disturbance.

In Chapter II.2 we introduce the "Autocatalytic Network". Autocatalysis also underlies many of the social and cultural phenomena that we discuss in Volume 2 of the book. The Autocatalytic Network may be a homeostatic loop or network, which in turn might be or become a control loop. Evolution tends to favour homeostasis over autocatalysis in general, but we will argue that homeostasis in the biochemical structure of a living organisms coexists and complements perceptual controlling through action on the external environment. Both autocatalysis and Collective Control, which we delay discussing until Volume 2, are basic conceptual structures, in much the same way that much of mechanics is based on the underlying properties of algebra and calculus.

Chapter II.2 discusses a nano-scale universe of dramatically oversimplified microbes and uses them to describe infrastructure, and to further develop the concepts of energy and entropy in order to explain their importance for the later parts of the book. We discuss them and their role in evolution of structure through the model of a trivial microbe modelled after Powers's "e-coli" process of reorganization. This cartoon e-coli evolves new capabilities through simple errors of production or of memory.

Chapter II.3 develops the concept of the autocatalytic loop or network further, considering a particular kind of autocatalytic loop, the homeostatic loop. One form of homeostatic loop is a control loop. We consider a version of the microbes introduced in Chapter II.2, showing how their membranes might form active surfaces that permit control loops to control external variables from within. We begin to look at real biochemical homeostatic loops and conjecture how they might be related to the perceptual control hierarchy in real organisms, including humans.

Many of the ideas and the general trend of the constructs presented in these chapters, including catalysis in a primeval soup, were foreshadowed by Oparin (1927). We come at them, however, from a different direction than Oparin, following Kauffman (1995) and separately Prigogine and his colleagues (e.g. Prigogine, 1947). The ideas parallel those of many experts in the chemistry of pre-life and early life in that they invoke autocatalytic processes in a "primeval soup". Indeed, as of the end of October 2018, the Wikipedia page on "Abiogenesis" contained over 300 references, many of them on the possible development of a RNA world through autocatalysis in a primeval soup more complex than the one used in Chapters 14 and 15.

The objective in those Chapters is not to propose anything new, but to show how autocatalysis leads to homeostasis and homeostasis leads to control by a very natural and highly probable set of events that are independent of the actual chemistry of the processes. If you want to learn the current ideas about the origin of life, you must look elsewhere. I assume one thing only, that life requires the existence of control, for entropic reasons. Control protects the structures of life from damage by the onslaughts of the exterior environment, where "exterior" means different things depending on the scale of one's interest. It could mean exterior to some component of a cell, exterior to a whole cell, to an organ of a body, or to the entire body of a living organism such as a bacterium, a tree, a bee, or a person. Later in the book we will consider whether it applies also to social structures.

An underlying autocatalysis may remain hidden when we observe phenomena, just as controlled perceptions are hidden when we see the environmental effects of controlling them. When explanations are required, however, what is hidden usually comes to the surface. We introduce autocatalysis here in the context of our nano-scale universe so that functionally identical processes may later be apparent in the workings of the vastly more complex organizations and interactions of contemporary animals and plants, leading to the infrastructure of human societies.

Autocatalytic processes are not control, but some autocatalytic loops may be control loops, because some autocatalytic loops are homeostatic, and some homeostatic loops are control loops. An argument from probability, due to Kauffman (1995), indicates that if there is a sufficient variety of elemental units of any kind that can combine to form complexes in the way that atoms of different kinds can combine to form molecules whose properties are not the sum of the properties of the atoms, then autocatalytic loops will eventually form if some of these combinations can catalyze the production of other complexes. The loops will produce an exponentially increasing variety of complex products, any of which may participate in further expansions of the original loop into an ever-growing network, from which independent loops may spin off to initiate their own family of descendant autocatalytic networks.

An autocatalytic network that consumes its elemental units — its "food" — in producing new products at an exponentially increasing rate will eventually run out of food, no matter how much there is available. But some autocatalytic loops are homeostatic, maintaining a constant production rate of product, stabilizing the concentration of product in their environment when the rate of decay of the product matches the rate of its production. Of the homeostatic loops, some are control loops, in which the concentration of one product is determined by the production rate of another, just as in a "standard" PCT control loop the rate of neural firings called a "perceptual signal" is adjusted by affecting some other variable whose value determined the perceptual firing rate.

So let us begin our delving into the substructure of social systems by using the micro- and nano-scale environment as a starting place.

Chapter II.1. Autocatalysis: Inanimate Creativity

This Chapter lays the groundwork for a form of self-organization that may be regarded as an analogy to Richardson's "Little whorls" that are perpetually being spun off the "Big whorls" in an avalanche of ever increasing complexity until they are collectively called "viscosity." The same functional processes can be seen at work in much that is discussed later in the book Another conceptual grandfather of the following Chapters is Oparin (1927), who described the characteristic of a primeval soup, and the formation of membranes and cell-like enclosures, which we come to eventually in Section II.3.4. If the reader wishes to skip the rather complex details, here is a brief abstract, which, with luck, will be sufficient to allow the principles to be evident in the following Chapters.

The basic process underlying these creative loops is called "catalysis" in chemistry. The chemical use of the word is well known, and is often imagined to be its only use, but catalysis means something more fundamental, an effect that will be called on throughout much of them rest of this book. In this more general usage, a "catalyst" is something — an event, an object, a shape, whatever — that makes more likely for something to happen that might have happened occasionally without the catalyst but happens relatively easily in the presence of the catalyst, which itself is unaffected by being so used.

For example, consider big city life in the mid-19th century. Some city-dwellers might have eaten fresh vegetables they grew in their gardens, but few poor people would have had that opportunity. At some point, a railway line was laid between the big city and a distant market town where local farmers had habitually come together to sell their vegetables to each other. Now, the railway line allowed the farmers to send their fresh vegetables to the big city, where they could be made available to most people. The railway line was unaffected by the transmission of vegetables in fresh condition to the city, but the ability of a city-dweller to enjoy a variety of vegetables straight from the ground was much eased. The railway line was a catalyst.

To develop the concept of autocatalysis, however, we use the chemical case, because it is the area in which the concept of catalysis is normally used. If two things come together to form a complex that has properties that could not be simply deduced from the properties of the individual things, we call the process a "reaction". Reactions have rates that depend on how likely it is for those two things to encounter each other and how likely they are to react when they do meet. The presence of a catalyst increases that rate. The presence of an "anticatalyst" reduces it.

A catalyst "C" is unlikely to be of the same nature as the reacting units "A" and "B". If A and B are simple cells, C could be the presence of light or of a particular chemical, or it might even be the existence of a solid surface along which A and B cells might sometimes move in two dimensions rather than three. C might be anything that has a physical shape into which particular elementary units fit only in a way that enhances their probability of combining. Folded proteins function largely because of their external shape, for example.

Any product created in a reaction has a chance of being itself a catalyst or anticatalyst for some other reaction. If there are enough different elementary components that can react, then it becomes almost certain that some complex product will catalyze some reaction. Eventually it becomes almost certain that a loop will form in which a chain of catalyzed reaction products catalyzing other reactions will return so that a complex product catalyses the original reaction (Kauffman, 1995).

At that point, the loop becomes self-stabilizing, and the rate of production of all the complex products of its components increases explosively. That is autocatalysis, and the loop is an autocatalytic loop. When there are sufficient instances of the product complexes, they also may participate in reactions to form even more complex structure that were never seen before — inanimate creations, or inventions.

The proliferations of newly “invented” complex products enhances the probability that other autocatalytic loops will form. All of the autocatalytic loops will develop complex branching and re-entrant structures, so the word “loop” becomes misleading. I call such sub-networks of mutual catalytic support autocatalytic “clusters”. Typically, at some point the whole soup is dominated by one large cluster, though there may still exist isolated small clusters and isolates that neither catalyze or are catalyzed by anything yet created in the soup.

Autocatalytic clusters can interact, both supportively and antagonistically, because all the complex products have the possibility of acting as catalysts or anticatalysts to reactions in clusters other than the one from which they formed. If a complex produced by cluster A is an anticatalyst for a reaction in cluster B, the productivity rate of cluster B will be reduced. If it is sufficiently reduced that B no longer contains even one autocatalytic loop, B ceases to exist. But this may not happen, and a complex produced by B may happen to be an anticatalyst for a reaction in cluster A, with the result that both stabilize at some intermediate value of productivity.

Furthermore, the “cores” of the two clusters will become more separate in whatever measure happens to be appropriate, just as we argued in Section 8.1 would be that case for mutually inhibiting perceptions. At the level of concepts, we see the increasing divergence between US Republican and Democrat perceptions of the World as an important example, both sides having internally consistent but totally incompatible perceptions of many facts, almost all of which they perceive only because someone they trust told them about it, and it fits with what they already perceived to be true. We explore this problem in more depth in Section III.7.1.

The opposite may happen. A complex product of cluster A may catalyse a reaction in cluster B and vice-versa, bringing their cores closer and enhancing the productivity of each. They react with each other in a way functionally identical to the reactions between pairs of their elementary units. If there are enough different clusters, they may act as elementary units in an autocatalytic network of clusters, creating a novel level of structure. Such level building may continue, much as does the building of control levels in the control hierarchy, as discussed in Section I.5.5 when we used the example of controlling a “chair location” perception instead of separately controlling perceptions of the locations of its parts.

Finally, catalysis can enhance the rate of a reaction so that it increases the steady-state concentration of the complex product only if the product is unstable. On average, a complex structure is of lower entropy than is the unstructured mix that contains the components. If it were not, the components would preferentially form the product whenever they come together, and the steady state concentration would be limited only by the availability of components. Catalysis could not increase this limit. Accordingly, each of the reactions in an autocatalytic loop absorbs energy. The loop can exist only in a through energy flow that allows entropy to be exported, typically in the form of heat.

Remember that an autocatalytic loop is not a control loop. The loop has no inside to complement an outside environment. It exists all on the same environmental “soup”. The “signals” that are the flows of material between stores of elementary units are not the information-conveying “signals” that are manifest in different forms in different places of a control loop. But it can be a homeostatic loop, if some of the “catalysts” actually inhibit reactions rather than promoting them. In Chapter II.2, we see that some homeostatic loops actually are control loops.

The entire autocatalytic network exists in a single physical space, with no “membrane” to distinguish inside from outside, and nothing at one point in the autocatalytic network is compared with or depends on anything at another point, in the way perception is compared with reference in a control loop and depends on a state in the environment on which the output of the control loop acts. They are different beasts — but not unrelated. As Oparin (1927) described, some reactions to combine to produce materials that form membranes of a kind. Section II.2.4 suggests how these membranes might not be passive, but might instead act as intermediaries that allow control of a concentration inside to stabilize a concentration outside — control of a “perception” by action on an environmental variable. Such membranes are permeable to information.

The autocatalytic network idea is a further extension of the tensegrity ideas developed in Chapter I.8, in which control effects flow through the hierarchy and the environment in a complex feedback network of many control loops. It is a true emergent, not just of control, but inclusive of control, building on tensegrity, as tensegrity builds on stiffness, and stiffness on the details of the control process. Though this may be true conceptually, the physico-chemical autocatalytic network itself cannot embody control, whereas the homeostatic loops and networks we treat late in this Chapter and in the next may do so.

II.1.1 Catalysis and Autocatalytic Networks

If our industrial and formal infrastructure exists because it has proved on average to ease rather than to inhibit control, what then of bacteria and other single-cell organisms? At one time all life was single-celled, each cell surviving to reproduce or not according to its luck in acting effectively in its chemical and social environment (the “soup”). What advantage might be found in joining together to form multi-cellular organisms in which particular structural relationships among cell types (motifs) are replicated many times over within any single individual and recur in many individuals?

The slick answer to this is always the same: the types of cells that combined to form permanent multi-cellular structures, such as the double helix of DNA that could be replicated, survived to reproduce better in those structures than did those same cell types when acting by themselves. This does not imply that the structures, as abstract patterns, reproduced themselves using the cells, any more than it implies that the cells reproduced themselves in the replicated structural patterns.

The units that reproduce might not be alive. We will argue that even structures of non-living elements could exist and reproduce. One such class, and perhaps the only such class, is the “autocatalytic network”, espoused by Kauffman (1995), and in a slightly different form by Powers (1995). In the early stages, there might not have even been a structural template, but the existence of one autocatalytic network automatically generates more autocatalytic networks that exactly replicate the original, or nearly so. In the “soup”, reactions such as A+B→AB go both ways, and settle on a ratio between A, B, and AB concentrations that depends on the probabilities that an A and a B will meet within some unit of time and that an AB will decay away during that same unit of time, having been created in a previous A-B interaction. There is, therefore, an expected ratio of A or B to AB in the soup.

A greater than expected occurrence of AB combinations might mean only that cell type A would be more likely to produce type A descendants when it was combined with a type B cell than when alone. Later, any mutated type A1 cells that attached more easily to type B would have even more A1 descendants in the soup than original type A cells. The connection mechanism, whether it be a jig-saw fit between the component cells or some chemical catalytic effect, would have been inherent in the underlying chemical and physical properties of the soup, without there being any template for the creation of the AB structure.

I said that this was a “slick answer”, and it is, because it offers no mechanism for the combination to work. But then, neither does a similar statement, that chemical C is a catalyst that promotes the interaction X+Y→XY over the reverse, the decay of XY back into X and Y, suggest a mechanism for catalysis. In neither case does the lack of statement mean that the effect fails to occur. If it does occur, then the ordinary processes of evolution dictate the future progress of the story.

If A cells reproduce better when associated with type B cells than alone, then B can be considered as a catalyst for the propagation of descendants of A, where the reverse "reaction" would be the lonely death of a type A cell. When As link with Bs, the result is more of type A in the future soup than would have been the case without any B type in the soup. Furthermore, any mutation to A1 that enhanced the ease with which the cell would find and link with a B would increase the enhancement of births of A1 type cells over their demise, increasing the density of A1-type over the ancestral As that simply linked to a B that came close enough in the right orientation. This is even more true of the A2 descendants that actively approached B cells, besides linking with them more easily than did the original A-type.

At this point, A2 — the B-seeker — may seem like a predator that has B as prey, but it may not be so. If A2 survives well in the company of B, at least long enough to reproduce, the mutation will not survive long in the soup if the linking process destroys the accompanying B. Indeed, the assumption that A2 in an A2-B linkage survives better than A2 alone carries the implication that a B1 that participates easily in the complex will survive better than a B that lives alone. A and B would be mutual catalysts, creating more numerous A1-B, A2-B and especially A2-B1 complexes, each of which is more stable than its predecessors.

Kauffman (1995), following Prigogine and colleagues (e.g. Prigogine and Lefever, 1975), described a cascade of catalytic reactions like this that produces ever more varieties of ever more complex molecules in a chemical soup. The units in his initial soup were conceived as simple chemical species, such as elements. In this soup, C might catalyze the A+B→AB reaction, and AB might catalyze a P+Q→PQ reaction that no single element in the original soup might catalyze. Figure q14.1 shows a short sequence of these catalytic effects, and suggests possible ways in which products might catalyze earlier reactions in the sequence, forming a loop. The biochemical control loop of Figure II.3.1 is an example of such a short autocatalytic loop.

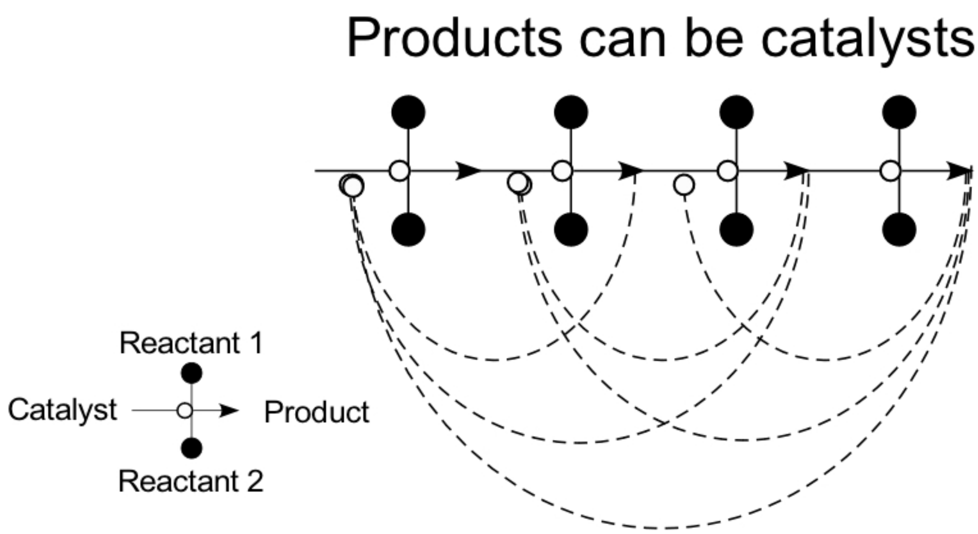

Figure q14.1 An autocatalytic sequence, in which the product of one reaction catalyzed by some component of the soup adds to the soup and catalyzes another reaction, and so on through four reactions, each adding a new product to the soup. Any of these products might possibly act as another catalyst for an earlier reaction in the sequence, forming an autocatalytic loop that includes any reactions in the sequence between the one catalyzed by the product and the one that produces the product.

If there were enough different elements in the soup and even a very small chance that a particular element might catalyze any specific reaction, Kauffman showed that an "autocatalytic loop" would almost inevitably occur, in which by a series of catalytic steps the existence of AB would catalyze the production of more AB so long as the soup contained a sufficient concentration to the basic As and Bs (Figure II.2.3). The same would be true of the product of every catalyzed product in the loop. The left panel of Figure II.2.1 suggests is simple example of a short (5-step) autocatalytic loop, while the right panel suggests how such a loop could induce the formation of complex products of reactions that would be possible but would seldom or never be seen in the original soup.

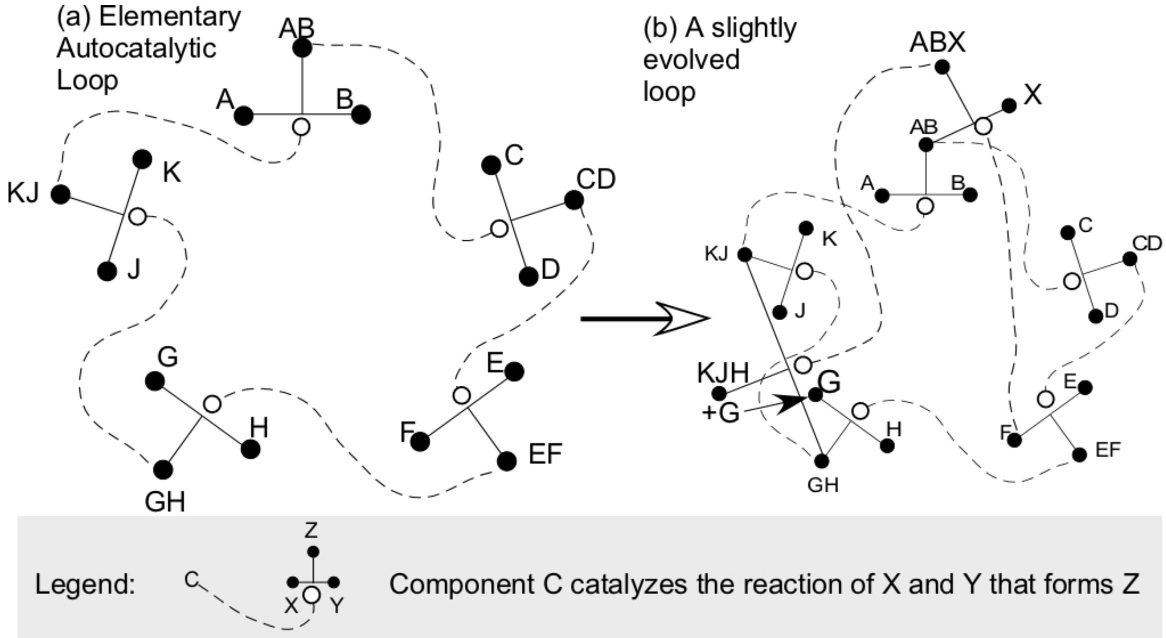

Figure II.2.1 (Left) A simple autocatalytic loop, in which each reaction is catalyzed by the product of a different reaction, and the random catalytic relationships complete a cycle so that each reaction in the loop effectively catalyses itself. (Right) any of the reaction products (and any of the basic elements) could participate in other reactions, producing more complex products. Here, AB can react with X to form ABX, and KJ can react with GH to form KJH, leaving a G isolated, which increases the concentration of G (this is a prototype of a purification reaction for producing pure G). The solid connecting lines do not imply a mechanical relationship, but rather, they imply a possibility for the indicated reaction to occur if an A and B happen to come into an appropriate juxtaposition to form AB. The catalyst merely changes the probability that the reaction will occur when they do come together, whether increasing it (catalysis) or decreasing it (anti-catalysis).

Each such loop increased the number and amounts of all the products of the reactions in the loop, at least until the supply of basic materials was depleted. Each new combination such as XY might catalyze a reaction that none of the basic elements could catalyze. That reaction might involve products that existed in appreciable concentrations only because they were the result of a catalyzed reaction in the loop. XY might, for example, catalyze a reaction AB+PQ→ABP + Q, in which both AB and PQ might have been rare and ABP even rarer had the parent reactions not been catalyzed within some autocatalytic loop.

Second-order creations such as ABP do not need to act as catalysts in the original loop, but nevertheless, in the soup they will exist in much greater concentrations because of the catalysis in the basic loop that produces relatively high concentrations of their components AB and PQ, increasing their likelihood of meeting and reacting together.

Figure II.2.2 places the evolved loop of Figure II.2.1 into a larger "soup" of potential reactions, which occur at one rate when alone in the soup, and at a different rate when catalyzed. It is important to remember that the word "rate" does not apply to the individual event in which one A meets one B and produces one AB complex, but to the number of A+B→AB reactions that occur per second in some volume of the soup.

Before the catalytic loop formed, occasional ABs might have been formed by the random meetings of an A and a B, as might occasional PQs, but they would have been much rarer without the catalysts that biased the reversible reactions in the direction of combination rather than dissociation. Quite probably it would be a very long time between instances when an AB would have met a PQ, and they might not have combined very often when they did meet. But once the catalytic loop formed, there would be many more ABs and PQs available to form ABP+Q, and since the product XY that was catalyzed in the loop was readily available as a catalyst, ABP might cease being rare-to-nonexistent and become suddenly common.

Each such chain of catalysis introduces a new or newly common product into the mix, which itself might catalyze some reaction among the ever increasing number of possible reactions. The number of complex products tends to increase exponentially over time (Figure II.2.2), because the more products there are, the more possibilities exist for new catalyzed reactions and different autocatalytic loops. In Figure II.2.2, the CD complex can be catalyzed by the G elementary unit to form many CCD complexes, which themselves would catalyze a reaction ABX+Y→ABXY, but ABX would almost never exist if the A +B→AB had not been catalyzed by KJ as part of the initial autocatalytic loop.

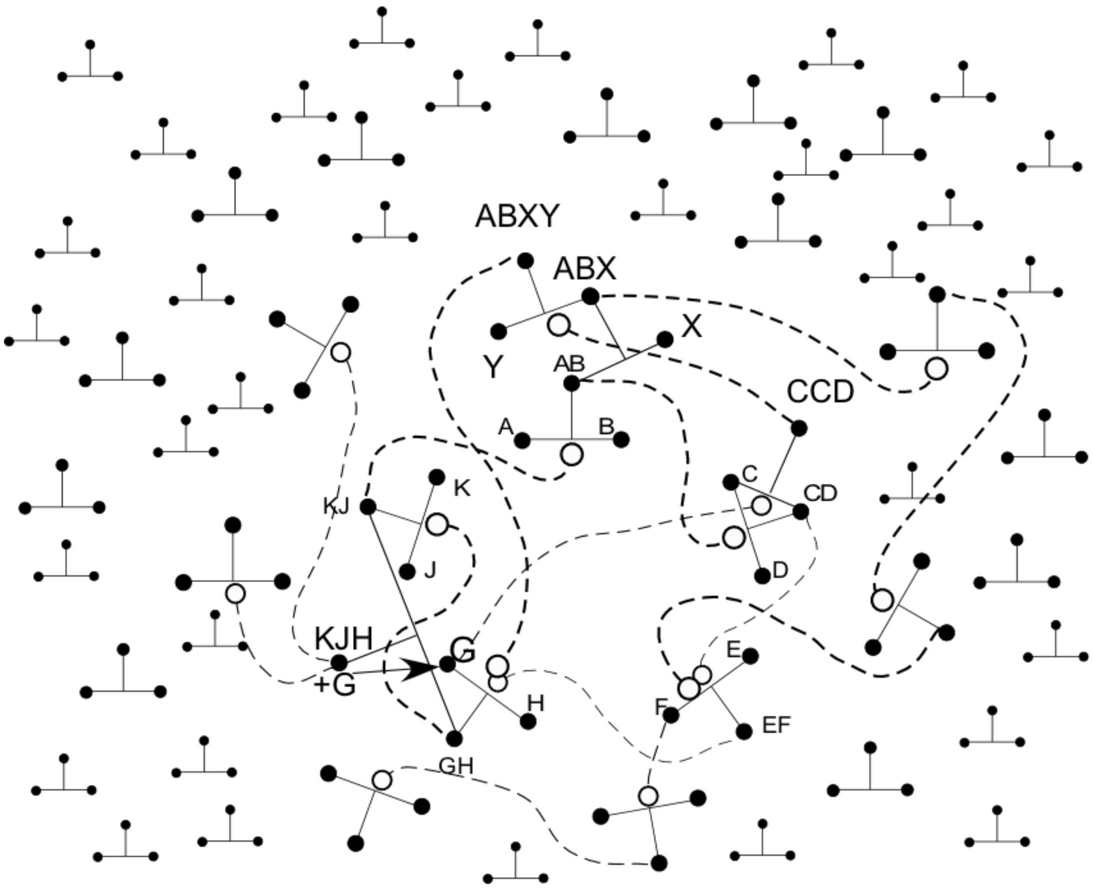

Figure II.2.2 Suggesting how quickly an autocatalytic loop can generate ever more complex structures, each of which might participate in or catalyze novel reactions. Here, the AB complex may easily combine with an X, which neither A nor B could do alone, while CD might join with another C using G as a catalyst. The resulting CCD might catalyze a reaction between the newly created ABX and a Y to create the novel structure ABXY, and so on.

Meanwhile. KJ and GH might react, with or without being catalyzed, to form KJH+G, increasing the stock of G, and by implication making G itself perform the role of catalyst (along with EF) in the production of KJH. Although G does not directly catalyze either the K+J→KJ or the KJ+GH→KJH+G reaction, nevertheless after KJH has been produced, no G has been used up, and almost no KJH could have been formed if G had not existed.

In another domain, the equivalent of "G" might be called "scaffolding" or "framework". Imagine building a Roman arch. The builders first create a wooden form in the shape of the underside of the arch, and lay on this form the stones cut for the arch. The final stone to be placed is the "keystone", and until it is placed the form cannot be removed because all the other stones would fall down separately. As soon as the keystone is settled in place, the builders remove the wooden form, indispensable during the building, useless when the arch is complete. In the analogy, "K, J, H" are stones of the arch, including the keystone, and "G" is the wooden form. Analogies can be found in a wide variety of domains, some of which will be mentioned in Volume 2.

The reason no KJH would have been formed in the absence of G is that the concentration of KH from uncatalyzed reactions between Ks and Js would have been rare, and unless KJ happened to be very stable, any that had been produced would probably have decayed back to K and J before the complex ever met a GH, which itself would not exist if G had not existed. Indeed, it would not have happened even while G existed, if the G+H→GH reaction had not been catalyzed by the easily made compound EF, which becomes even more common because of its participation in the autocatalytic loop. Like KJ, little or no GH would have existed in the soup at any given moment. Accordingly, we might call “G” a second-order catalyst for the production of KJH from the soup. The argument of improbability gets much stronger when we talk of even larger complexes such as AABXYG, which probably would never have formed in the lifetime of the Universe without autocatalytic loops and networks.

The converse is true when any one autocatalytic network has formed using products constructed by reactions among elementary units that have moderate to high concentrations in the soup. Then, the production of larger complexes becomes almost inevitable rather than almost impossible, as Figure II.2.2 suggests.

Figure II.2.2 also illustrates the possibility of autocatalytic loops developing parallel structures, which would strengthen the loop against both change in the local soup concentrations and damage to a path component. Two such parallel paths are highlighted, though at this point in the development of the network, neither is part of a complete loop, because both use as a catalyst or as a reaction element the uncatalyzed product ABX. Nevertheless, if ABX is a high-rate reaction with a relatively stable product, while neither AX nor BX have this property, its presence in the soup enhances the production rate of the products in the two side-branches, which eventually catalyze G+H→GH and E+F→EF respectively, both of which are components of the original loop.

Both branch extensions of the original loop of Figure II.2.1 enhance the production of ABX and of ABXY, but not in an ever-expanding way. We might see a burst of production of ABX, because of the increased production of AB in the loop, and the production of ABXY, which could not exist without the earlier production of ABX. If ABX exists at a sufficient concentration in the soup, then it will get used up in producing ABXY in a catalyzed reaction. Whether this matters depends on the availability of X. If the supply of X is limited, the concentration of ABX will fall despite its relative stability. The production rate of ABXY will be similarly reduced, while the main autocatalytic loop will continue its high rate of production of the various product complexes.

Such bursts of productive activity around new complexes are, I hypothesize, the “before their time” precursors that often occur an appreciable time before evolutionary or technological revolutions. As the number of complexes produced by the initial loop grows, possibly one of them might catalyze AB +X→ABX, at which point the branches begin to participate in the autocatalytic action, adding redundancy to the operations of the original loop. When that happens, what was “before its time” will have come into its time. Of course, no such catalyst may ever be produced for the AB+X→ABX reaction, and both ABX and ABXY might eventually dwindle to nothing, as ever more complexes are built that use the same elementary components. Their species will have gone extinct.

Each individual catalytic process might belong to more than one loop, so the parallel branching catalytic reactions might themselves form parallel branches that eventually create entire loops that include none of the original elementary units. Because of the redundancies inherent in the parallel structures, a breakage of one arm of a fork would not destabilize either the new or the old loop. These substructures would themselves be increasingly complex self-stabilized autocatalytic networks. Autocatalytic networks are homeostatic, at least in that their existence tends to increase the likelihood that their components exist and vice-versa, at least until they run out of “food”, their elemental components.

A Reflexive Autocatalytic Foodset (RAF) is a structural motif that has the emergent property of creativity. The environment in which a RAF does its work is what Gabora, Beckage and Steel (2021) call a “foodset”, and that I call a “soup” throughout this book. I shall continue to use “soup” despite the “F” in the acronym. This Chapter is about autocatalytic loops in general, self sustaining structures that exhibit creativity. The RAF motif offers a specific mechanism that applies to every link in a autocatalytic loop of concepts and in reaction networks that contain such autocatalytic loops.

The basic RAF structure has three components, two components of the soup that can react to form a product that becomes a new component of the soup, and a catalyst from the soup that now contains the product as a new component. Their soup consists of consciously perceived concepts, models of something about the natural and social environment of an individual living control system. We would call their “concepts” “perceptions”, but they are necessarily perceptions in imagination, so we will continue to use the word “concept” when we need to identify that they are in imagination.

Gabora, Beckage and Steel distinguish two different kinds of components of the concept soup, concepts derived from the experiences of the individual acting in the world and those described to the individual by others. We would distinguish them differently, as perceptions either fully or partially reorganized into the perceptual control hierarchy, and perceptions that consist of environmental situations not previously incorporated. Concepts introduced by others are themselves possible new components in the soup, and can be used as such, so long as they do not conflict with and inhibit the action of other components in existing RAFs.

II.1.2 Population and Autocatalysis

What, after all that speculation about the use of complexes as catalysts, is the probability that autocatalytic networks will arise in some arbitrary collection of elementary units? Is it vanishingly unlikely, almost inevitable, or somewhere between? This is a general question. Its answer depends on the number of different kinds of elementary units available to participate in reactions, which we call the “population” in this section.

If there are too few varieties available, there will be insufficient “invention” of new complexes, and the probability is near zero that an autocatalytic network will come to exist. If there are more than some critical number of different kinds of elementary units, the appearance of an autocatalytic loop and thence an autocatalytic network becomes almost inevitable (Kauffman, 1995).

We ask this key question here in connection with chemical elements, as did Kauffman, but later in this chapter the units will be cell types, and then different elements of technology, and in the next Chapter they will be the side-effects of individual or collective control of a variety of perceptions. The same underlying construct of enhancing combination rates over dissolution rates — catalysis — exists in all these situations, together with others we will not examine. The same mathematics applies to the base question, no matter what the character of the elementary unit, or of the means by which those units may form complexes from two or more elementary units.

What follows is Kauffman’s argument for the case of the chemical elements and their compounds. We will call them “elemental units” to avoid the implication that the mathematics works only for chemical compounds. Among an arbitrary array of N elemental units, assume that there is a probability P that one unit can catalyze a reaction between a particular pair to form a complex. The probability is (1-P) that a selected unit does not catalyze a specific reaction chosen at random. We have two questions. The first is how probable it is that in a “soup” of N elementary units at least one reaction is catalyzed. The second is how probable is it in a soup of N elements the product of one reaction will catalyze another and that reaction a third, and so on until a loop is formed. These questions are not trivial to solve, but we can solve the first, and put some kind of limits on the second. Kauffman reports the results of simulations that demonstrate the effects.

We next investigate the probability that precursor autocatalytic loops and clusters might occur in the soup. Since there are N^2 possible reactions among the N elementary units, and on the order of N^2 different kinds of unit in the soup (complexes and elementary units together), there are on the order of N^4 ways some unit might catalyze some reaction. The probability that the actions of at least one of the group will catalyze at least one reaction is $1-(1-P)^{N4}$. To see how rapidly this probability approaches unity, if P is 1/10,000 and there are 80 elementary units, the probability that there is at least one catalytic link is 0.5, if there are 150 units it is 0.9, and if there are 300 units over 0.9999 (Figure II.2.3). In the Figure, the example of catalysis is a side-effect of control that eases control of some unrelated perception, a control loop taking the place of a reaction among simple chemical elements.

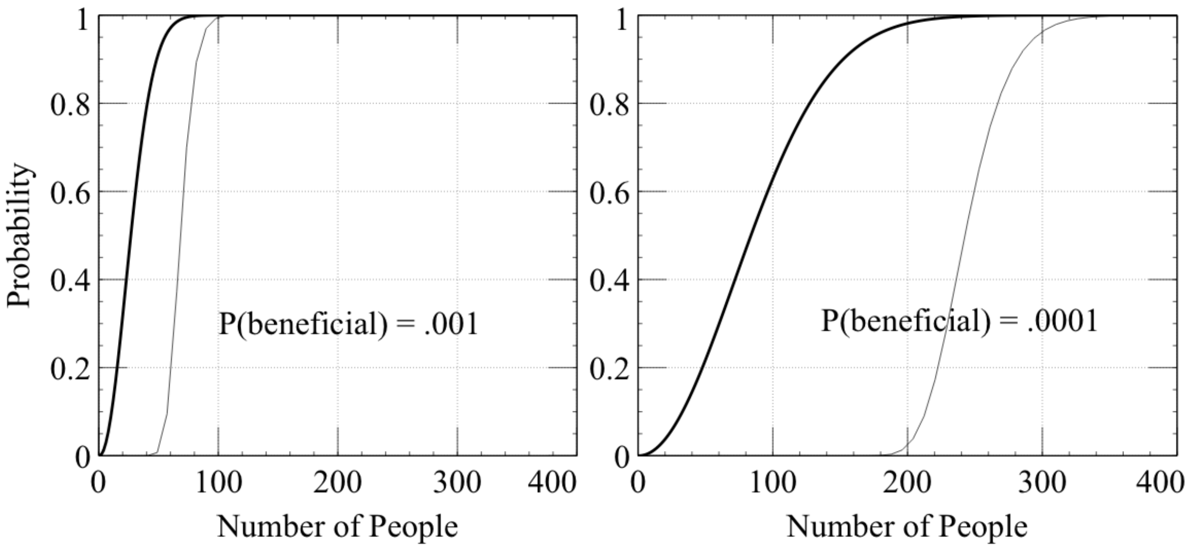

Figure II.2.3. The probability that at least one side-effect interaction will be helpful (heavy curve) and that there will be at least one loop (light curve) in a population of given size, if the probability is (left) 1/1000 (right)1/10,000 that a random specific interaction is helpful.

Figure II.2.4 shows on a log-log scale the number of units required to ensure that the probability that there is at least one catalyzed reaction exceeds between 0.5 (lowest symbols) to 0.99 (upper symbols), as a function of the base probability that a randomly chosen unit will catalyze a randomly chosen reaction.

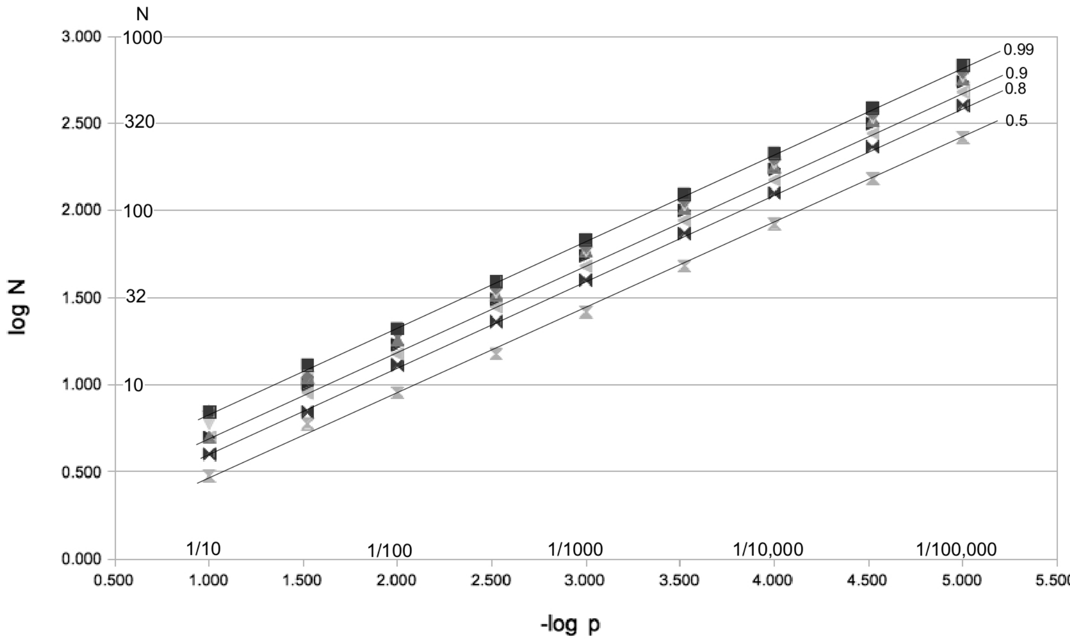

Figure II.2.4 The log of the number of control units (N) required for the probability to be 0.99 to 0.5 that at least one beneficial side-effect exists as a function of the probability p from 0.1 to 10^{-5} that a randomly chosen effect is beneficial to a randomly chosen other control unit. N varies as the square root of 1/p (the straight lines)for all probabilities of a beneficial side-effect existing.

The relationship on a log-log scale has a slope of about 0.5, meaning that N, the number of units required, is proportional to $1/\sqrt{P}$, the base probability, regardless of the value of P the probability a particular elementary unit (side-effect) will catalyze a specific reaction (will benefit control of a specific perception). With a few tens or low hundreds of different perceptions being controlled, it is almost certain that there will be at least one beneficial side-effect, even if there is only a one in ten-thousand chance that a randomly chosen interaction is beneficial.

When there is a high probability that at least one interaction is inadvertently helpful in a population, there are probably more. The probability that there are at least two, for example, is almost the square of the probability there is one (almost, because one of the N*(N-1) possible interactions has already been shown to be beneficial).

The light curve in each panel of Figure II.2.3 shows the probability that there will be at least as many catalyzed reactions as there are units, for two different probabilities that any one unit will catalyze some reaction. When there are as many catalyzed reactions as there are units in the population, it is inevitable that there will be at least one autocatalytic loop.

In the soup there are N elementary units, and therefore N^2 possible reactions of two units (which may be of the same kind) to form a complex. Some of those reactions will produce a stable complex, some will not react at all, and some will produce a complex that breaks down slowly or quickly under the influence of variations in the local environment. At any time in the soup the stable and slowly decaying complexes will persist for long enough that they themselves might participate in further reactions, creating larger complexes. With enough individual units, it is highly likely that most of the unit reactions in the soup will catalyze or be catalyzed by the products of other reactions. At this point, a large catalytic network of forms, leaving only a few isolated small autocatalytic loops functioning independently, plus a few reactions that are as yet uncatalyzed by and not catalyzing any other reaction that can occur in the soup.

No matter how the catalyst or anticatalyst functions, the reaction product decays at the same rate. Its concentration in the soup depends on the catalyzed production rate and the decay rate in exactly the same way as in a control loop the asymptotic output of a leaky integrator depends on the ratio g/r between the gain rate and the leak rate of the integrator.

When we are dealing with concentrations of chemicals involved in a reaction, what corresponds to "g" and what to "r"? The value of "g" depends on three things. If A and B are the concentrations of the reacting chemicals and X the concentration of the catalyst, "g" depends on the product of the three concentrations, the mutual reactivity of units of A and B, and a "catalytic multiplier" that is the influence of X on the reactivity of the components.

As for "r", it is a simple decay rate of the reaction product AB, in the local environment. In a high temperature environment, the AB product is hit harder by stray environmental molecules, which might break the A↔B bond, add the impinging molecule to the complex, or have no effect on AB. The strength of the bond is another measure of the reaction rate, since the greater the affinity of the two species for each other, the more reactions per second will occur for a given concentration in the soup. Similarly, the greater the strength of the bond, the less likely it is that stray environmental shocks will dissociate the complex.

This being the case, "r" is a decreasing function of bond strength. Since "g" is an increasing function of bond strength, the concentration of the product in the soup, determined by g/r, is strongly dependent on the reactivity, which is a measure of the probability that a reaction will occur when an instance of A comes near an instance of B, regardless of whether that encounter occurs in the presence of an instance of X.

If a reaction is positively catalyzed and the source units and the catalyst remain in good supply, the concentration of the product complex in the soup will stabilize at a higher level than if it were not catalyzed. The words "remain in good supply" are important, as they are the basis of Darwin's selection by competition for limited resources. Catalyzed complexes decay at the same rate as do those same complexes produced without catalysis. If their constituent elements are not available to produce more of them, the species of complex will go extinct, or persist at a very low level when the decay products loop back around to provide more food — "recycling" in everyday speech.

For our purposes, the isolated clusters and loops that are expected to form well before $P \times N^2$ reaches 0.5 are important to the story of the initial growth of N, the total number of different kinds of elementary or complex units available to participate in reactions that produce larger autocatalytic portions of the large network of possibilities. Following Kauffman, we will use the word "cluster" rather than "loop" to represent a catalytic group that may be a simple loop but that could also be a highly complex structure with many loops and branches, but is not autocatalytically connected to any other part of the larger set of possible reactions.

Why are these precursor clusters important? Because by catalysis each increases the concentration of the complexes formed in all the reactions of their cluster (which may include a loop), and therefore available to participate in further reactions. The "phase change" that occurs as $P \times N^2$ passes 0.5 completely changes the community structure. In the new structure after the phase change, almost all of the clusters are interconnected, which leaves behind only a few self-supporting isolated autocatalytic clusters or loops.

As a planetary-scale example, we may see the similarity of cities throughout the world as products of an interconnected and autocatalytic set of ideas and beliefs, while occasional isolated tribes, such as in the Amazon region, go their own way until they find that their ability to control the perceptions involved in their way of life is affected by an encroaching Global culture.

This approach, in which all pairwise reactions are equally likely and all units are equally likely to serve as catalyst, is very unlikely to be realistic. What might be closer to the truth, remembering that we are not talking only about chemical elements connecting to produce complexes, but also about a wide range of other domains, from the combination of simple cells to form multicellular complexes to the combination of vocal noises to produce complex words that form meanings when combined with environmental states, to the combination of ideas to form novel creations, and beyond into the complex mixtures that we call “politics”?

The catalytic network can be seen as two different networks of relationships, a reaction network that represents the affinity of two units, and the network of the degree to which elementary or complex units catalyze those reactions. So far, we have been calling “P” a simple probability that a unit will catalyze a randomly chosen reaction with no regard to the strength of that catalysis. When all the probabilities were equal to P, and all the reactions were equally likely to occur and their products dissociated at the same rate. Now we have the same network in which the nodes are unit kinds while reactions and catalysts form two types of link (solid and dotted lines in Figure II.2.1 and Figure II.2.2), and the links have different weights.

What do these “link weights” mean in functional terms? The reaction weight is a measure of the affinity of two neighbouring nodes in the network, which translates into the likelihood that two instances will combine to form a complex when they meet, and into the rate at which instances of the complex would dissociate. Typically, the easier they combine, the slower they disintegrate. Figure II.2.5 suggests the functional relationships involved in one production component of this ever-growing network.

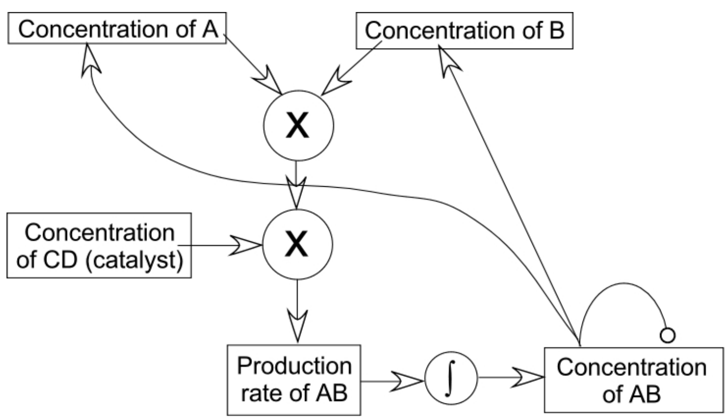

Figure II.2.5 Numerical influences on the production of a single complex in the soup. The decay rate of the AB complex is ignored. The “X” symbols mean multiplication of the input numbers to produce the output number.

The catalytic weight is the multiplier that the presence of the catalyst has on the reaction rate. The ratio between the total reaction rate in the soup and the dissociation rate for the complex determines the asymptotic concentration of the complex in the soup. But the total reaction rate in the soup depends on the concentration of the possibly reacting units in the soup, since the more A-type and B-type there are in any small region of the soup, the more opportunities there are for them to react and produce an instance of the complex. The whole network of relationships is very complicated and not amenable to a simple analytic solution.

In the larger network, A, B, and AB might all be catalysts for some reaction or other. One of them might catalyze a reaction C+D→CD, or some intermediate reaction whose product catalyzes that reaction. Suppose AB is the one that catalyzes the production of CD, as in Figure II.2.6. In the original basic soup that contained mainly the elementary units, with no catalysts[1], the concentration of AB and of CD might be very low, so AB would not be able to catalyze many C+D→CD reactions per second. As the catalytic effect of CD took effect, the concentration of AB would increase, increasing the rate of the C+D→CD reaction, in a positive feedback loop. The two multipliers must increase together at the same proportional rate.

Even such a simple two-component autocatalytic loop has several feedback loops that must be taken into account in any analysis that determines what the proportional rate might be. No matter what the result, however, their joint growth will initially be exponential until the disintegration or decay rates of the AB and CD products, or reductions in the supply of the basic "food" units (A, B, C, and D), limit their concentrations to some asymptotic steady state. The same will be true of all the productions of complexes in any autocatalytic loop that has but a single path with no optional branching.

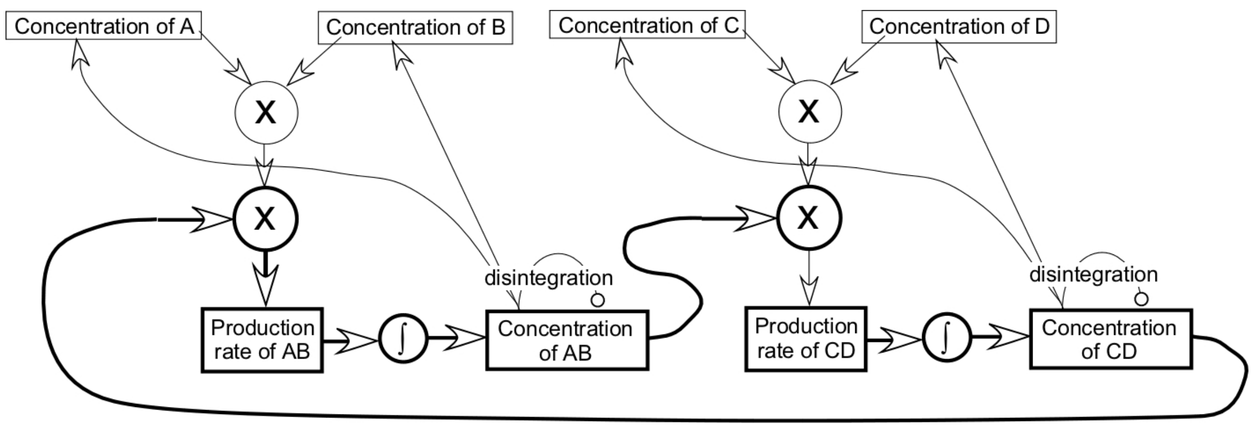

Figure II.2.6 Numerical influences in a minimal two-component autocatalytic loop.

The proportional growth rates of the products around the loop may be initially identical, but that does not necessarily mean that they stay identical, because their decay rates are all different. A product that starts out with a very low concentration because its elementary units have a low concentration may nevertheless have a low rate of disintegration, and may build to a high final concentration, while another that started out with a high concentration may not gain very much in all because it has a higher decay rate. Eventually, all of them will approach asymptotic values if they participate in no further reactions — but at least some of them will participate further, extending the network and introducing branching structures into the asymptotic loop.

---

1. As when atoms began to form after the "Big Bang" and there were effectively only hydrogen and helium atoms, which then began to form more complex atoms inside the first stars, and then to form even more complex atoms in the supernova explosions of stars. Only after an early star died as a supernova did atoms more complex than iron appear in the Universe. Until there had been enough supernova explosions, there might not have been enough variety of elements in any one region of the Universe to permit the development of the autocatalytic networks we are talking about. Life could not have emerged on planets of the early stars.

II.1.3 Anti-catalysis: Inhibitory effects

The effects of one element on the reaction between others is not limited to catalysis, if catalysis is defined only as increasing the rate of the affected reaction. We will use Kauffman's argument again in Chapter III.2, when we treat loops of control side-effects and show how they may be homeostatic. On a social scale, the inherent creativity of such loops results in building and sustaining an ever more complex infrastructure for a society. There, we will also use control of perceptions of those side-effects as a mechanism for further stabilizing the infrastructure, creating a culture in the social sense, as opposed to the "culture" of cells or molecules in our "soup". But right now we use it to argue that "what's sauce for the goose is sauce for the gander" — there is no a priori reason that catalytic processes should be statistically preferred to anti-catalytic ones. We ask what may happen when a loop incorporates an anticatalytic link, or when a product that exists in the "soup" has an anticatalytic effect on a reaction in an existing autocatalytic loop[2].

Anticatalytic effects have long been known. In a conference presentation as early as 1930, Vibrans (1931) began by referring to observations of anti-catalytic effects dating back as far as 1797, and continuing through the 19th and 20th century. The "anti-oxidants" that are so commonly mentioned in contemporary health-related advertisements and reports are good examples of anti-catalysts, if they actually work as advertised.

Early in the history of life on this Earth, the atmosphere was very different from the way it is today. Oxygen was a poison, excreted by most life-form as does vegetation today. The early life-forms had to excrete the oxygen because it disrupted important homeostatic loops and autocatalytic processes by reacting with their components.

Over time, the build-up of this poison in the atmosphere eventually led to an atmospheric concentration to which life had to accommodate, which it did by producing new kinds of autocatalytic loops that used oxidation reactions as an energy source. We presumably use some of these old oxygen-using loops or similar descendants of them even now. Nevertheless, many of the old homeostatic loops persist in our bodies. Oxygen could still damage them, leading to the requirement for anti-oxidant anti-catalysts to protect them, while the new oxygen-based reactions were incorporated in new kinds of autocatalytic loops that produced ever greater numbers and varieties of novel structures[3].

The kind of decay suppression performed by an anti-oxidant is different from an anticatalyst that participates directly in a reaction by reducing the over reactivity of the components. The anti-oxidant acts by preferentially reacting with oxidizing components, thus shielding the vulnerable AB products from strong environmental effects. The result is the same for all practical purposes, since it reduces the decay rate "r" of the AB product, as also would depressing the gain of an $AB+O \rightarrow ABO$ reaction.

Direct anti-catalysis seldom gets the attention that catalysis does, perhaps because an anti-catalytic

2. In the diagrams of Figure II.2.5 and Figure II.2.6, if the catalytic multiplier is less than unity, the network node shown as having a catalytic function is actually an anticatalyst. This suggests that in any consideration of link weights in the network, the weight assigned to a catalytic link should be the logarithm of its multiplier. All the numerical values shown as "signal values" on links, such as reaction rate, should also be logarithmic, and the multipliers shown as adders. An elementary or complex component with no effect on a reaction rate would have a link weight of zero, equivalent to a multiplier equal to unity.
3. Presumably the same kind of dramatic restructuring of life will occur as a consequence of the build-up of carbon-dioxide in today's atmosphere. But discussion of that possibility is for a different place.

influence on the productive A+B→AB reaction simply means that AB is less likely to be observed. Only if AB is a product of interest to an investigator will its low concentration be studied and a potential anti-catalyst identified.

When we consider the mutual influence of control systems in the next Chapter, however, we will find chemical anti-catalysis to be more wide-spread than catalysis, just as inhibition tends to have a wider spread in neural interactions than does excitation (Chapter I.9). The question to be addressed then will be how and when catalytic effects produce autocatalytic networks in the presence of widespread inhibition. This problem becomes even more central by Chapter IV.7, when we consider "The Madness of Crowds" (Mackay, 1841) and separately the intractable opposition between nearly homogeneous political or religious groups who each try to destroy the cohesion — the glue provided by a commonly held autocatalytic network of perceptions or beliefs — of the other.

The details are unimportant. The point is that autocatalytic loops and networks can be disrupted by factors that inhibit at least one of their component reactions, though as we will see later in this Chapter only loops that incorporate at least one anticatalytic effect can be stable (homeostatic) when there is an unlimited food supply, and only homeostatic loops can lead to the kind of perceptual control that lies at the heart of everything in this book.

This issue is not limited to the internal biology of the body, but will be seen everywhere that autocatalysis may occur. In technological systems, for example, the key reaction is the combination of disparate ideas to form a new product, which might be a physical product or a conceptual one. But that reaction might be inhibited by another existing idea, such as, for example, that one of the component ideas is unsupported by current conceptions of science — meaning that it does not participate in an existing technological autocatalytic loop — or by one's religious beliefs, which means that it must be a false idea.

Perhaps this is why such a high proportion of innovations, especially in mathematics, are made by young people. Their perceptions of "the way the world works" may be less embedded in quasi-stable networks of autocatalytic loops than are those of older people. They have fewer fixed structures, self-sustaining and augmenting loops, of ideas and are therefore more amenable to participate in reactions that involve ideas not incorporated in the existing loops. We are here returning to the development of complex perceptions first discussed in Section I.5.3. There, the loop structure was treated from an entirely statistical viewpoint. Now, though we maintain the thought that the perceptions that persist are the ones that can be stabilized by control, we also add another source of stabilization, homeostasis, which relies on the feedback loop of autocatalysis.

Returning to the issue with which this section started — that not all elementary units have an equal probability of participating in reactions or of catalyzing arbitrary reactions — we now have oxygen as an example of a reaction network node that has many links that include it as a component. To early life, this fecundity tended to break catalytic reaction networks by destroying their products before they could catalyze the next reaction in a loop, but now, those very same reactions themselves produce complexes that act as catalysts or energy sources in autocatalytic loops.

Hydrogen also forms compounds with many other elements, critically including oxygen to produce water. Water catalyzes or provides an environment for many reactions. Sodium and the other halogens are frequently seen in the company of the fluorine-chlorine set, just as male and female combinations in the social world are more frequent than male-male or female-female ones. To these examples of reaction components, we might add gold or platinum as nodes that have many catalytic links. These are just examples to show that the concept of unequal probabilities applies to both the reactions and to the catalytic part of the network.

Whether the probabilities for a particular type of network connection are the same or different, formal analysis of the network as a whole is very difficult. The actual behaviour of such complex systems is usually chaotic, meaning that small changes in initial conditions sometimes can lead to completely different futures. What we can say, with some confidence, is that a network with such a characteristic range of influences among its nodes will have much the same properties as a "small-world" network, in which the predominant pattern is for localized modules to be connected sparsely into larger modules at all scales in a quasi-fractal structure.

We saw this modularizing tendency when we discussed the structure of the control hierarchy in an individual organism. Here we see it as a more general phenomenon. In this same vein, we can expect that short loops will tend to involve more reactive elements rather than ones with fewer affinities, and that the longer loops, which are more likely to incorporate rarer and less reactive species, will be at the same time less stable and more productive of "inventions" of new species. Variety matters.

Returning to the world of A, B, and X, in which X may catalyze a reaction between A and B, Kauffman's mathematical argument applies equally to the development of multi-cellular structures, in which the As and Bs represent different types of cells. If A and B cell types accompany each other with higher probability in the presence of X, then the A+B→AB product will appear with higher probability in the regions of the soup with higher concentrations of X. In a molecular reaction A+B→AB, the reverse reaction, AB→A+B, is also possible.

How much of AB remains in the soup depends on two rates, the rate at which As and Bs meet and join and the rate at which an AB dissociates into its parts. The relationship can be described by the same equation as we use for a "leaky integrator" output function in a control loop, using now the letters to represent concentrations in units per litre, say: $g\int A(t)\times B(t)dt - r\times AB(t)$. The result is a steady state in which the ratio of A and B to AB is proportional to g/r, which we might call the "reaction gain".

If X catalyzes the reaction, it increases g without changing r, increasing the "reaction gain" and the concentration of AB. If it is anticatalyst, it reduces the concentration of AB in the soup. We will see that effect in Volume 4, when we discuss the political construct of the "Big Lie" that is being and has been so effectively used by elected would-be autocrats of many nations, from Hitler, who described the concept in Mein Kampf, to so many modern politicians all over the world. The "Big Lie" acts as an inhibitor for the productive autocatalytic loop of ideas that constitute the basis of what Eurocentric countries call "democracy", such as "free speech" and "open inquiry".

Catalysts that form part of the autocatalytic loop are products of the loop reactions, so these are likely to be of the same nature as their elementary constituents. If the participants in the reaction are molecules, the product will be a molecule. If they are ideas, the product is likely to be an idea. The catalytic action, however, need not depend on the likeness of kind around the loop, The AB complex resulting from the catalyzed A+B reaction might happen to have a configuration into which cell types P and Q could fit in a way that allowed them to enhance their survival together, even after they float off as a PQ complex. AB, having been produced with the assistance of a C, might then catalyze a possibly rare P+Q→PQ linkage, and Kauffman's argument would proceed unchanged.

This example leads us to consider how an existing autocatalytic loop or network-cluster might grow. Much of the preceding discussion has been based on the idea that the raw elements of an autocatalytic loop or network are in unlimited supply. In practice, this is never true. Use of A in producing an AB complex means that there is one less instance of A to be used in a different reaction or as a catalyst. This is not true of the use of an instance of C as a catalyst. The same instance of C can be used many times over to make more AB, but that instance of A cannot be re-used to make another AB until the original AB is recycled into its A and B constituents.

In the "soup", as Kauffman pointed out and as was mentioned above, many small "clusters" will arise before the main body of the soup gels into a coordinated autocatalytic network. Each cluster is at least an autocatalytic loop, if not a more complex autocatalytic network. Within each such independent cluster, new complexes are being formed, and old ones are serving to catalyze reactions within the local network. But nothing prevents a complex that is a catalyst in cluster A from being an anti-catalyst, a poison, for some reaction in cluster B. If this happens, and the inhibition is strong enough, cluster B will at the least lose a pathway, and at worst will have its autocatalytic loop broken, and will dissolve, allowing its more complex products to decay without the rapid replacement provided by autocatalysis.

On the other hand, if cluster B does survive the inhibition, either with a reduced rate of production of its catalyzing complexes or by using different network branches to complete its loop, there is always the possibility that one of its own productions will serve to inhibit some process in cluster A. The Kauffman argument suggests that if cluster B does not die, the probability increases over time that such an "A-inhibiting" complex will come into being. In its turn, this inhibition will reduce the inhibiting effect of cluster A on cluster B.

The two inhibitions would complete a loop joining the two clusters, not in a larger conjoined autocatalytic network, but in a positive feedback loop that stabilizes either when one cluster dies or the concentrations of the important products stabilize at some asymptotic levels. The effect is the same as was discussed in Chapter 9 on Lateral Inhibition, except that the flip-flop elements are autocatalytic clusters rather than control loops. With gentle mutual inhibition, the positive loop gain will be less than unity, and the two clusters can both survive. if the loop gain of the inhibitory processes exceeds unity, however, one cluster will dominate.

The above argument depends on the idea that one A-type cluster interacts with one B-type cluster, and both are fixed in space relative to each other. But in the soup, neither of these premises is likely to be true. The A and B clusters interpenetrate spatially, except for the effects of their mutual interactions. Even if there were only a single instance of each type, they could move apart, like conflicted controllers, perhaps by altering the precise complexes that are involved in their reactions, but perhaps also by moving physically, as we now show.

In the case of the autocatalytic clusters, this happens almost automatically. If the "soup" is more or less homogeneous initially, its homogeneity will be reflected in the existence of formally similar autocatalytic structures of, say, type A in different regions of the soup. The complex products that are part of the autocatalytic process will be created in the region of the cluster, but they will diffuse over time, since being catalysts they are not destroyed by participating in the reactions they catalyze. Their diffusion allows the reactions they catalyze to occur with increased speed everywhere in the neighbourhood of the initial type A autocatalytic cluster. It spreads not only its products, but also its processes — its reactions — through the soup in all directions. Each occurrence of a sequence of type A reactions enhances the likelihood of a type A loop becoming active nearby because of this diffusion of catalysts.

But the same is true of cluster B, some of whose products inhibit the loop reactions of cluster A. In the soup near an instance of cluster B, the cluster A reactions will be slowed as compared to their speed on the side of A away from cluster B. Cluster B tends to sweep its neighbourhood clear of instances of cluster A, and A also tends to sweep its neighbourhood clear of B. Between them, depending on the strengths of their anti-catalytic products on the "opposing" cluster, the region may be a dead zone or a zone with intermingled reactions from both clusters, varying in concentration from the core of one cluster to the core of the other. The core region of highest activity of each cluster has moved away from the core of the other, but in the interzone there may arise a new type of loop, type C, that involves some reactions from A, some from B, and some that occur in neither loop but involves products or waste from both.

We may ask "Where have we heard something like this before"? It was in Section 9.2, where we noted that lateral inhibition would move the peak sensitivities of two mutually inhibiting perceptual processes away from each other, creating contrast effects and certain illusions. The mechanisms are very different, but the principle is the same, and the movement of the peak sensitivities away from each other is for the same reason that the highest density, the core, of the mutually inhibiting autocatalytic clusters will separate geometrically in the soup. In Section 9.2, however, we did not propose the possibility that a new perceptual process might arise that partakes of properties of the original mutually inhibiting ones to make a new kind of controllable perception. Whether this actually occurs, or if it does, under what conditions, are open questions.

II.1.4 Constructive Revolutions: Levels of Autocatalytic Loops

Separable autocatalytic networks or loops form when the larger network contains two cycles that share no common elements. This does not mean that the operations of one cycle have no side-effects that influence the other, as we discussed above. In Chapter III.2, we investigate networks that consist of the side-effects of control of different perceptions. There the same level-building process builds from simple loops to loops of loops, and thence to loops of those superloops. For now, however, we restrict ourselves to the effects of increasing diversity inherent in the creativity of an autocatalytic network, and see how superloops become just as inevitable as do the lowest level autocatalytic loops.

An autocatalytic network produces increasing numbers of different kinds of product, each of which might combine with others and catalyze new or old reactions. By doing so, it provides ever more opportunity for a new and independent autocatalytic cluster to emerge. The new autocatalytic cluster itself produces an ever-increasing number of new complexes. which are likely to have a family resemblance that differs from the family resemblance among the complexes produced by other clusters. When there are enough reaction products in the soup, the mathematics of considering them without considering the original "atomic" units leads to the same conclusion. Clusters and loops are highly likely to form that incorporate those new creations, without or nearly without any of the earlier catalysts.

Another situation in which eventually there will be enough distinct autocatalytic *networks* to allow the same mathematics to be used, in which complete networks are the elementary units. A "superloop" at a new level may come into being, functionally identical to the original autocatalytic network, but in which networks of networks are the product. Indeed, the mathematics argues that this is an almost inevitable consequence, once the number of complexes produced by the individual autocatalytic clusters exceeds a critical value. Each low-level network in the new superloop produces product complexes that enhance the activity of another network in the superloop, just as each reaction in the low-level network produces a product that catalyzes the next reaction around the loop. In this way, the existence of one low-level autocatalytic network catalyzes the operation of the next around the superloop.

When there are distinct and separable autocatalytic networks, one network may combine with another to form a "product" double network, as discussed at the end of the previous Section. We will see a social example of such a double network (of side-effects of control, not of catalytic processes) in Figure II.12.8, where the "product" double network is strengthened not by synergy, but by conflict. Eventually, with enough such level transitions, we come to the concept of a level transition being manifest in the form of a revolution in which multiple varieties of more complex forms arise and become stabilized. In evolutionary history, the phenomenon is called "punctate evolution", because long periods of apparent stasis are punctuated by short periods of rapid change and of increasing complexity of novel species.

A botanical example of this kind of punctate evolution is offered by the last 450 Million years of plant reproduction methods (Leslie, Simpson & Mander, 2021). Spore-producing plants of various types appeared around 430 Million years ago (Ma). They were joined by a few pollen producers between 360 and 350 Ma, but few are found in the fossil record over the next 30 million years or so, before a rapid burst at around 320 -310 Ma. After that there was no significant change until the seed-producing plants appeared and proliferated about 120 Ma. Today, all three kinds of reproduction are used, with appreciably fewer morphological types among spore producers than among pollen producers, and among pollen producers than among seed producers.

In these pulses of rapid evolution, we can observe that spore dispersal is normally performed by the plant alone. Pollen must be delivered to the female plant by some mobile creature, usually a flying insect or one of a few bird species such as hummingbirds, while seeds are ordinarily dispersed by some animal that eats the fruit in which the seed is encased, before dropping the seed along with its soil-fertilizing dung, Each pulse involves a homeostatic loop of greater complexity than the last, but the new method of reproduction does not supplant the old. It adds possibilities.

In referring to human technologies we call such periods of dramatic change “Revolutions”. From ancient history, one clear example of such a transition was the Agrarian Revolution of some ten thousand years ago. Another might be the Industrial Revolution of the 19th century, while we seem at this writing to be have had a recent technological revolution based on electronics, and may be transitioning to another, based on robotics and artificial intelligence.

We have had transitions or revolutions based on communication, which allowed nations and empires to incorporate city-states that had previously been sufficiently isolated to act independently. The Roman Road held together the Roman Empire, as the Mongol postal courier service did for its even bigger empire and the telegraph cable network for the 19th and early 20th century British Empire. But these empires fragmented, so their qualification as indications of a level change based on communication may be invalid. Nevertheless, we may well ask ourselves whether the easily available global “instantaneous” wireless communication might herald a single global empire of the 22nd century.

Let us look at the Agrarian-to-urban Revolution in a little more detail. There is much archaeological data available for a large number of regions over the last ten or twelve thousand years, and considerable analysis relevant to our levels of autocatalytic networks was done by Turchin et al. (2018b). They constructed a massive database of archaeological data and interpretation called “Seshat” (Turchin et al., 2015; Turchin et al. 2018a) from large parts of the world, from which they sought to discover whether societies in different parts of the world evolved similarly from isolated hunter-gatherer groups to the complex technological societies of today.

The Seshat group used mathematical methods that involved no archaeological or sociological preconceptions other than those inherent in the data, and they took care to eliminate bias effects due to those. Their main finding was that exactly one single principal component of the data accounted for nearly 80% of the variance in social complexity, because the evolution of various measures occurred closely in lockstep — just as the autocatalytic level-building process by the completion of newly completed loops of pre-existing independent processes would anticipate. The nine indices they used all had very similar loadings on the main principal component, the highest by a small margin being “Polity population”, the number of people in a political unit. This is a number critical to the Kauffman analysis.

The autocatalytic level-building process suggests that the population of interacting controllers would be the driving variable for a technological revolution, the rest being its consequence. The more complex the civilization, the greater the variety of perceptions any individual would have available for possible control. The number of individuals in the population also increased over time, to multiply further the number of controllers that could participate in autocatalytic loops. As the theory would anticipate, each of the societies examined showed long periods of stability, followed by at least one jump to a new level of complexity, and more if the archaeology was sufficiently extensive.

After describing the Guttman scale in which traits may be ordered by their tendency to appear earlier or later in the record, Peregrine, a member of the Seshat group, (2017) says:

> *But there is also an interesting pattern of jumps in which several scale items appear together. Indeed the pattern appears to be fairly regular in several ways. First, there seem to be similar rapid leaps from societies having none of the traits to having agriculture and/or villages, implying that these traits appear together or in rapid succession. There appears a second common leap to a state form of government, with intervening traits appearing together or in train.*
>
> *The step-like rather than smooth accumulation of traits suggests that the unidimensional process underlying the Guttman scale is not uniform in its effects. Rather, traits often appear in clusters or groups […]. It is interesting that these clusters of traits map onto existing typologies of recurrent social formations. Cluster A is similar to what are commonly called chiefdoms—sedentary, inegalitarian but non-state societies. Cluster B encompasses states, some large and bureaucratic (Cluster D), some smaller and lacking scribes, money, and other elements of bureaucracy (Cluster C).*
>
> *The Guttman scale analyses thus suggests that recurrent social formations may be the result of a step-like or punctuated process in which a critical state is reached followed by a transformation, or, alternatively, that intermediate states are unstable. The transformed states are relatively stable and appear as recurrent social formations, although each evolves independently through the same transitive process. I suggest that what we see as recurrent social formations are not "stages" of development or societal "types", but rather are the results of an autopoietic process of convergent evolution acting across societies through time.*

Peregrine does not describe this autopoietic process, but I suggest that it is autocatalytic level building. In support of this idea, after presenting a different analysis, Peregrine notes:

> *There appears to be interdependence between scale and technology such that neither can grow without the other growing in roughly parallel fashion.*

"Scale" here is more or less the same as the population of interacting individuals, the first principal component of the analysis of Peregrine et al. (2017). All of this is exactly what would be expected from the development of successive autocatalytic "revolutions".

No matter how you identify them, the rate at which "revolutions" occur seems to be increasing rapidly, perhaps exponentially. Thousands of years intervened between the Agrarian Revolution and each successive localized complexity jump that signals a new autocatalytic revolution, and then to the putative road-based Communications Revolution, possibly another thousand or two to the Industrial Revolution or its precursor Age of Enlightenment, then a couple of hundred to the Electronics Revolution and only tens of years to the incipient Robotics-AI Revolution. As much as anything, these numbers appear to validate the idea that there is an ongoing exponential increase in the variety of autocatalytic networks at different technological and social levels of complexity.

The same can be said about biological revolutions, in which about three billion years elapsed between the first life that we can discern in the geological record to the first multi-cellular life, whereas the rest of biological history with its mass extinctions and increases in ecological complexity is contained in little more than half a billion years.

These revolutions usually have precursors, novelties sometimes characterized as "before their time". The philosophers of Ancient Greece might be considered precursors of the Enlightenment that led to a Scientific Revolution that was one of the components of the Industrial Revolution. But the Scientific

Revolution was not a descendant of those philosophers, any more than the steam turbines of the Industrial Revolution were descendants of the steam turbine toy demonstrated by Hero of Alexandria a millennium and a half earlier.

Similar examples occur in many fields, science itself not least among them. Only when the right ones have support and support each other mutually do the old ideas and concepts seem revolutionary and prescient. The concept of the autocatalytic network seems to account for these effects very neatly. The precursors merely indicate that some of the catalytic support structure that eventually became self-sustaining was beginning to exist. At any stage, it all takes a flow of energy to generate the reduction in entropy involved in creating the product structures. Eventually, the clusters that produce "before their time" product merge to become self-supporting, and a new autocatalytic network has come into being.

Even social revolutions, which may or may not have the same underlying autocatalytic processes as the technological revolutions, are often preceded by localized instances of unrest. The democratic changes of the century beginning in the 1830s in England had precursors going back at least to the time of Richard III in the 13th century, in the form of local marches to London asking for justice for the poor, each of which was put down by the authorities, except in the Civil War of the 1640s, and even then the result was a fairly quick return to a political system much like the one that had been overthrown.

Going back strictly to the autocatalytic networks, the homeostatic loop does not appear full-blown all at once. The elementary units exist and participate in reactions all along. Some of those reaction are catalyzed, producing more of the product than would be in the basic soup. Some of those products catalyze reactions among the elementary unit, with the overall effect that bursts of new productivity are seen from time to time, but die away because of insufficient infrastructural support from their environment. In the social environment, these "bursts" are the novelties that are "before their time".

If all of this seems to be almost inevitable, since it seems to apply at levels all the way from microbes to nation states, a natural question is why we do not seem to see similar revolutions in the development of large-scale technology by non-hominid species. We do occasionally see invention by crows, apes, and a scattering of others, and we do see hunting teamwork and social play in many species from schooling fish to hunting wolf-packs and sea-mammals. One day, I watched baby seals that seemed to be playing a mixture of hide-and-seek and tag around a rocky pool. But these are isolated phenomena, which show no sign of being precursors to species-wide or even local group-wide autocatalytic breakout to a new level. Why do they not?

I suggest, with some trepidation, that the answer may be the different complexity of our language as compared to theirs, coupled with the stigmergic property of writing. With the possible exception of dolphins and whales, we seem to be the only species that can describe an indefinite range of patterns in such a way that another individual can imagine or reproduce them. Since an autocatalytic loop requires a variety of catalysts to catalyze a variety of reactions that put "this" together with "that' to make an entirely new "thingamajig" and tell someone else how to do it again, language may be a medium of catalysis, much as complicated molecular shapes may be catalytic at some biological levels.

Insufficient variety and insufficient scale of interaction may combine to deprive other species of advanced technology. In another context, we might call it "inadequate education" that does not allow a sufficient variety of concepts and perceptions to interact in the brain of a growing child of whatever species, including our own. The stigmergic property of writing that allows the effect of an idea to be rediscovered years later, like Mendel's discovery of discrete "genetic" inheritance, greatly enhances the likelihood that the idea might form part of a much later creative autocatalytic loop of innovation.

This seems an appropriate point to stop and ask ourselves whether, if inadequate variety in their education is a sufficient reason for the failure of non-humanoid species to have a technological breakout along the lines of the first cells that banded together to create multicellular organisms, then what might be the effect of inadequate variety in education on a human society? I leave that question hanging for you to ponder, until Volume 4.

II.1.5 Autocatalysis to Homeostasis

In several places I have stated or implied that some autocatalytic loops are homeostatic — its signal values are self-stabilizing, like the perceptual signal value in a control loop if the reference value does not change — without showing how this must be the case. Now is the time to remedy that omission. Figure II.2.7 shows a five-element loop with simple variables, together with a visually similar loop in which reactions of two input components are catalyzed by the product of another reaction. Visual analogies are often dangerous, and should be justified. In both loop diagrams, the "a" and "b" labels represent "signal" values passed into the loop while the "p" or "c" values represent signals passed around the loop. They are, of course, all functions of time. The black disks represent nodes at which the inputs from outside the loop are processed together with the "signal" produced by the previous node in the loop to create a "signal" sent to the next node around the loop.

In the kind of two-node loop we call a "control loop", the two inputs from outside are called resepctively "reference" and "disturbance", while the signals passed around the loop are called "perception" and "output". The nodes are not simply the CEV and the comparator, but contain also the processing stages that contain the input and output side processing in both the hierarchy and the environmental feedback path. The only thing that distinguishes a control loop from a generic two-node homeostatic loop is that the multiplier gain is much greater from the reference input node to the disturbance input node than on the other leg of the loop.

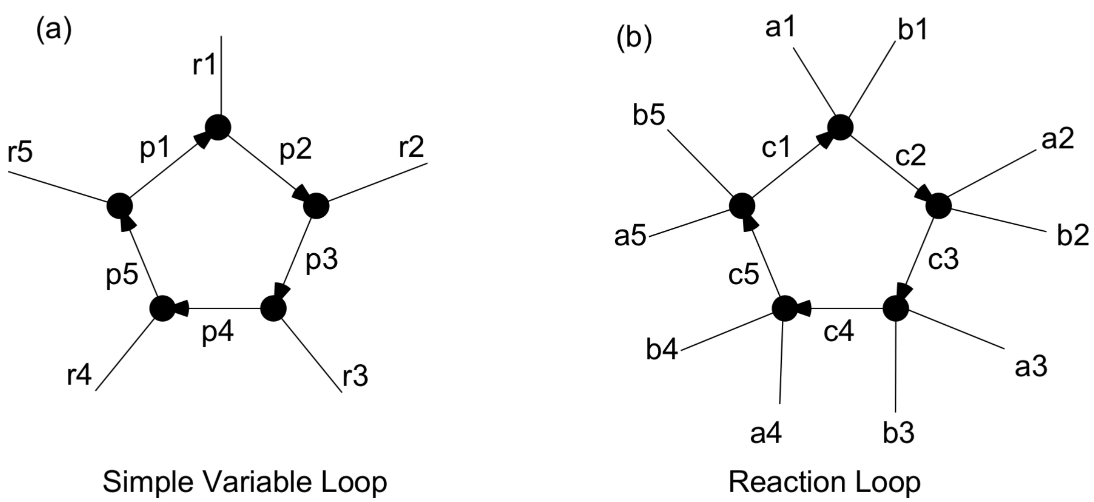

Figure II.2.7 A visual analogy between a loop of simple variables (left) and a reaction-based loop (right). In the simple loop, "r" is analogous to a reference in a control loop, whereas "p" with the same number is analogous to the corresponding perceptual value. In the reaction loop "an" and "bn" are the elementary units in a reaction, while "cn" is the catalyst for the number "n" reaction. "cn" is the reaction product of the preceding reaction in the loop, as for example c3 is the product of reaction 2 and catalyses the a3-b3 reaction.

Figure II.2.7 shows one visual difference, apart from the change of labelling, between the simple variable loop and the reaction loop, which I will now call the autocatalytic loop because of its essential use of catalysts. The nodes of the simple variable loop have but a single input from outside the loop, whereas in the autocatalytic loop each node has two external inputs. Each of the two inputs represents the concentration of one of the elementary units that react together to form the product that catalyzes the next reaction around the loop. This apparent difference in the diagrams is not functionally significant, since the value of any "r" in the simple loop could be the the result of combining several different variables.

Let's look a bit more closely at the simple variable loop first. The values labelled "p" seen from one viewpoint take on the role of the perception in a control loop, a perception that has a reference value provided by the following "r". From another viewpoint "p" is the output of the previous process. It takes on the role of the output value of a control loop, opposing the effect of the following "r" input. In a control loop, perceptions do oppose their reference values, just as outputs oppose the incoming disturbances. The difference between them is only that the gain between the comparator node and the CEV node is appreciably greater than the gain between the CEV and the comparator. Our homeostatic loop diagram indicates no gain anywhere, since all the processing, which includes any gain, is hidden in the node. Figure II.2.8 shows the gain as an asymptotic equivalent multiplier separated out as a component of the node that follows the combining property of the node.

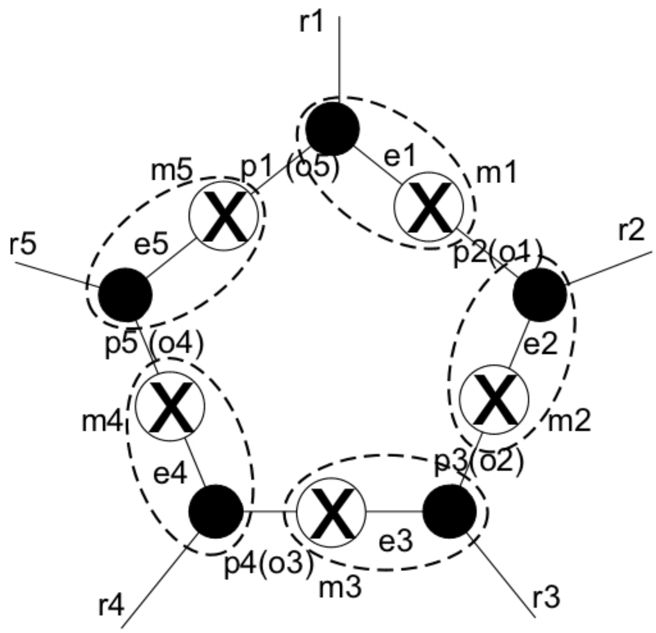

Figure II.2.8 The homeostatic loop of Figure II.2.7, showing each individual node as two processes, the first being the combining of the two inputs to produce an "error" value "e", while the second , the "X"in each stage, represents any other processing performed on e to produce the output "o" from the node, which is also the "p" input to the next node around the loop. Asymptotically, "X" is a simple multiplier, though dynamically it can represent any kind of processing.

You may have noticed a significant omission in this discussion, the sign of the overall loop gain. If all the nodal outputs, the "p" variables in the simple variable loop or the "c" variables in the autocatalytic loop, vary in the same direction when an external variable changes magnitude, the loop gain is positive. If it is greater than unity, any change in the values of the variables will be exaggerated and the variables will all grow until the loop hits some limit such as the range of variation of which some node is capable. On the other hand, if an odd number of the "p" or "c" variables reduce their magnitude when the others increase theirs, the loop gain is negative, and the loop can be stable — homeostatic. In a control loop the error value is usually taken as reference minus perception, with the result that the error becomes more negative or less positive when the perceptual value increases while the reference value stays constant.

Concentrating only on the asymptotic values for now, we can represent all the processing that might be performed by the operations of the loop between one node and the next by a simple multiplier "m". Dynamically, the values may stabilize smoothly, they might go into a repetitive oscillatory patterns, they might be chaotic, or they could explode toward infinity insofar as the physics and chemistry allow, depending on just what processing is represented by the different "X" processors.

Since we are talking only about a homeostatic loop, we know by definition that an explosion toward infinity will not actually happen, because an odd number of the nodes acts asymptotically as a negative multiplier. A broad definition of "homeostasis" as the maintenance of statistically stable characteristics over time and in the presence of disturbing influences allows for the other possibilities. Of them, smooth stabilization is the most likely, repetitive oscillations might occur as time-keepers such as the diurnal clock that alters the concentrations of many molecules over the day-night cycle but are not of particular interest here, and chaotic oscillations are unlikely to be of much biological value when they do occur. Accordingly, for our present purposes we can use the simple multiplier equivalent ("X" in Figure II.2.8 interpreted as multiplication) and asymptotic values to represent the processes and variables in homeostatic loops.

How does the loop gain play out for the autocatalytic loop of Figure II.2.7? Before this Section, we have been assuming that an autocatalytic loop must have the explosive exponentially increasing production of product that we just said will not actually happen in a homeostatic loop. But now I have claimed that even an autocatalytic loop can have negative loop gain. Can it? Can an autocatalytic loop ever be a homeostatic loop? Let us look a bit more closely at the origin of an autocatalytic loop, network, or cluster.

As Figure II.2.6 illustrates, the output of each reaction is a complex product that happens to act as a catalyst for the next reaction in the loop. The catalyst acts as a multiplier on the rate of the following reaction. That rate is integrated to produce the concentration of the next catalyst around the loop, so the integrated result of the catalytic multiplier on the reaction rate provides the next multiplier.

The mere existence of a catalyst "C" does not determine the multiplier. The multiplier determines how much more likely or more quickly single units of the two basic entities "A" and "B" will react to form an "AB" complex if they meet in the presence of a unit of the catalyst as compared to when they meet without a catalyst. The rate of an A+B→AB reaction in the soup is an aggregate of individual reactions, some of which occur when a unit of C is present, and some when they are alone. The effect of C on the rate depends on the concentration of C. The more there is, the more likely it is that a meeting of a unit of A and a unit of B will be catalyzed. If C is an anticatalyst, the multiplier is less than unity, and the more of it there is, the less likely it is that an A and a B will react when they meet.

The rate of a reaction between two components depends on the affinity of the components to each other and on the effectiveness of the particular catalyst, but in a reaction between specific components aided (or inhibited) by a particular catalyst, these are invariant. What varies are the concentrations of the three participants in the soup. Without the catalyst, the reaction rate of A+B→AB is $k_{affinity}$×(concentration of A)×(concentration of B). With a catalyst C this is multiplied by $k_{catalytic}$×(concentration of C). Using lower-case letters to represent the concentrations and combining the two "k" multipliers into one fixed multiplier k, we find that the rate of A+B→AB catalyzed by C is k×a×b×c.

For the product k×a×b×c to be negative in any particular interaction, one or three of the variables must be negative. Neither concentrations nor reaction rates can be negative, though their time derivatives may be. The affinity between the A and B components cannot be negative, because we are talking about a reaction that produces a non-negative amount of their AB product in the absence of a catalyst. The only remaining possibility is that $k_{catalytic}$ must be negative. An odd number of the products in the loop must be anticatalysts for the next reaction around the loop.

How can this be? Not long ago, we argued that one way to kill an autocatalytic loop was for another autocatalytic loop to produce a product that was an anticatalyst for a reaction in the vulnerable loop, acting as a poison by breaking the loop. Now we are saying that for an autocatalytic loop to be stable, meaning that its rate of producing its reaction products approaches an asymptote, the loop itself must generate at least one product that is anticatalytic for another reaction in the loop. Is this not a contradiction?

No, it is not a contradiction. In the case of independent "antagonistic" autocatalytic loops, the anticatalyst is produced outside the vulnerable loop, and its concentration is independent of what happens to the vulnerable loop. As the vulnerable loop reduces its output all round under the influence of the external anticatalyst, the effect of the anticatalyst continues at the same strength.

This is not true when the anticatalyst is produced by a reaction within the loop, because as the reaction rates around the loop decrease, so does the production rate and the concentration of the anticatalyst. Overall, the catalytic multipliers in the rest of the loop may easily be large enough to overcome the anticatalytic effects of the maverick, but they can never change the sign of the loop gain. If the other catalytic multipliers are not large enough, the loop may die. If they are, the loop is homeostatic with some stable production rate of the products, rather than being in a state of continual exponential growth.

There is an old saying: "one man's meat is another man's poison." Here we have the concept a product that is a catalyst within one loop being a poison for a different loop. Moreover, we have the possibility that the same poison in low concentrations may be essential to the loop. How many such products do we know in everyday life? Without salt we die. With too much salt, we die. For salt, we can substitute many of what are called "essential minerals" or "vitamins". If the concentration of these entities in the body is insufficient or excessive, we die. If we consume too much, we die. How much of the control hierarchy of an organism is evolved or reorganized so that we control perceptions in ways that eventually tend toward getting the right concentrations of these anticatalyst "essential poisons". All of it? Quite possibly so.

The question then becomes of whether it is likely that loops containing an anticatalytic stage are likely to come into being in the ever-expanding complexity of autocatalytic products and loops or clusters. Here, Kauffman's probability argument once again applies. If there are enough different kinds of elementary or complex units that could participate in reactions either as components of new kinds of product or as catalysts, any one unit is as likely to be a catalyst as an anticatalyst for any specific reaction. This is as true for the elements of the initial soup as at any later time. Hence, in any of the loops that form, each stage is as likely to catalyze as to anticatalyze the next stage.

Loops in which the complex product of each reaction anticatalyzes the next all the way around the loop are as likely to form as are ones in which all the links are catalytic. The result, however, is spectacularly different. The former results in fewer products of the participating reactions than would be the case in the absence of catalytic effects, with the consequence that their products are unlikely ever to become parts of more elaborate complexes. In effect, such a loop is purely conceptual, because its own existence decreases the quantities of anticatalysts of which it is formed. Its existence is a mirage.

In dramatic contrast, a loop that has only catalytic connections produces an exponentially increasing concentration of its products, and thus a proliferation of an ever increasing variety of increasingly complex products. These are the autocatalytic loops we would see if they exist. They must be few, relative to the homeostatic loops that incorporate both catalytic and anticatalytic processes. If they were not, their reactions would produce vast quantities of product that would swamp the world, if their food supplies (the reacting components) held out. A pandemic virus may be an example of such a process. Autocatalytic structures without internal inhibition tend to deplete the resources on which they depend, and then they die or are drastically reduced in representation in the "soup".

How probable is it that a loop is autocatalytic, and how probable is it that a new autocatalytic loop will be homeostatic? Assuming that catalytic effects and anticatalytic effects are equally likely when any one elementary or complex unit is paired with any particular reaction, we can say that about half the links in a single-path loop will be anticatalytic, with a binomial distribution of actual numbers. Also, in half of all loops, the catalytic multipliers will be large enough to make the loop gain exceed unity numerically, and in half of those there will be an odd number of anticatalytic states, making the loop gain negative and the loop homeostatic. A quarter of all incipient autocatalytic loops will be homeostatic, but these are the loops that provide stability on which other reactions and loops may depend.

We conclude, then, that Kauffman's argument applies regardless of whether anticatalytic effects exist in the soup. The mathematics assures that if there is enough diversity for a first autocatalytic loop to form, other loops will follow. Of these, half will have more anticatalysis than the catalytic effects can overcome. These loops will only be found by careful analysis, because their products will be seen only rarely.

The other half will produce appreciable concentrations of novel products and therefore of opportunities for those novel products to join into new loops that could not have existed without the older loops. Almost all these loops will be homeostatic, but their rates of production of product, and their corresponding consumption of "food" as source material and for energy, will vary greatly. Very occasional loops will occur that have no inhibitory anticatalytic stages. These continue proliferating product until they run out of food, and then they die or at least experience a population collapse.

II.1.6 Homeostatic Loop Development

A homeostatic loop as so far described does not consist of individual instances of the elementary units reacting and then reacting again. Each such instance corresponds to the firing of a single nerve when we consider the workings of the neural networks in the brain. Just as the feedback loops of the neural currents in the neural bundles discussed in Section 1.3 can be seen only in many firings of many neurons, so the homeostatic loop can be seen only in multiple instances of elementary components combining to form a complex.

The homeostatic loop is a functional loop, not a loop defined by fixed connections. The control loop to which the homeostatic loop is mapped in Figure II.3.2 is, on the other hand, usually treated as a set of processes that communicate with each other by way of fixed connections — wires — even when it is read as a functional diagram of what happens with multiple firings of many nerves. The concentrations of chemicals in the soup are analogous to Powers's "neural current" in the "wires" of the control loop, and the individual occurrences of a reaction to the firing of one of the neurons.

A single-loop circuit connected by physical wires can be broken by cutting any single wire or by shutting off the energy supply to one of the processes. A homeostatic loop in which the changes are not propagated along a wire but consist of concentrations of elementary units in a "soup" cannot so easily be damaged in that kind of way. The links are highly redundant, as are the nerve firings within one of Powers's "neural bundles". If any of the many instances of AB catalyzes any of the many possible instances of $C+D \rightarrow CD$, and around the loop any XY similarly catalyzes $A+B \rightarrow AB$, you cannot stop the process by removing an instance of any of the components or intermediate products. The loop is less like a single wire ring than like a braided chain. All that is needed for the loop to continue operating is a continuing supply of A, B, C, D,...X, and Y. If the supply of one of them is totally cut off, then the effect is the same as cutting a wire in a single-circuit kind of loop, but you have to reduce the concentration rather dramatically to stop the loop from acting.

When a loop has been operating for a while, the new products built from its AB, CD …, XY complexes together with the original elementary units will have begun to react to form yet more complexes, some of which will catalyze other reactions, not only the new ones, but old reactions, too. If now ABPQ turns out to be a catalyst for A+B→AB, then if the supply of one of the elementary units R, S, …, X, Y was cut off, AB would continue to be made, perhaps not as prolifically as before, but sufficiently that the loop would continue to produce all the same old products except for those that contained the missing elementary unit, say "T". Products containing T would decay, and their decay would restore some supply of T, allowing the old loop to be reconstituted, though only at a rate determined by rate of recycling of T.

Figure II.2.9 shows a similar situation using the old loop of Figure II.2.2 (slightly amended), in which the critical component "E" has been depleted, with the result that EF is no longer being produced, and only the existing supply of it can continue to catalyze the G+H→GH reaction. If that had happened to the initial loop of Figure II.2.1b, all the catalytic processes around the loop would have effectively stopped very soon. However, the loop has been in action long enough to produce several complex products, and now ABXY is available to catalyze the G+H→GH reaction. ABXY itself relies on a new loop, one that uses none of the original catalysts.

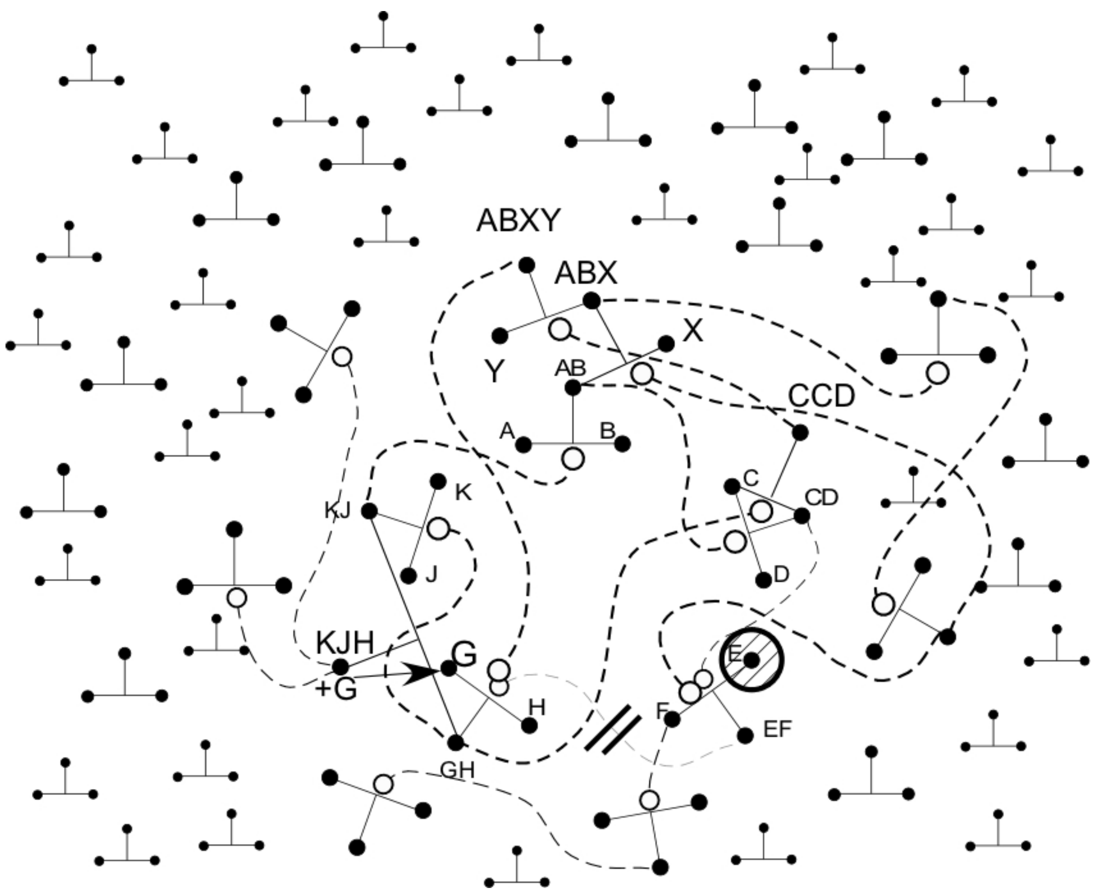

Figure II.2.9 A partially self-repairing autocatalytic process. An original catalytic link (EF catalyzing the G+H→GH reaction) has been stopped by depletion of the supply of a critical component, E, but will be restored when some of that component is recycled by the degradation of products containing it.

In this new loop, ABX catalyzes one unidentified reaction, whose product catalyzes another, the product of which catalyzes AB+X→ABX, and ABX is an essential component of ABXY, not a catalyst in the loop that produces AB. The new loop requires a supply of AB, a product of the old loop, while the old loop keeps most of its reactions functioning so long as the new loop produces enough CCD as a catalyst for the ABX+Y→ABXY reaction. The loops run separately, but use each other's product. Another such separate loop depends on the dissociation of a reaction product.

Without EF, a new loop supplies a functionally equivalent catalyst (ABXY) that keeps the loop working. Another elemental component "G" is continuously replenished by the sequence of reactions G +H→GH, K+J→KJ, KJ+GH→KJH+G. Elemental "catalysts" in this kind of loop will be depleted only if the loop's reactions are too prolific in their supply of complex products, leaving not enough supply in the "soup" to provide for more production in one of the links. The catalyst "G" is exempt from this restriction because of the new invention of the reaction sequence that replenishes the supply of "G".

The main loop, the ABXY-producing loop and the G-reproducing loop are quite different in their processes and products, but they do not run totally independently; once each reaction in the catalyst-reproduction loop has been formed into the complete loop it would keep running so long as any AB (or H) remained unused, and the same applies to the old loop so long as any CCD remained. But each would eventually die without the other's continued operation. The old loop has, in function, repaired itself, not directly, but by creating a separate process that resupplies a functional equivalent for the component that was lost.

In the evolution of biological species, we see this quite often. For example, humans have lost the ability of our ancestors to make Vitamin C, but we can live quite well if we substitute for that ability eating fruits and vegetables that contain Vitamin C. Without the right fruits or vegetables, we do not live so well, as the sailors who died of scurvy during the long voyages of exploration might have testified during their fatal illness. Vitamin C is an essential part of at least one important homeostatic loop, but the replacement "short circuit" goes through the perceptual control hierarchy.

A further point may have been missed in the preceding, which is that proliferating loops tend to get shorter over time, as new products create catalytic "short-circuits". Our ancestral (Figure II.2.2) loop's evolutionary followers short-circuited the missing E+F→EF reaction by using a functionally equivalent catalyst supplied by a newly completed loop (Figure II.2.9). The intervening set of processes that produce the replacement catalyst may be quite long when seen from outside, but that whole other loop is but a single step as seen from within the broken loop. A product happens to act as a catalyst that results in a process elsewhere in the loop being catalyzed.

The construction of a short-circuit in a loop does not mean that the earlier path ceases to function. The Vitamin C example illustrates that parallel paths allow for the loop to continue to function if a link is broken, but it does not imply that the link will be broken. A loop with many short-circuit paths simply has many ways to survive damage to one path or another. We could call such a loop a "braided" loop or simply a network, depending on which aspect of its functioning we want to emphasize. When, in Volume 4, we come to talk of ideas rather than chemicals as the reacting components, we will have a reason to anticipate the resilience of an *idée fixe* opposed to countervailing facts, and the autocatalytic escalation characteristic of "*The Madness of Crowds*" (Mackay, 1841).

The same short-circuiting process is quite clear in the development of new ideas from combinations of old when, say, producing new technology or mathematical proofs. The first prototypes often are produced by a long, complicated procedure, but shorter and easier paths are often found later, in the form of new intermediate stages or mathematical theorems that can take the place of earlier independently derived complicated analyses. It is easier to write “XYZ follows from ABC (Ramanathurian, 1927)” than to go into a five-page re-derivation of Ramanathurian’s 1927 that you can use in the new theorem that you are explaining. It is easier to buy a flibbertewidget that forms part of what you want to build than it is to make one from scratch, no matter how simple or complex it is for someone else to make flibbertewidgets from scratch.

Chapter II.2. Control: an Example of Homeostasis

This Chapter works from autocatalysis and homeostatic loops to control loops, in an effort to determine whether the evolutionary development of control systems was as inevitable as was the onset of autocatalysis. First we see how homeostasis maps onto control, and then how homeostatic systems might evolve into the simpler but more powerful control loops of PCT.

II.2.1 Homeostasis and control

Ignoring the wild autocatalytic loops that run amok and may fatally deplete their energy and material supplies, we continue to look at homeostatic loops and their relationship to control loops. Control loops are a special subset of the simplest possible form of homeostatic loop, the two-node "there and back" loop. What is true in general about homeostatic loops must be true also of control loops, but the reverse is not necessarily true. We will concentrate first on autocatalytic homeostatic loops and look at some conditions in which they might be stable.

To be stable means that if something changes at one of the inputs, other variable values around the loop change to compensate, keeping the value in the loop immediately preceding the changes input more or less constant. It does not mean that when an input changes, the loop continues to retain all its old variable values. In a control loop, if the disturbance value changes, so does the output value, though if the control is good, the perceptual value will change very little. If the reference value changes, so will both the perceptual and the output value. Likewise in a larger homeostatic loop.

As an example, consider the two-reaction loop of Figure II.2.6 as a control loop. To do this, we make the AB product rather stable, and quick to react. AB gets created rapidly and is destroyed slowly, so its concentration builds to substantial levels and declines slowly if the supply of A or B is cut off. CD, on the other hand, is created with some difficulty, though more rapidly in the catalytic presence of AB, and decays fast. The concentration of CD responds quickly to changes in the supply of its components. Finally, we make CD an anticatalyst for AB.

A control loop has two input variables, a reference value and a disturbance value. The autocatalytic loop has four, but they can be treated as just two, the products of the concentrations of A and B, and of C and D respectively. We can label the former R (reference) and the latter S (disturbance, a.k.a. "sensed change"), leading to the diagram of Figure II.3.1. The anticatalytic action of CD on A+B→AB is shown by the grey shading of the long connection curve.

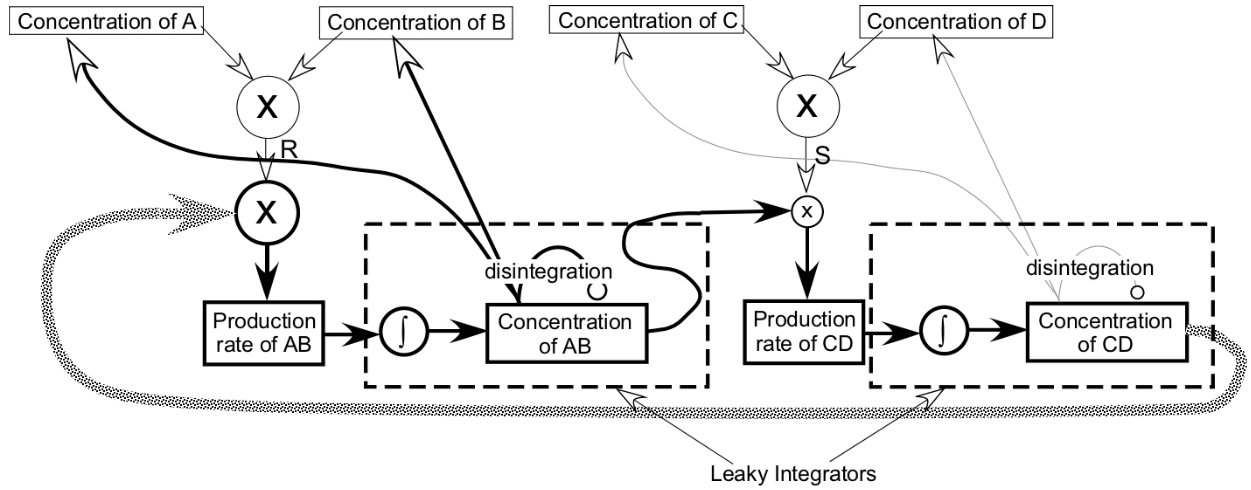

Figure II.3.1 The autocatalytic loop of Figure II.2.6 as a control loop. CD is now an anticatalyst for the A+B→AB reaction. Dashed rectangles delimit a leaky integrator process.

Suppose the system has reached some stable values because the supply of the components has been steady for a long while. Now one of them changes. Let's say that there is now more of A, meaning that it is easier for an instance of B to find a partner. The production rate of AB increases, so the concentration of AB slowly rises. AB catalyzes the C+D→CD reaction, so the concentration of CD also rises. But CD is an inhibitor for the A+B→AB reaction, reducing the rate of production of AB.

We consider the production rate of AB as analogous to the error value in the standard control loop, and the production rate of CD as analogous to the sensory input. Both are shown as being processed through a leaky integrator, but a control loop is not symmetric, and neither are these leaky integrators. The asymptotic gain of a leaky integrator is the ratio of the integration gain rate to the leak rate. Here, those values are respectively the production rate and the decay or disintegration rate of the complex, which we can set as parameters in the analysis.

Figure II.3.2 shows Figure II.3.1 overlain by a standard control loop that indicates the components of the control loop in terms of the catalytic processes and products.

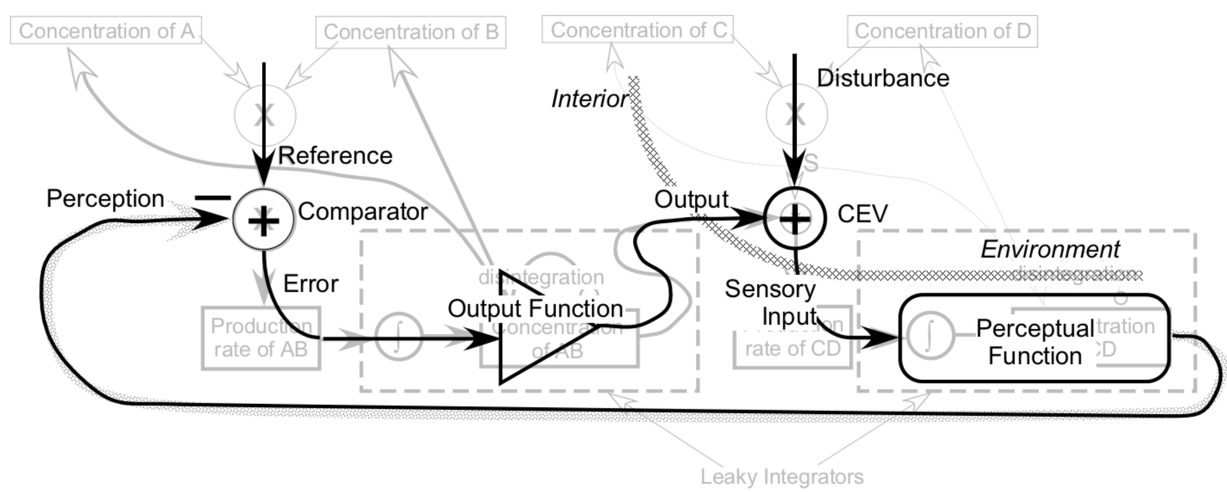

Figure II.3.2 The autocatalytic homeostatic loop of Figure II.3.1 as a control loop. The processes and variables are those usually represented in most basic perceptual control diagrams.

The overlay suggests two apparent anomalies. Firstly and least importantly, the overlay separates the interior from the environment. By the environment, what is usually understood is that part of the Universe outside the skin of the organism. However, to an elementary control unit that consists of perceptual input function, reference input function (in this case a simple connector), comparator, and output function, the environment is everything else. Its output goes blindly into its environment, and what returns is sensory input that its perceptual function converts into a perceptual value — a perception.

The second anomaly is less obvious and more significant, but in the end it turns out not to be an anomaly at all. It is that the comparator of a control loop is usually shown as an adder with a negative perceptual input and a positive reference input, so that it produces an error value that is reference minus perception. The autocatalytic diagram shows a catalytic multiplier at that point. At the CEV, too, a multiplier is replaced by an adder. Nevertheless, this autocatalytic loop does function as a "standard" PCT control loop, as I now describe.

Typically, if the components have a high affinity for each other, their reaction rate will be high and the decay rate of the product will be low. The integrated amount of the product, the "output" of the reaction process considered as a control loop, will be high, as in the output stage of a "standard" control loop such as those with which we introduced the concept of control at the start of this book. The output stage "gain" will be high. In contrast, if the C+D→CD reaction is slow and CD decays rapidly, the concentration of CD will vary rapidly when either of the component concentrations change, and so will its inhibitory effect on the production of AB.

The concentration of A and B provides a reference value for the concentration of the CD product, while variation of the availability of C or D acts as a varying disturbance to the concentration of CD. Those changes are opposed by changes in the concentration of the AB catalyst, stabilizing the "perception", the concentration of CD in the soup.

In a tracking study that asks a person to use a mouse to make a cursor on screen follow the movements of a target, the perception of the target-cursor distance can follow erratic changes much more rapidly than the person can move the mouse to compensate for them, so the extent that the time it takes to create a perception is usually ignored in any later simulation of the human behaviour. This difference is the reason we map the high-affinity AB reaction to the output side of the mapping and the rapidly decomposing CD product to the perceptual side. There is, after all, no anomaly.

II.2.2 What is the Environment of Control?

Figure II.3.2 should raise another issue in the reader's mind. The separation between interior and exterior is illustrated in Figure II.3.2, but it was there dismissed as unimportant. But is it really so unimportant?

Both in the earlier part of the book and in what follows, we consider control through the environment as central to our theme. Is the importance of the environment outside our skin important only because of my own interest as an author, or does it have an importance based on something else? Many control loops and longer homeostatic loops probably do not involve any variable outside the skin, but deal with the levels of internal variables such as hormone concentrations. Surely they also must be important? If they did not exist or continue to function properly, would we continue to survive? Probably not.

The dilemma posed here is false. The barrier shown in the Figure between and "Interior" and an "Environment" need not be the skin at all. What it really signifies is a distinction between the variable that is made to match a reference value (the perception) and the real reality variable that is influenced by the actions dependent on the reference-perception deviation (corresponding to the CEV) whose value in turn leads to the value of the perception. As we saw in Section 10.1, the main flow of energy that permits export of entropy from the loop occurs through the path from reference to CEV, with very little in the form of increased entropy returning from the CEV to the perception.

In the AB-CD loop of Figure II.3.1, this reduction of entropy by using energy is implemented by the production of the long-lasting AB product, not in creating the CD product, since CD decays rapidly, returning the energy and regaining the lost entropy.

From the physics viewpoint, the whole function of control is to limit the import of entropy from the environment into the system that contains the control loop. There we have another word "contains". The use of that word implies the existence of some boundary between the "system" and the rest of the Universe. Perhaps we had better investigate it a little further.

In the homeostatic loop onto which the control loop is mapped in Figure II.3.2, the entropy export occurs when energy is used to create a reaction product AB from two component units A and B. In many biological situations processes are powered by the destruction of complex components, increasing the chemical entropy of the system and delivering power where it is used for other purposes. The mitochondria within eukaryotic cells do this. But we ignore it here, because we are interested in a different situation. We are exploring the mapping of homeostatic loops onto control loops and the development of the latter from the former.

Figure II.3.3 shows a single node from either Figure II.2.7b or Figure II.2.8, the Figures that illustrated five-node loops. Each node receives "food" from the soup and returns complex product and waste to the soup. The concentrations of food, product, and waste are the variables of interest so far as the operations of an autocatalytic, homeostatic, or control loop are concerned.

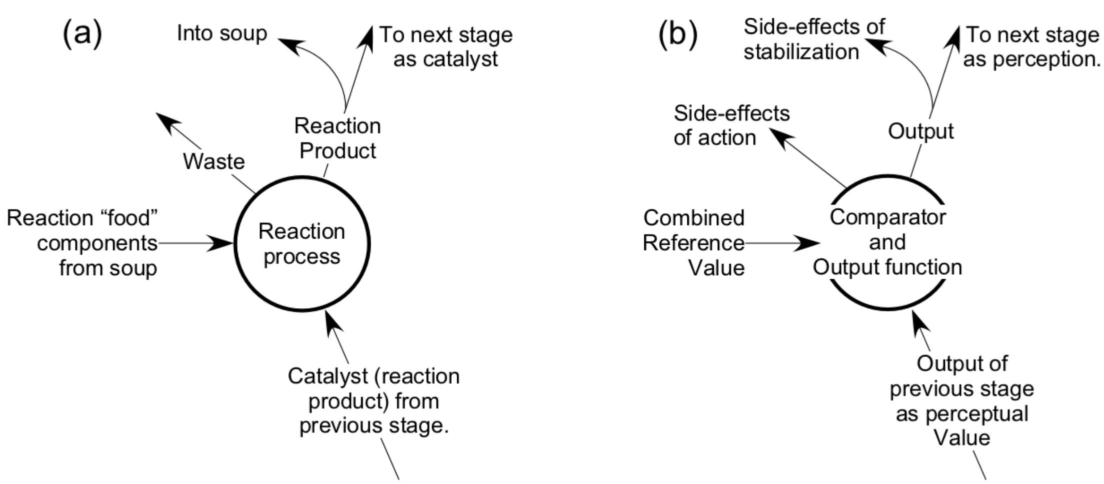

Figure II.3.3 (a, left) One node from the homeostatic loop of either Figure II.2.7b or Figure II.2.8. The node represents a reaction that combines "food" inputs of different kinds (or even more than one instance of the same kind) to produce a complex reaction product, some of which is used to catalyze the next reaction around the loop. "Waste" represents components of the input that are not part of the reaction product. (b, right) The same labelled as part of a perceptual control loop.

The soup is not homogeneous, even if it initially might have been, which is a rather unlikely prospect, no matter what the elementary units happen to be at any particular level of analysis. As the soup metaphorically or actually swirls and eddies under the influence of some external energy source, concentrations of all its components in any particular place or in relation to other components vary greatly. Those variations are the context in which the loop maintains homeostasis or control. In the loop of Figure II.3.2, for example, despite variations in the concentration of C or D, the concentration of CD is maintained near some ratio to the concentration of AB.

In Figure II.3.3, the same node structure is shown twice, once labelled as a reaction in a loop of chemical processes, once labelled as though it formed part of a perceptual control loop. There is a difference in the diagrams that may not be apparent. In the reaction loop the stage Gain is in the effectiveness of the catalyst, as was shown in Figure II.2.8, whereas in the control loop the Gain that is produced by the output function is combined with the comparator into a single node, as was done in Figure II.2.8 by the dashed ellipses.

Waste, in the reaction loop, occurs when a reaction such as AB+XY→ABX+Y produces two products, one of which serves to catalyze the next reaction around the loop while the other takes no further part in the autocatalytic process and is returned to the soup. With more complex reactions, both the ongoing catalyst and the "Waste" may be complexes that exist in the soup only because they are formed in the reaction, such as AABC+XYZ→ABX+ACYZ, in which ACYZ might be a waste product that is essential food for some other process, while ABX is both catalyst for the next stage and available for other uses in the soup.

In both parts of Figure II.3.3, the output of the node is shown as dividing into two streams, one proceeding to the next loop in the process, the other as contributing to the environment. In the control node image (Figure II.3.3b) the latter branch is labelled "side-effects of stabilization". These side-effects of stabilization are included among the effects considered by McClelland in the introduction to Chapter II.4. However, what is usually treated in PCT discussions as "side-effects" are not the stabilized product of control but the influence of the control actions on the rest of the environment, unrelated to the variables of the control loop. They are "waste".

In Volume 2, we show how both kinds of side-effect have the same importance for the development of society as the reaction product side branch into the soup does in providing the ingredients for proliferating complexity of products by the development of ever more autocatalytic processes. As we pointed out earlier, in the social context this proliferation is responsible for such "revolutions" as the Industrial Revolution or the currently ongoing Technological Revolution.

Earlier, we discussed how one autocatalytic loop might bud off another one that uses some of the products or pathways of its progenitor. Consider now what might happen to the waste product of a reaction. It might become food or a catalyst for some other reaction in the soup. Certainly, because a planet is a closed system for all practical purposes, at the very least it will eventually decay away into its elementary components. Kauffman's probability argument suggest that it is more likely to be used as is, or that when it decays, some of its parts will be used. Either way, the waste of an autocatalytic process is a stable product of the process, and can participate in other reactions that might well become links in an autocatalytic loop, perhaps homeostatic, or even a control loop.

Thinking ahead to suggest how this might apply to social processes, imagine a farmer who produces a crop as an atenfel for controlling his perception of his bank balance. The bank balance is the CEV of this process, and, strange as it may seem, the existence of the crop is part of the action to perform control of the farmer's perception. So far as the farmer is concerned, provided that the money goes into the bank, the crop might as well be dumped at the North Pole. In Figure II.3.3, it would be a side-effect of action. But the crop is, in the most literal sense, food for many other control loops, supplying energy, material for building and maintaining the parts of organisms that eat it, and probably providing many catalysts such as vitamins.

Putting all these possibilities together, the farmer might come to perceive that the crop he sells is "healthy". He then might begin to control for a property he might call the degree of healthiness of his crop, and to grow the crop for more than the single purpose of controlling his perception of his bank balance. In the spirit of Powers's many-to-many perceptual control hierarchy, crop-growing now supports more than one controlled variable because what was once "waste" now participates in an entirely different loop that has nothing to do with the bank. The farmer would grow the crop to control his perception of the benefit it provides to those who eat it, even if there was no buyer (provided he had some other source of income so that doing so would not involve a conflict in his control hierarchy).

When we come to language and the development of meaning at the beginning of Volume 2, we will follow through an example of side-effects becoming the CEV of a new perception, or of becoming atenfels in the control of a newly controlled perception. We will do this by examining how a new mother and her baby might learn to communicate meaningfully with each other without using existing language, in the process creating a new language.

II.2.3 Variety from Uniformity

Earlier, we talked about the multiplicities of autocatalytic loops that might well be expected to come into being before the whole structure freezes into one autocatalytic network of the kind described by Kauffman (1995). Later we will see how such a proliferation of autocatalytic loops that function totally independently may engender a higher-level autocatalysis, in the form of loops of loops. Now, however, we have a different process for the development of interdependent loops, each of which has a different set of properties in that it produces different kinds of product, even though both are descended from the same progenitor autocatalytic loop.

Though they have quite different properties, in the example of Figure II.2.9 each of the two independent loops depends on the other for its continued existence. This interdependence would be unusual, as most of the new loops would produce products that, at least in the first few generations of increasing complexity would be unlikely to catalyze any reaction in the old loop. To make Figure II.2.9 in my role as creator-deity for my imaginary world, I controlled for illustrating the possibility that the new loop might provide a catalyst that would function to repair and support some part of the old loop. By Kauffman's argument, this kind of interdependence is almost certain to occur on some occasions, if only because the new loop. like the old, produces an exponentially growing number of complex products out of the elements of the basic soup.

We may have here a prototype of the start of specialized function in a community, whether that community be of cells, or of people. Compare this to the slime-mould, which most of the time consists of a set of pseudopods of apparently identical cells that grow from a central "seed". A cutting from one part of such a slime-mould will act as a seed for a new one. When, however, a nutrient is in short supply (as with "E" in Figure II.2.9), the "community" produces two kinds of cell, one of which forms a stalk for a "sprouting body" to make a structure shaped somewhat like a mushroom from which spores are distributed to distant parts through the fluid environment. Or compare it to a company that starts by producing a simple product from scratch, gets big, and splits into project divisions that produce components, divisions that assemble them into products while also selling excess components, and divisions that coordinate the sales and the acquisition of supplies. Or, if others already produce functionally equivalent components or services, the company may start to buy them elsewhere rather than continuing to make them in house.

A budding new loop probably incorporates at least one reaction product of the original loop either as food or as a catalyst in one of its reactions. In Figure II.2.9, the AB product of the A+B→AB reaction is an essential food in the new loop, which incorporates a reaction AB+X→ABX. Of course, in the soup, as we have emphasized before, if there is enough variety of elementary or complex units to produce one autocatalytic loop, there probably will be several, all created without regard to the existence of others. We are not talking about them. They are in any case relatively unlikely as compared to a new loop that buds off an existing one, as we argue below. The developing buds that have formed into progeny loops are the ones that are likely to require food from a parent.

But we can look at it in another way. Apart from basic loops, ones that incorporate only reactions among the basic elements, which share no reaction in common — if any such independent pairs are likely to exist[4] — all autocatalytic loops require as food complex products produced elsewhere. No loop feeds on its own products, though all produce their own catalysts. When we think of our own bodies, we require a few basic elements such as carbon, hydrogen, oxygen, nitrogen, iron, zinc, selenium, and so forth, but in most cases the elements are of no use to us in elemental form. We need some kind of molecular complex that incorporates them, and such complexes are produced only by chemical reactions, mostly in other biological processes, each of which (including those in a human) is at its core a self-organized autocatalytic network.

Every autocatalytic loop that is created using any of the complex products built by the first loop is a descendant of that first loop, because it would not exist without those products. Only those that formed complete loops of catalyzed reactions using just the initial variety of elements in the soup would be truly independent. So what is different about the loops that bud off the original loop rather than just being built from its products? And is the creation of a second and third autocatalytic loop from the basic soup elements even likely? Let us think about that for a moment, because if it isn't, then it follows that life on any planet or isolated place in the Universe will belong to a family tree with but a single primordial progenitor, but on different planets or in different places, those locally unique progenitors might be very different.

Kauffman's argument suggests that where one loop can exist, more will follow, but it does not say that the following loops will share no reactions with the first one. Indeed, it is much more likely that they will share reactions that are already being catalyzed by another product in the first loop. Furthermore, they are likely to incorporate entire sub-sections of the original loop, building on reaction-catalysis-reaction pathways that are already producing complex products. The new networks very probably will incorporate these complex products as food or as catalysts.

Why are these outcomes more likely than that the new autocatalytic loops would use a completely distinct set of reactions catalyzed in a loop by their own products? After all, if there are 100 elementary units, there are potentially 10,000 different reactions between them forming 10,000 different reaction

4. Though it is quite possible that different first loops could have started in different locations such as different tidal pools or freshwater ponds and come into contact later. Patel et al (2015) and Pearce et al. (2017) both make this suggestion, but their analyses do not produce totally independent chemistries in the different ponds.

products. All of these reactions occur in the soup, producing their unique products. Any one of the products might catalyze any of the reactions, making 100 million chances that some product might catalyze some reaction. If there is a one in a million chance that one product will catalyze any randomly chosen reaction, one would expect there to be around 100 catalyzed reactions among those ten thousand, more than enough room for two separate loops to be formed.

But this back-of-a-small-envelope calculation leaves out a critical factor, that different elementary units will have different reactivities individually and different affinities between pairs. Some of the ten thousand theoretically possible reactions will occur never or hardly ever, while others will occur whenever their constituent elementary units come together, producing copious quantities of their product. These are the units most likely to participate in further reactions, either by being catalyzed or by producing a product that acts as a catalyst to some other reaction. They are also the units most likely to be involved in the first autocatalytic loop to form.

If the elementary units are chemical elements, for example, one should take into account that hydrogen, lithium, sodium, and oxygen are highly reactive, and that carbon has a high affinity with not only itself but also with many other elements. Such elements are likely to produce the most reaction products in the soup without being catalyzed at all, so they and their reaction products are the most likely to participate in the first autocatalytic loop. Indeed, all of the independent reactions proposed by Patel et al. (2015) as being likely to happen in separated ponds depend on hydrogen cyanide.

The next loop also is more likely to include these high-concentration products as food or as catalyst than to avoid all of them in favour of reactions and reaction products of less reactive elementary units. The general effect of these varied probabilities is to indicate the relative unlikelihood of the soup engendering sets of independent loops as compared to the likelihood that all loops seen in the soup after a while are descendants of buds sprouted by one original loop. No matter how many levels of loops of loops develop over time, and no matter the complexity of the product, all are almost certain to be the descendants of one progenitor, no matter where in the Universe life may in future be found. Furthermore, if Patel et al. (2015) are correct, then wherever life that depends on liquid water is found, it will have the same basic chemistry as here on Earth.

II.2.4 Membranes, Tubes and Cells.

Over the last two or three Chapters, the word “concentration” has been bandied about rather freely. But many commentators (***REFS***) have pointed out that in a primeval ocean, the concentration of any reaction product is likely to be very low. There is a lot of hydrogen and oxygen, but they are tightly bound together in the form of H_2O (water), and there are few other elements in high concentrations. A solution frequently suggested is that reaction components are more likely to meet in two dimensions than in three, so the surfaces that bound the ocean, tide-washed shorelines, are the most likely place for useful rates of reaction to occur. A surface acts as a kind of general catalyst, and on a surface, things may grow that would not grow in a free-floating space.

For our argument, it does not matter whether the initial places of relatively high concentrations are surfaces or “warm little ponds” (Pearce et al., 2017). Some such place of enhanced concentrations will be important to increase the probability that reaction products will meet. In the next few paragraphs, we will assume that the place is a tidal pool, but the same kind of argument can be applied to a pool fed by rainwater or a relatively small stream and depleted by evaporation or overflow in times of flood. We will find that the same necessity for regions of concentration of complex products holds in the vastly different realms of innovation in societies of organisms and of human invention. “Silicon Valley” is an example.

We may assume that growth initially occurs on a surface, but when the autocatalytic loops begin to be important, a two-dimensions space might be restrictive. A single instance of an A and a B may react and leave an AB behind, but this particular AB molecule will not catalyze a C+D reaction unless it happens to be close to a C and a D. The same is true all around the loop. It is quite true that such loops could physically form ring structures in which no component is far from places where its products could catalyze the next pair of components. But that would be a "one-shot deal", that would leave one instance of each reaction product lying on the surface with nothing more to do. Everything must move over the surface, and at concentrations sufficient to sustain an autocatalytic loop of any significant productivity, they would be likely to interfere with each other quite badly.

Furthermore, the reactions in an autocatalytic loop require a through flow of energy, so the reacting components require to be near an energy source and an energy sink. The surface of the ocean exposed to the air is a place where things move around, and where energy from the sun can be dissipated into the cooler water. Too much motion would tend to disrupt long complex loops, while too little would tend to restrict the access of the components and catalysts to each other. We might suspect, as have many others (***REFS***), that the earliest autocatalytic processes might have been built in shallow sheltered pools on the sea shore. When the autocatalytic process began to produce complexes of sufficient complexity, it is easy to suppose that the surfaces of such pools would be essentially covered with the complex products, many of which would have a carbon backbone. The surface of the pool might be what we could call a scum.

Scum or not, in such a pool, the three-dimensional water content has little exchange with the greater ocean on the time-scale of which we speak. The time-scale of variation in molecular processes may be nanoseconds and of complicated suites of processes perhaps milliseconds or even seconds, but the exchange of water in sea shore pools is on the time scale of the tide, which even in those early days when the Moon was closer and the Earth day was shorter would have been hours rather than seconds. The surface scum might well have been stirred into the depths of the pool by the wind. Even there they would have been far more concentrated than when swept into the greater ocean at high tide.

One can see such a semi-landlocked shallow pool covered by scum as a kind of grotesque large-scale cell bounded by a membrane that keeps its complex molecules at concentrations much higher than in the environment outside the cell, but that is not the line of thought I want to follow at this point. Instead, I want to think about natural selection as being a process that has nothing intrinsically to do with life, but that would work no matter what the structure. Of structural types that tend to recur, those that recur more often in their environment are the ones that will dominate the population over time, if the two structure share a limited resource. Natural selection will apply automatically to the autocatalytic loops in the pool.

On every tide, some, but probably not all of the "stuff" that was in the pool will be swept into the wider ocean, and the water in the pool will again have a low concentration of the food for the autocatalytic loops that might have remained. A new scum may form before the next high tide, but it, too will be swept away when the high tide comes. Each tide, different components of autocatalytic loops are likely to be left behind in the pool, but not necessarily always the same ones. Some of these loops may produce complex products that physically stick together as though they were part of a loop, though they are not. The waves of the high tide would be likely to break them up. But some would be "stickier" than others, and they would survive, even after being swept out to sea. Those would eventually form more of the population than would the fragile ones.

This process would have been going on in millions of pools scattered around the planet, in the hot equatorial regions, near the cold poles, and everywhere in between. Any one of them might be destroyed utterly by each high tide, but not all of them every tide through thousands, millions, or billions or years. Some will continue developing through many tides, and in some parts of the world the tides are less extreme than in others. In some of them, the "sticky" complexes might begin to form platelets or sheets, that when bound to the rocky perimeter of the pool, might form a roof to the pool.

But there is another possibility, which is that the complex molecules formed with carbon rings or backbones are seldom flat. They have extensions in all directions. We know that pure carbon will form the tubes we know as nanotubes in just the heat of a flame. Is it too much to suggest that more complex molecules containing carbon might also form tubes from complex carbon-based molecules in one or more of these millions of autocatalytic pools some time in the millions of years and billions of tides? Pearce et al. (2017) suggests that the reactions they propose in the "warm little pools" would need to be complete orders of magnitude more quickly than that, and those products could well wash down into seashore pools.

Tubular structures would be built of molecules that participate in autocatalytic reaction loops, so their interiors would include products of those loops, protected from dispersion by the tides, even if the tubes themselves were swept out to sea. Some tubes at some time might happen to have one or both ends closed. They would be containers of a mixture of complex components that participate in autocatalytic networks, prototype cells, and such cells would tend to be survivors along with tubes.

When you add that many of the autocatalytic loops that include the molecules that constitute the membranes of the tubes and cells would have extended both within the cell and into their exterior environment, and that some of these loops would have been control loops, you have a possibility for stabilization of an interior variable by action on a variable in the environment, an early, chemical, form of perceptual control. The "cells" at this point are not at all like the cells in our bodies, having in common little other than that they are capsules bounded by semi-permeable membranes that have distinct relationships between the chemistry in their surrounding environment and that in their interior.

Cells that were able to control internal variables and increase the local chemical stability of their environment this way would have a survival advantage over cells that merely had an autocatalytic or even a homeostatic inside-to-outside connection. Having a survival advantage, they would be the type that later constituted almost the entire population. The "waste" side-effects or the stabilization side-effects of one cell type might easily be the food or the catalyst of a reaction in another loop that is partially inside a different cell, in the manner of a budding-off loop. Such connections lead almost inevitably to the specialization of cells in multi-cellular clusters. And that is as far as I want to push this extreme speculation.

These speculations are extremely far removed from evidence, and are totally unjustified as scientific descriptions of early pre-life. I introduce them mainly to suggest the possibility of linkages between the early autocatalytic structures that Kauffman has shown almost certainly would have existed and the complex physical structures of cells, tubes, and fibres of which animals and plants are constructed in the present. The important linkage is from autocatalytic loops of which some would have been homeostatic to homeostatic loops that have only two nodes — control loops, of which some would act through a membrane to stabilize the local environment by way of stabilizing internal variables.

But is it merely a coincidence that Sallan et al. (2018) find that the ancestor of all vertebrates developed in the kind of environment described above? And that the ancestral forms of many different groups of vertebrates, such as those with jaws, for example, also started in such pools or lagoons, and not from forms that had dispersed to the wider ocean. And further, that such tidal marine environments are even now hot spots of species diversity (Pimiento 2018)?

II.2.5 What's in it for the Cell?

The title of this Section was inspired by a message to the CSGnet mailing list by Bruce Nevin

(2018.09.11) in which he wrote:

> *This question has long been in my mind: "What's in it for the cell?". It refers to the magic of embryology, maturation, learning and healing generally in any multi-celled organism, but here particularly to the cells that constitute the neurological embodiment of a hierarchical perceptual control system.*

Living things visible only in the most powerful microscopes are very much simpler than are multi-cellular animals and plants, and yet are subject to many of the same physical constraints as are we. Much of what we discuss in the rest of this chapter will apply equally at multiple levels of size and complexity, from individual cells and their internal structures to identifiable cell assemblies, to organs and body parts, to entire animals and plants, and even to social structures called ecologies in which entire animals and vegetables are the elementary units, and onward to the largest structures that our one world can hold in multiple versions. To the microbes, we are an infrastructural environment that they have built, and they help to maintain us as we maintain our social infrastructure.

So, let us take a brief look into a social structure millions of times smaller than us, that of the cells within any organism. A cell is rather simpler than is the organism of which it is a part, but as with social organisms from slime molds, schooling fishes, hunting wolf-packs and dolphin packs to human families, teams and nations, many individual entities of similar structure may work together to achieve results that they could not achieve by themselves. One person cannot move a Stonehenge-type thirty-ton stone, but a few hundred people can transport it over hilly terrain. The billions of cells in an organism may vastly outnumber the tens of individuals in a social group, but the principles that govern their interactions are close enough to being the same that cells can serve in some important respects as an effective model for human society.

Nevin's question: "What's in it for the cell?" has a simple answer that fails, and a more complex answer that succeeds. From the point of view of the cell, it would be bathed in the same nutrients and could receive the same chemical and electrical signals whether or not it participated in the control operations of the organism acting through the organism's external (possibly social) environment. So the simple answer is that there is nothing in it for the cell. Nothing would (apparently) change for it if it ceased to participate in control or even if its actions proved detrimental to the performance of other controllers. It would be a "cheater" by taking benefits it had not earned by its contribution.

That simple, selfish, answer is the answer that would be seen from the viewpoint of the cell, and it is the same answer that would be given by a person who claims that his or her success was entirely due to their own efforts, while giving no credit to the supporting influences of other people in his or her social environment. A naïve view of Perceptual Control Theory has led some contributors to the CSGnet mailing list to assert that PCT supports such an ultra-libertarian political viewpoint. It is a viewpoint often espoused by very rich people.

Consideration of single isolated control loops, ignoring environmental influences apart from combining them to form "the disturbance", is certainly consistent with the idea that the individual will — the reference value, applied to the environment with sufficient loop gain eventually through the forces applied by the muscles — is all-powerful. Leni Reifenstahl's masterful 1935 Nazi propaganda film "The Triumph of the Will" (in the usual English translation of the title) was intended to make this point, though in this case, the "will" was only that of the Great Leader, or Führer, allowing the masses none of their own.

This first answer to Nevin's question takes the "Controller" viewpoint from inside the cell, but it is just one of the available viewpoints. Two other viewpoints that could offer a different answer are "Observer-Experimenter" and "Analyst-Theorist", or "See what happens" and "Figure out why things happen". We have earlier called them just "Observer" and "Analyst" and I will continue to do so. In this case, the Analyst viewpoint comes to a very different answer to Nevin's question, which the analyst answers "*Nothing*" and substitutes a slightly variant question: "*What's in it for the cell is survival of cells like itself into the future*." It's the same answer as one could give to an equally puzzling question: "*What's in it for martyrs who willingly allow themselves to be killed*". That answer is "*What's in it for the martyr is survival of people of their kind into the future*."

This does not contradict the simple answer that there is nothing in it for the cell. Rather, it extends it. What's in it for the cell happened long ago, to its ancestors who survived when others of similar kinds did not. The cell does what it does because way back in history, cells like it, doing what they did, survived *because they and helped organisms like them to propagate more*. The cell benefits by the very fact that it exists, which it does because its properties are like those of its ancestors. If in important ways they are not, it may have few descendants. A cause that inspires martyrdom is likely to be one worth joining. The martyr may inspire more new adherents than the death of the martyr dissuades.

The Analyst can see that a cell that contributes to the ability of its organism to maintain its structure by controlling effectively in its environment lives in an organism that has a higher likelihood of survival than would an otherwise identical organism in which that cell exists but has random effects. Those effects might even detract from the ability of the organism to survive. Such a cell may itself live happily in the environment of its organism, but its organism is less likely to survive to produce descendants, because of the energy used in maintaining the cell that does not contribute. Cells that contribute to the welfare of the organism, meaning that they assist in maintaining its structure (their environmental infrastructure), will leave more descendants in the progeny of the organism in which they live.

Even a virus that kills the host organism that provides its necessary infrastructure may prosper individually by having many direct descendants, but at the expense of its kind if it is so virulent as to deprive its future descendants of hosts. Such a virus type either must find other hosts in which it is less virulent or it will become soon extinct along with its host species.

A popular child's sing-song taunt is (at least it was popular in my childhood) "Cheaters never prosper", but experience suggests that in modern Western society this is not true. Like a virulent virus, cheaters often prosper individually; it is the society that includes and allows cheaters that does not prosper. Likewise, a non-contributing cell may save its own energy by not contributing to the survival and prosperity of the organism, but the organism uses energy to maintain that cell. The organism is worse off than it would be if it simply killed the cell, and a lot worse off than in would be if the cell were to begin contributing to its ability to control — enhancing the organism's "worth" in the language we have been using.

How the society that constitutes the organism achieves this effect is a question about how "policing" systems might have evolved. For the organism, the most effective "policing" is the reconfiguration of the cell to do something useful to the organism, but failing that, removal of the influence of its actions by encapsulating ("imprisoning") it or killing it also reduces the organism's energy costs of maintaining the unhelpful cell and of retaining effective control. These policies have their equivalents in most, if not all, human societies and many non-human ones as well.

II.2.6 A cartoonish "e-coli" and its descendants

Imagine a single-celled organism, a much simplified version of the e-coli bacterium that was the inspiration for Powers's "e-coli" method of reorganization[5], a hill-climbing optimization technique that assumes no knowledge of gradients. The Powers "e-coli" entity is a marker that occupies a point in a multidimensional scalar field. The term "scalar field" means that every point in the space has a numerical value associated with it, which in the present context we might as well call "fitness" or "quality".

The Powers e-coli marker point moves in one of two ways. Either it continues going in one unvarying direction, which it does so long as the "fitness" of its location continues to improve, or it "tumbles" (changes to a new random direction). For reorganization purposes, the location in the space is defined by the values of connection parameters among the control units in a control hierarchy, while the "fitness" is some function of the health of the intrinsic variables, or more commonly the quality of control of the units in the hierarchy.

Our Powers-type pseudo-e-coli or "p-coli" organism is more than a point in a space. It is a cartoon version of a real e-coli bacterium that moves through a fluid "chemical soup" in three dimension. Its "fitness" is the concentration in the soup around it of a nutrient that it uses as food. The p-coli is conceived as a cylinder longer than it is wide, which has only two important properties. One is that it is able to produce descendants at a rate that increases with its rate of food consumption. If it doesn't "eat" enough, the p-coli eventually starves to death, because it needs energy continuously to maintain its structure against the entropic decay that would be implied by the Second Law of Thermodynamics.

A small number of the descendants of a p-coli have a mutation that might repurpose some structure in the original or might even incorporate a new structure using the old as a component. We will concentrate, but not exclusively, on mutations that duplicate an existing structure. Later, we suggest why this is reasonably likely to happen. The second property of the original "p-coli" type is that it can sense a longitudinal gradient of the concentration of the nutrient chemical in the "chemical soup" in which it lives, and can act to move ahead in the way it was going or to tumble.

Figure II.3.4 suggests the functional organization of the original "p-coli" and a possible mutant descendant "di-coli" in which the single sensor-motor array structure has been entirely duplicated and the two arrays have been oriented so that the two motors act in somewhat different directions relative to the long axis of the cylinder. Each sensor-motor array of the di-coli acts exactly like the single array in its p-coli ancestor, driving ahead or causing tumbles depending on the direction of the concentration gradient. The resulting catalogue of motions of the organism, however, is not exactly the same.

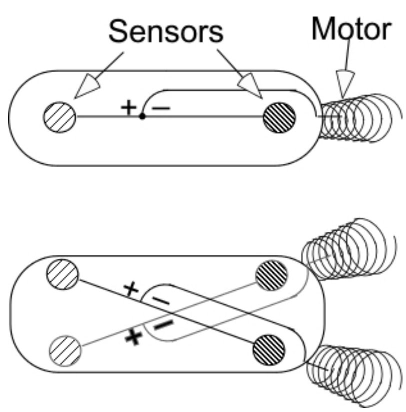

5. The actual organisms that might have evolved are not irrelevant, but are both much more complex than would be useful in making the point, and (more importantly) are not known to me. I present an entirely hypothetical and trivially simple evolutionary process that clearly is entirely different from what actually happened, though I believe that in principle much the same is what did happen.

Figure II.3.4 (top) The original p-coli. When the head sensor (left-end) senses a higher concentration than the tail sensor, the motor spins one way, driving the organism forward, but if the reverse is true, the motor spins the other way, causing the organism to "tumble" in place. (bottom) A possible mutated descendant, "di-coli", in which the sensor-motor array has been duplicated. The motor driving the organism more directly up the gradient will tend to slew the travel direction into the direction of maximum steepness, and if either motor causes a tumble, the other will bias the direction resulting from the tumble.

The mutated "di-coli" organism with two motors will "steer", in the sense that if one motor drives harder than the other, the heading of the organism will change. If the di-coli motors drive harder the greater the sensed gradient, then di-coli will be quicker in finding regions of high concentration than will its p-coli ancestor, since it will be preferentially slewed toward the direction of the steepest gradient. If one motor causes a tumble, the other will bias where the di-coli points, until both motors are driving up the gradient again.

In a soup that contain nutrients (food) shared by the many p-coli and initially few di-coli, the survival and the rate of producing descendants of any one organism is enhanced by its spending time in regions with high food concentration. In a liquid "soup", such high-concentration regions do not stay localized very long. They move about in the flows and eddies of the liquid. Accordingly, the ability to move quickly to regions of higher nutrient concentration enhances the number of di-coli descendants compared to those of the ancestral p-coli. Over time, the soup will be dominated by di-coli, with far fewer p-coli. This is simple Darwinian "Survival of the Fittest," and it will occur even if there is an infinite supply of food.

The type of di-coli shown in Figure II.3.4 is far from the only configuration that could result from a simple reduplication of the sensor-motor array as a unit. Here are some more possibilities, one of which (type c) also contains a connection error that occurred in the duplication process.

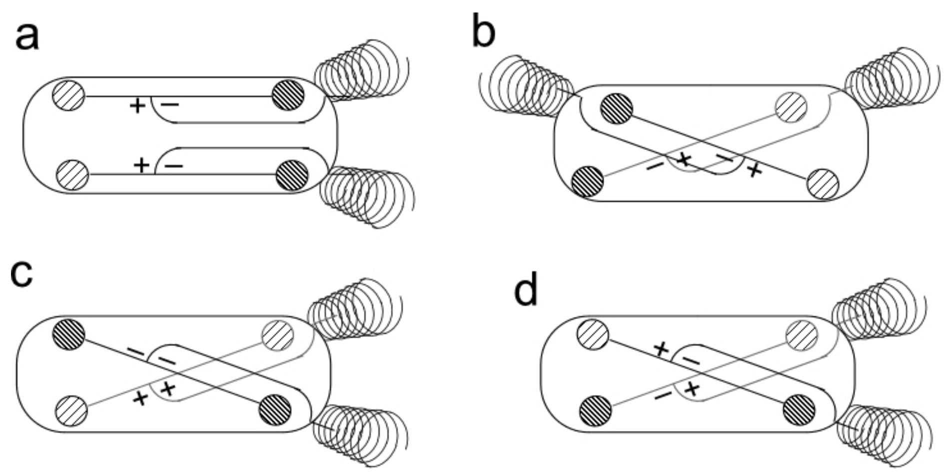

Figure II.3.5 Some di-coli configurations that do not work very well in getting the organism into regions of high concentration of nutrient. (a) behaves the same as p-coli, so is no improvement but wastes energy, (b) usually tumbles but sometimes moves sideways a little, (c) since the two ends are compared, one motor is always tumbling while the other is always driving, and (d) almost always has one or other motor tumbling, though it can move directly sideways up the gradient of concentration.

Apart from (c), these configurations differ only in how the two sensor-motor arrays are oriented with respect to each other and to the cylinder axis. In (c), the comparisons have got themselves cross-wired rather than producing two identical sensor-motor arrays. In none of these mutations do the new arrays contribute much, if anything, to the speed with which the resulting di-coli finds high nutrient concentrations, as compared to the original p-coli. Type (a) is as good as its ancestral type and is perhaps a little faster moving, but to construct and maintain the configuration requires more energy than does the simple p-coli. If the extra nutrient gained by the added speed more than compensates for the excess energy requirements, this variety will outcompete the p-coli type, but not by as much as would the “good” di-coli of Figure II.3.4. In all the other cases, the addition of the second sensor-motor array actively detracts from the survival probability of the organism.

All these “mis-configured” potential di-coli types will soon be dominated by well-configured di-coli in the environmental soup, and although examples of their type may continue to be produced and to produce descendants, they will contribute little to the population in the long term. On the other hand, well-configured di-coli will soon come to predominate in the soup. Initially, though, none of the di-coli may be able to pass the relative orientations of the two sensor-motor arrays down to their descendants in their design template (read “DNA”). All of them will have di-coli descendants in which the relative orientations of the two sensor-motor arrays are random. Nevertheless, there will ordinarily be more surviving “good” di-coli than ineffective ones, because most of the bad ones will starve quickly, while the good ones live in “food palaces”.

The mere fact of duplication of the sensor-array in the di-coli descendants implies a change in their design template from that if the ancestral p-coli, and it is not very radical to suggest that one of them might at some point acquire a mutation that adds its orientation to the template. Because there will be more good than bad p-coli in the soup, and bad p-coli that inherit the bad orientation will starve as quickly as the parent, the “good” orientation is the one that will be incorporated in the most di-coli in later times.

The “steering” function of the di-coli is not control, since there is no interval variable whose variation would affect the relative power applied to the two motors. Each motor drive depends entirely on the gradient in the direction in which it drives. There is no link in a di-coli that senses the difference between the two gradients, and nothing in either of the two duplicates of the “drive-motor genome” suggests the possibility that one might evolve. Nonetheless, among the myriads of instances of replicate di-coli, some mutations are inevitable, including a possible reduplication of the already duplicated part of the genome relating to the motor mechanism that determines the gradient (Figure q15.6)

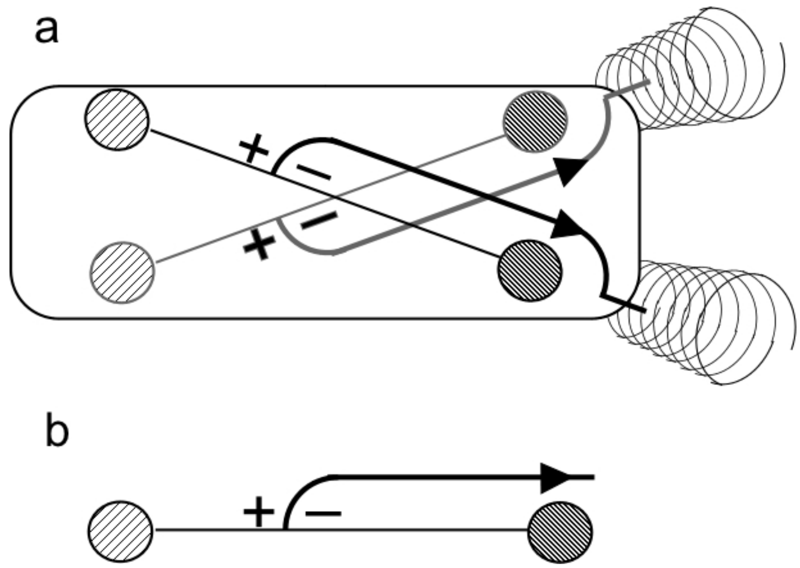

Figure q15.6 A di-coli (a) and a further duplication of the part of the genome that generates the drive mechanism of the original p-coli (b). The black triangles represent output that specify the drive power to wherever it goes. In the p-coli and di-coli, it goes to the motors. In the further mutation, the extra drive mechanism is not yet connected.

You may notice that the motor mechanism has a familiar shape, that of an elementary control unit with a reference value fixed at zero. The gradient sensor is a perceptual function that reports the difference between the food concentrations at its plus and minus ends, while the output function produces the signal sent to the motor drive, which cannot go negative because any action by the motor in the surrounding environmental "soup" would either impel the p-coli or di-coli forward or would prevent the gradient sensor from doing its job.

We considered this "one way" problem in connection with the inability of a nerve to fire at a negative rate in Section I.4.6 on the comparator function (Figure I.4.7), and in Section I.8.1 on control stiffness (Figure I.8.2). In both those cases the circuitry was already built and we simply described it. Now we are asking about plausible mutational routes from p-coli, which produces "tumbling" when the gradient reverses, to a descendant of a descendant that produces "steering toward the food supply".

The first p-coli descendant by gene duplication, di-coli, still is expected to have motors that cause tumbling if their sensed gradient goes negative, but the new "steering" descendant has an different and new function, which does not involve tumbling. Its gradient detector does not receive input from an environmental sensor, as do both ends of the gradient detectors in p-coli and di-coli. So what gradient is it likely to measure? Where do its inputs and outputs connect and what might the new set of connections do?

Without going into the probabilities of this and that useless configuration, so long as there is an accessible useful configuration that gives the mutated di-coli even a very small fitness advantage over the original di-coli, the mutated version will tend to become more dense in the "soup". One fitness advantage might be if the new drive mechanism sensed the difference between the two gradient detectors of its parent di-coli, and used the value of that difference to augment the drive to one motor while inhibiting the drive to the other.

A "tri-coli" with this modification to a di-coli template would drive more directly up the food concentration gradient than would any of its ancestors. By the same token, it would be more agile than its ancestors in following swirls and eddies in the soup that moved the high concentrations around. It would have created a two-level control hierarchy.

This example illustrates a canonical form of answer to questions of the kind posed by Nevin "*What's in it for the cell*?" In this example, the question answered has been the more generic question "*What's in it for the structure*?" This is a more generic question, since "the cell" is just one kind of structure among many, from the internal cytoskeleton of a real-life cell to cell assembly patterns that are repeated throughout the mammalian brain, and from there through structures of social roles. What's in it for the structure is "nothing", except that those structures that help their host organism or society to survive are more likely to be found in the future than are those that don't.

There is more to be said about p-coli and di-coli and their descendants. Take the structure of Figure II.3.5a, in which two duplicates of the p-coli seem to perform the same function in parallel, giving no apparent advantage other than travel speed to the di-coli over the p-coli ancestral type, and incurring an energy cost for the construction and maintenance of the extra copy. But that disadvantage is small, if it even exists, and might even be overcome if the double sensor-motor array allows it to survive damage to one of the arrays by limping along with the other motor while the first is being repaired. Repairs, as we shall argue below, would use the (genetic) design template that allows the di-coli to produce descendants.

So long as some of the "parallel" di-coli exist in the soup one of them might mutate further, repurposing one of the arrays so that it responds to the concentration gradient of a different chemical in the soup. If that chemical also was relevant to the welfare of the di-coli, the ability to seek out regions of the soup with a high concentration of both would give that mutation a better probability of survival to reproduce than its ancestral di-coli would have. Likewise, a mutation might create a sensitivity to a noxious chemical, and if the mutation of the chemical structure of the sensors reversed which end happened to be the more sensitive, such a mutated di-coli could simultaneously approach nutritious concentrations while avoiding noxious regions — even if some region contained both temptation and poison (Figure II.3.6).

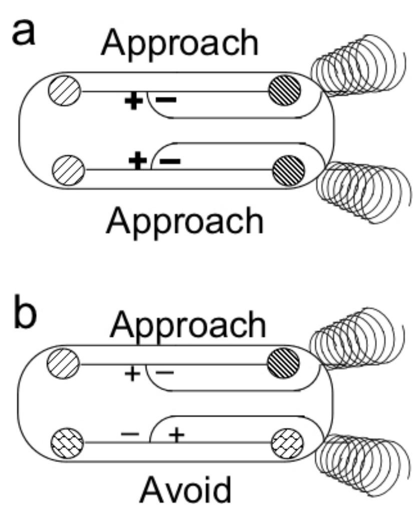

Figure II.3.6 (b) A di-coli second-generation mutation descended from the (a) type, that is attracted to one kind of chemical and avoids another kind, rather than both sensor-motor arrays being the same.

In such ways, reduplication, mutation and repurposing are capable of producing complex structures that have a variety of behaviours that enhance the survival of distant descendants in soups that contain many chemicals. At the same time, structures that do not contribute to the survival of the organism will tend to be eliminated from the population mix over time. What matters is the structure. The same structure is used for avoidance as for attraction in the example of Figure II.3.6, using different sensor elements.

The elements are not irrelevant, but they are not the be-all and end-all. They need the structure, as the structure needs them. We will see many examples of the importance of the same stable structure built with different types of elements throughout Volume 2. As we described in Chapter 8, the structural motif, not its elements, produces the emergent property.

A template description for a structure that is replicated is rather more concise than the combined descriptions of two patterns that by chance happen to be the same, but an over-precise description of any structure requires far more information than does a description that allows a few loose ends. Furthermore, to maintain the structure against random destructive changes is much more costly for a precise structure than for one with a little more tolerance.

In Section I.4.7 we saw some beneficial effects of tolerance in active control. Now we see that the same is true when the control systems in question do not act through the environment, but act to restore structure according to a template description. Tolerance implies the possibility of leaving uncorrected some small changes — mutations. And we will argue that the same is true of social structures, which do not evolve by intergenerational change through mutation, but are maintained by collective control according to learned reference templates (Volume IV). Some tolerance allows this collective control to function without overt conflict, but it allows also for cultural drift in a Lamarkian kind of pseudo-inheritance.

Tolerance saves energy, as any mother knows who has first tried to make her children behave perfectly but later relaxed her standards a little. With this thought, we are edging toward consideration of Prigogine's work on self-organization in energy flows, but we are not there yet. There is more to consider.

Societies do not create descendants and then die, as do the p-coli and di-coli, and as do all real living organisms. They change over time, split and merge, so that over a few generations many of their features change dramatically, which has much in common with the "propagate and die" way that the features of individual organisms are manifest in later generations. Both have internal structures, some of which contribute to the maintenance of their overall structure, some of which actively disrupt the structure. These latter die out, or the society changes dramatically to accommodate them, but society and disruptive individual types usually do not both remain the same for very long.

The longevity of a society in a recognizably stable form will depend on the balance between the influences of maintenance and of destruction, as well as on the way the society as a whole interacts with other societies in the world.

II.2.7 Tightening the homeostatic loop: Eukaryotes

In Chapter III.2, we discuss side-effect loops in social situations. Here, we note that the production of a complex that acts as a catalyst is a simple side-effect of the reactive potential of the progenitor component units. Side-effects are the same, whether at the scale of single molecules in a soup or the scale of interactions among social complexes. The phenomenon that forms the focus of this Section has its analogues at many scales, some of which we touch on in various places later.

Side-effect loops are inherently fragile, and the longer the loop the more fragile it is. A break in continuity anywhere around the loop destroys the self-supporting property of the processes in the loop, whether they be simple reactions that each produce a catalyst for the next process around the loop, or the complex social interactions that have side-effects that affect apparently independent social structures. In this Section we use a hypothetical two-process homeostatic loop, from deep in evolutionary history.

Life on this world has three main branches that diverged early in evolution. They are archaea, bacteria, and eukaryotes. The first two were and are unicellular, while eukaryotes eventually formed multicellular organisms, including us. Eukaryotic cells are strange, in that they have two sets of DNA, one compacted in a nucleus, the other encapsulated in a separate membrane-enclosed object called a mitochondrion, several of which may exist in a single eukaryotic cell. The nucleus and the mitochondrion have independent DNA contents, the mitochondrial DNA being contributed solely by the mother in sexually reproducing species such as ourselves, the nuclear DNA being provided equally by each parent.

The mitochondrion supports the eukaryotic cell in a variety of ways, the details of which are irrelevant to our story, but which include providing energy from its processing of chemicals absorbed through the cell membrane. The containing cell in return supports the continued existence of the mitochondrion it contains, creating at least one homeostatic loop.

It has long been assumed that the mitochondrion is descended from a bacterium that somehow got incorporated into a cell that did not originally have one. One suggestion for the genesis of eukaryotes from archaea and bacteria has been provided by Imachi and co-authors (Imachi et al., 2020, summarized by Schleper and Sousa, 2020). Imachi et al. accept that the ancestor of the present-day mitochondrion was a free-living bacterium species. They argue that the ancestral host is likely to be a member of the Loki family of Asgard archaea, which nowadays live in the cold deep ocean sediments.

What the host cell and the mitochondrion do for each other as side-effects is not really relevant to this story. In Volume 2, we will be interested in perceptual control in social situations, and I want to use the mitochondrial example later when we discuss interacting social structures. There is only one known example of the biological mitochondrial embedding in a host cell. It is often claimed to be the only one that ever occurred, but is this necessarily true? Can we say any more than that this example that was ancestral to all eukaryotes is the only one to have survived the test of evolutionary time long enough to have produced fossil remnants for us to observe and be assured that they were our ancestors?

Baum and Baum (2014, as described by Dance (2021) proposed what may have happened to create first eukaryotes from a bacterium engulfed by a tentacled archaeon that used the same nutrient. Both benefited from the arrangement, if the waste product of each could be used by the other. Here I follow Imachi et al. (2020).

What actually happened might be quite different, but one thing probably has not changed: the generation time of bacteria is thousands or millions of times shorter than the two to four week generation time of the archaeon living in the cold deep ocean that Imachi et al. think likely to be close to the ancestral type. So, for the purposes of the example, we can assume that a single engulfed bacterium would generate many copies inside the archaeon, and that essentially all the evolutionary change to turn the bacteria into mitochondria was in the bacterium, not the host cell, though the host type must have evolved to allow the bacterium to be encapsulated. Figure q14.4 suggests four possible early stages in the evolving relationship. Figure q14.4c is inspired by Figure 1 of Schleper and Sousa (2020) in their summary of Imachi et al. (2020).

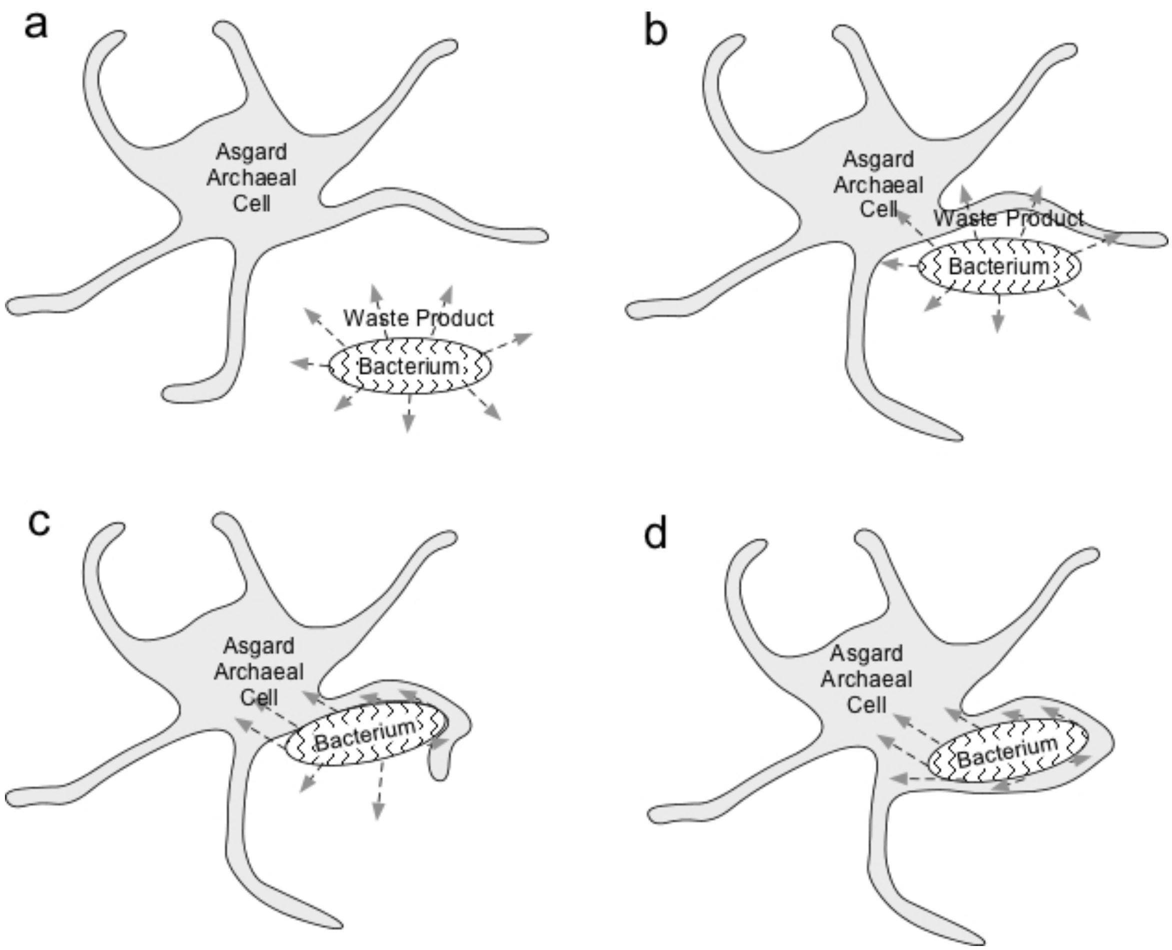

Figure q14.4 A proposed evolutionary beginning for the eukaryotic cell line. A bacterium produces as a side-effect a waste product that improves the control of the internal processes of a Loki-type Asgard archaeon, which produces side-effect waste products beneficial to the bacterium (not indicated), creating an autocatalytic loop. (a) bacterium and archaeal cell float freely past one another, benefiting each other only slightly. (b) The bacterium line evolves to move toward source of the Archaeon's food-like waste. (c) The bacterium evolves to maintain membrane to membrane contact with the archaeal cell, maximizing its supply of "food" as an independent entity. (d) the bacterium manages to enfold itself within the archaeal cell, guaranteeing both a constant food supply.

When the bacterium is free-floating, as it is shown in panel a of the Figure, it gains the benefit of the side-effects of the control processes of the archaeal cell only when near to a member of that cell class. The relatively sessile archaeal cell likewise receives the benefit to its control processes only when a bacterium of that class swims nearby. There is a side-effect loop by which the two types of cell survive better in the proximity of each other, each providing the other with a waste product that as food improves its ability to control some important variable.

The evolution of the bacterium species would lead an individual of the species to control for approaching any member of the right type of Asgard cell. The evolution of the Asgard cell would tend toward assisting the bacterium to adhere to the surface of the cell, following which both would have their survival further enhanced by direct trans-membrane communication of the useful waste products. The final and most difficult phase would be the encapsulation of individual bacteria by individuals of the Asgard species, followed by reproduction of the encapsulated bacterium and Asgard cell together as eukaryotic cells. The encapsulation of individual by individual is not itself difficult. We eat food that benefits our gut microbiota in ways that enhance our health, but we do not incorporate the gut microbiota into our genome.

The closer and the more reliably a member of the right type of bacterium is close to the right type of Asgard cell, the more interchange of these waste products will occur. To the engulfing archaeon, the bacterium supplies a food that at least complements the food it might acquire from its environment. Over time, the multiple copies of the bacterium evolve into what we now know as mitochondria, supplying energy to the surrounding cell, which evolved from the body of the archaeon. This new energy source allows the archaeon freedom to increase its motility, and to produce more descendant chimeric cells, the first eukaryotes. The summary by Dance (2021) describes how the chimera's descendants might have evolved to use a common mechanism for dividing the component archaeal cell body and each of its included bacteria-descended mitochondria.

This hypothetical scenario can play out only if the host and the bacterium have several distinct properties. First and foremost, the final eukaryotic encapsulating host cell that includes the bacterium must either propagate faster or produce more descendants per generation than its archaeal ancestor type, while at the same time incorporating the included bacterium in its descendants. If it did not, each of its descendants would need independently to catch and encapsulate a free-living bacterium passing by. For this coordinated propagation to be possible, the process that copies DNA to a descendant must include in its functioning the DNA of the bacterium, as if the bacterium had become a subroutine in a program for DNA replication.

We don't know how mitochondria first became associated with other cell types such as the archaeon proposed by Imachi et al. (2020) in such a way that they could produce descendants together in the form of eukaryotes. We don't know whether mitochondria began to be descended only from the maternal line when sexuality first started, or whether the restriction happened much later. Such questions do not matter for this book.

There has been a lot of research over the last couple of decades about how common descent happens now, some billions of years later, when eukaryotic cells may contain lots of mitochondria, but all of it is irrelevant to the way that I want to use it as an example to which I can refer when I deal with encapsulation in various social conditions later in this book. My simple-minded version serves that purpose better, I think, than a complex discussion of biochemistry would, even if I were an expert in the topic rather than a rank amateur.

II.2.8 Cellular Clusters and Conductive Tubes

Among the "revolutions" that can occur from the development of autocatalytic loops of loops, we must include the Industrial Revolution of the 18th and 19th centuries. In part this revolution was one of communication, railways making possible travel and the transport of goods between areas of Britain (initially) in hours that had taken days by horse-drawn vehicles on rutted roads. Communication by good roads had enabled the development and maintenance of the Roman Empire. Communication by the equivalent of the "Pony Express" of the US "Wild West" held the enormous Mongol Empire together.

The invention of the electrical telegraph along wires, followed soon by electromagnetic communication through the "ether", speeded communication by more orders of magnitudeAll of these enhancements in communication speed allowed ideas (and goods and people) to interact more and more easily over longer and longer distances. None of this development, however, superseded the older means of communication face-to-face.

We may guess that a similar sequence of developing communicative capability may well have occurred in the evolution of living organisms. In Chapter II.2, we used the idea of the primeval "soup" as the environment of the early autocatalytic processes. In such a soup, reaction rates depend on the concentrations of the reactants and catalysts. Concentration is a local property. To take advantage in one place of a complex being created in another place, the complex product would have had to diffuse between the places, a slow communication process.

In Section II.3.3 we proposed (as had Oparin and many others) that the space for diffusion might have been delimited by a permeable membrane, restricting the distances over which communication between the reaction types in an autocatalytic loop had to diffuse, thereby speeding the loop production of complex product.

The concept of tubular membranes was introduced as a means of protecting the autocatalytic environment from moderate physical disturbances such as the waves of the open ocean, but it has another effect. It limits diffusion to one dimension rather than three, maintaining the concentration of the diffusing product as it progresses from where it is produced to where it is used in the autocatalytic loop. The tube speeds communication within the loop better than does the encapsulation of the products in a bubble or vesicle. But it is not fast enough for a mobile organism of any size to do more than leisurely float along taking advantage of or avoiding slowly varying properties of its environment. A new invention was required — electrical communication.

One may presume that the rhythm of waves might produce pumping actions that affected the diffusion of reaction products and speeded communication in systems that were appropriately configured, but no matter how complex a mesh of pumping cells and connecting tubes might become, the speed of processes was severely limited by distance. The waste products produced by reactions within the autocatalytic loops inside a cell might be food or catalysts for loops in another cell, creating a homeostatic loop that might or might not have the asymmetry that distinguishes a control loop from homeostatic loops in general.

If the cells floated freely through the soup, such control loops would act at speeds limited by diffusion and distance, and would be, like the "neural bundle" and "neural current" concepts used by Powers, generalized. A cell of type A would diffuse its waste to all cells of type B in its neighbourhood, and they would in turn diffuse their waste to all cells of type A to complete the loop. Any one type A cell would receive input from all the nearby type B cells, not distinguishing one from another. How strong the feedback loop, and hence how active the two autocatalytic loops, would depend on the average proximity of A and B cells. Autocatalytic loops that worked by communicating materials among cells would be effective only in small clusters of cells.

Going back to our cartoonish e-coli of Section II.3.3, type A and type B e-coli would tend to aggregate or clump together — and the cells in the middle of the clump would starve for lack of other food, catalysts, and energy throughput. Cells much simpler than the cartoonish e-coli would also tend to clump together if one provides food for another, if only because if an A cell provides food for a B cell, a B cell that happens to be near an A cell will survive better than one that has no A cells in its neighbourhood, as in our Archaeal cell and bacterium that evolved together to become a eukaryote.

Without going through all the stages, we may simply recognize that the whole autocatalytic process necessarily spawns many more types than just A and B, no matter what level of product we are talking about, whether it be molecules, cells, or neighbourhoods within cities. The better the communication between reacting units and their catalysts the more rapidly the complex product will be produced, and the higher the steady-state concentration will be. Since the catalysts are the complex products of other reactions, the better the communication between reaction types, the more productive the autocatalytic loop.

No autocatalytic loop purposefully increases its rate of production, but just as with the free-floating cells that will tend to clump in the soup, so also autocatalytic loops that produce more product will be less sensitive to outside disturbances that might tend to inhibit (anti-catalyze) any of its reactions. Natural selection therefore favours rapid communication in the world of autocatalysis, and that implies that autocatalytic loops will tend to become functionally shorter, and would tend to cluster their reactions in a space physically bounded by a permeable membrane

Diffusion through tubes is faster than through a three-dimensional "soup", especially if the tubes are subject to a pumping action. We should therefore expect that autocatalytic loops containing cells would come to be threaded with communicative diffusion tubes that connect cells having mutually compatible food and waste requirements for their internal autocatalytic processes. These tubes would tend to branch out many ways so that tubes from A type cells would contact B type cells, and vice-versa (though the return path might well be by way of C-type, and D-type cells around an autocatalytic loop of cells.

The functional reason for expecting a complex of threaded tubes to develop is the same as the reason ancient street patterns in many villages threaded their way among the houses so that pedestrians such as merchants could pass from one neighbourhood to another, rather than passing over the house roofs in a diffuse manner as was the case in the very early Anatolian town of Çatal Höyük 9000 years ago.

And so we come to Alessandro Volta and the voltaic pile or battery (Volta, 1800). Volta discovered that when certain different metals come into contact with a brine-soaked pad between them, a current (manifest in a spark) would flow in a wire that connected the two metal pieces. From this came all the discoveries about or that used electricity until Faraday's invention of the dynamo began to be used in the 1870s.

How does this relate to the foregoing? Electrical signals can be used to transmit what previously required face-to-face presence, either for talking or for transmission of written material. The spark of a telegraph signal or later of Hertz's radio transmission conveyed information. At some point in evolution, some reaction products or the products of their decay were ionized, some positively, some negatively. If one kind was produced in a reaction on one side of a membrane, while the other was produced on the other side, a voltage difference would exist across the membrane.

Liu et al. (2020) even found that a protein membrane laid on a conductive surface can create a sustained voltage and can supply continuous electrical current by its interaction with air of normal humidity. We might expect evolution to have found suitable membranes to do the same thing in an environment of a mix of different chemical liquids to create sustained voltage differences across the membranes, as we now see to be the case in biological cells such as nerve cells.

Voltage differences usually result in electric currents that dissipate the difference if the source of the difference is not renewed, as it is in a voltaic pile. We are not discussing autocatalytic loops or networks, which do renew their products so long as they have food and an energy supply. The voltage would be unlikely to build up if the membrane were conductive, so we assume that it is at best permeable to electricity as it is to material products — by way of specialized apertures.

Electrical currents and voltages influence chemical reactions, and the reactions that participate in autocatalytic processes are no exception. The Kauffman probability argument applies here, just as it did in the formation of autocatalytic loops. If it can happen that some process is speeded by a current produced by a voltage difference between some product of its autocatalytic network and a product of a different autocatalytic loop, then the first loop will survive better if it locates a "specialized aperture" or terminal near the other loop's region of ionization, and provides a conductive path for current to move between the production place of positive ions back by a different route to the reaction that creates negative ions.

It does not take too much imagination to see that we are talking here about the substitution of electrical current for the transmission of material waste among mutually supportive autocatalytic networks. When two networks intercommunicate, as we saw in Chapter 12, they may merge into a single network, or they may become parts of a higher-level loop of loops. Either way, we are talking about nerves when we talk about tubes that propagate voltage differences to many different branch terminals — synapses. Furthermore, although the words were not used, we have described both Hebbian and anti-Hebbian synaptic modification, the latter being an aspect of autocatalytic loops that are homeostatic, as most must be in a stably functioning organism.

Is this a surprise? It should not be. All we have said, in a very long and roundabout way, is that electrical communication is very likely to occur when any kind of repeated process generates ionization. Electrical communication is orders of magnitude faster than is material diffusion, so electrical communication permits autocatalytic loops to have vastly enhanced physical extent as compared to loops that do not use electrical communication. It also allows for control loops that react quickly to changes in external variables whose effects permeate the divisive membrane, and therefore it permits the existence and survival of mobile rather than sessile organisms. What plants use nerves in the way that a simple mobile organism does?

Here we leave consideration of individual organisms that have tightly interconnected systems, and begin our discussion of organisms (mainly humans) that must communicate through their shared external environment. We must consider the role of the observer, since for any individual organism the Universe consists only of what can be perceived. This seems like a simple proposition, but from just this proposition Einstein arrived at his revolutionary Theory of General Relativity and the famous formula $e=mc^2$.

By following the implications of doing as Powers often advised, thinking from the viewpoint of the controller, we soon grow the scope of our enquiry, metaphorically following Alice down the rabbit-hole into a world that initially seems very complex. It turns out that in this world Perceptual Control Theory makes an unexpected kind of sense, simplifying our understanding of much that has seemed unreasonably complicated. In the process of the analysis PCT suggests, as PCT has a habit of doing, that the conventional approach to many topics is likely to give the wrong answers or no useful answers at all. We follow that thought beginning in Part 6 after working through the metaphorical jungle to the clarity of the other side.

II.2.10 The Environment as Infrastructure

We would seem here to be at an impasse, because the first organism, before p-coli, could be not much more complex than our AnBn pseudo-molecules. But we ourselves are much more complex. The way out of the impasse is to note that the “soup” in which the initial As and Bs floated was not a uniform mixture, but contained everything that was required to allow an original A and an original B to coalesce into an AB. The soup was the environment and its properties constituted the infrastructure that allowed the coalescence to occur.

That same environmental infrastructure was sufficient to create and maintain the cyclic networks of relationships, the autocatalytic loops. Those networks of relationships in one way simplified the soup by reducing the uncertainties of the relationships between dyads when the members coalesced, and in another way increased the entropy associated with the ever increasing possibility space of more and more different dyads and larger complexes.

More networks meant more catalyzed product types that would be available to form different species of homeostatic networks. Only when higher-level networks that have those different species of free-standing networks as elemental units begin to occur do we get the possibility of mechanisms for the template-based construction of more of the basic homeostatic networks, as opposed to their self-organized structuring from the elementary units initially available. Template-based networks need not be restricted to those elementary units as their reacting components, but can also include any complexes so far produced.

At no time was the Second Law of Thermodynamics violated in the soup as a whole, though it was very much violated locally every time a new combination was synthesized out of the elements that floated in it. The infrastructure initially was the entire soup, or a sufficient part of it to provide a high probability that the elements involved in a catalyzed reaction could occasionally be found sufficiently near each other to participate in such a reaction. This, I think, is what Prigogine realized intuitively in 1947 and fleshed out into a full theory of self-organization in the subsequent decades.

If the infrastructure for the autocatalytic loop network was the soup itself, by the time p-coli had emerged, much of the infrastructure had been encapsulated inside the cell membrane. In the soup without catalysis, thermodynamic equilibrium is maintained because the reaction A+B→AB is balanced by the inverse decay reaction AB→A+B, the two reaction rates sustaining a constant ratio of the low-entropy complex to the higher entropy separated units.

Functionally, this is identical to the effect of the leaky integrator in the output function of the "standard" PCT control loop. In a biological control loop it might also be the mechanism. The integration rate is the rate at which the forward reaction produces the AB complex from the A and B ingredients, while the leak rate is the breakdown probability per second of the AB complex. The catalyst supplies the equivalent of the integrator gain rate multiplier. In the presence of the catalyst, a considerable excess of AB will exist, and that takes energy to sustain. As the McClelland quote says, maintenance takes work.

To do work takes energy, and p-coli needs energy to maintain its internal structure. This energy could be obtained by the chemical decomposition of the kind of complex chemicals built by the descendants of the process that built the autocatalytic networks, that is, by digesting food. To digest food requires getting food. A precursor of p-coli might have simply floated in the soup. It might have contained internal autocatalytic networks that its membrane held separate from the soup at a higher internal energy density than the surrounding soup. The internal autocatalytic networks interacted with the autocatalytic networks through its membrane, accepting some constituent chemical complexes from the soup and excreting waste elemental components in a spatially concentrated version of the AB⇄A+B two-way reaction.

The membrane would not be a prerequisite for the process to occur, but it drastically changed the rate at which it would occur, by concentrating the chemical ingredients in the small space enclosed by the membrane. For the larger and more complex combinations, the probability of their occurrence in the soup might have been so low as to expect one or two instances in a few million years, whereas inside the membrane that rate might be one every few milliseconds.

The decay rate per instance, though, would remain unchanged, so that a lone complex molecule that did occur in the soup a million years after another one had occurred might last a thousand years before decaying to its constituents. The same kind of molecule created as one of a thousand created every second would also have a thousand-year life span, so there might be very many of them available to participate in other reactions that would almost never occur in the initial soup.

As Prigogine and Lefever (1975) note, any cell must be complex (heterogeneous) enough to maintain a template of a correct internal structure, and a mechanism for correcting the structure when the inevitable damage caused by external events began to occur. To create the repairs would require acquisition of the materials to remake the broken structure, which would be done by ingesting food that contained them, so that they could be extracted as the food was digested to produce energy. That template must have been sufficient to create complete new copies of itself.

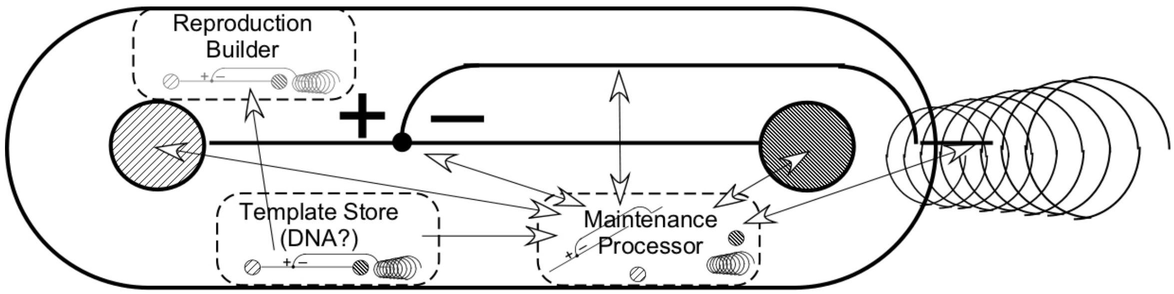

Figure II.3.2 Template-based maintenance and reproduction modules in an p-coli or di-coli. Each module is shown as informationally permeable. The template provides reference values for the perceptions controlled by the maintenance processor module and the reproduction builder. This p-coli is informationally efficient, compared to its ancestor.

Figure II.3.2 suggests how this "template and processor" permeable modularization might look functionally in some descendant of our initial p-coli. The original p-coli had available only its own structure as a template, remaking each part individually, but this descendant has a memory that contains instructions for how the sub-assemblies are made and how they are combined to create copies of itself, or to maintain its own structure against decay. From one point of view, these are reference perceptions, from another, they are programs and sub-programs or "objects" in object-oriented programming. This informationally efficient p-coli is appreciably more complex than its ancestral version, but it should both need less energy to maintain and survive longer to make more descendants.

We argued that if structures recur, their description may require less information if encapsulated into two distinct processes: firstly, a description of the structure, and secondly, a description of when to invoke the structural description and actually make an instance of the structure. In Section 10.7 we introduced a stupid furniture-maker we called Ted, who was asked to make a chair by fitting together four legs, a seat, and a back, all of which were structures made by someone else. Ted did not know what a chair looks like, and so fitted the pieces together at random. We computed that the probability that Ted's production would be a chair was about one in 23,040.

Suppose Ted had been told "Take the seat and put the legs on that way round one at a time on the same side of the seat, and then put the back on". He would have made a good chair the first time. Furthermore, if Ted, who had no memory, was asked to make a second chair, he would have had to go through the whole business again, being told once more to take the seat and add the legs, etc. Compare Ted's memory-less procedure with what he might do if he had a memory of the earlier instructions, or had a picture of what a finished chair should look like. He would make a chair after being give the simple instruction "Make a chair".

Informationally, a memory or a template picture requires no further information once it has been accepted (disregarding maintenance to correct decay). All that is required for Ted to produce another chair is the instruction to make one, if the ingredients of leg, seat, and back, are available. The template is part of Ted's supporting infrastructure, just as is the storeroom where already formed structures — legs, seat and back — are concentrated. With a good template and a well-stocked storeroom, Ted can produce chairs as fast as he can put them together, 23040 times faster than he could do it without a template and without instructions beyond "Make a chair from these parts."

Ted can't keep making chairs unless he gets food. It takes energy to put the pieces together, if only the mechanical energy of moving them into their places. It took energy for someone to make the pieces, and all that energy has to come from somewhere and go somewhere. And the completion of the chair has reduced the entropy of the set of parts. These things combine, as the energy exported in the form of heat is the highest entropy way in which the organism (Ted, perhaps) can interact with its structure. As we have said before, every control system is a refrigerator. Ted's completed chair is "cooler" than the separated parts, not in thermal temperature, but in the information required to specify where all the parts are with respect to each other.

Any structure as complex as a bacterium, or even a virus, could use the surrounding "soup" infrastructure both to sustain itself and to create copies of itself. But it would take less energy if it had a template and a "stockroom" of parts obtained through its food. Not all its copies would be exact, and the mistakes in a complex structure would usually be damaging.

Suppose Ted made a mistake and added two more leg pieces to his chair to make a rudimentary kind of armchair. Maybe another apprentice, "Sam", sees Ted's proto-armchair in a reject pile, likes it, and makes more of them. If people buy them, his boss will ask Sam to make more, and perhaps stop paying Ted, who continues making simple chairs despite occasionally making a mistake. Likewise in the evolution of the cell, what works is what survives. It is not hard to follow the evolution of simple floater cells in a soup infrastructure to the basic p-coli that is moved through the soup by its motor, and contains what is actually a primitive control loop that acts on and in conjunction with the infrastructural soup to move preferentially to places with more food. The p-coli control loop does nothing for the p-coli except allow it to use the environmental infrastructure to get the necessities for maintaining (and replicating) its internal structure.

Powers called the mechanism by which the control hierarchy is structured to keep the intrinsic variables healthy "reorganization" (Chapter 10) , but in his writing he emphasized the reverse relationship, that "reorganization" enables ever improving abilities to control perceptions of the environment. Now, looking at reorganization through Prigoginic glasses, we see a different third side of exactly the same process. We will take a quite different view in the next three Sections, as we discuss a novel measure called "rattling". More complex organisms evolved better ways of maintaining their own structure by developing ever more complex relationships with their environmental and internal infrastructures.

Internally, this is called reorganization, externally it is called control of perception. Both involve reduction and export of entropy. Figure 11.3 (reproduced here as Figure II.3.3) showed how it is possible for the mechanism of reorganization to use the same structure as does a control loop that acts through the external environment. The foregoing discussion suggests how using duplicated forms might save entropic costs by reducing the entropy of the template required for structural maintenance and reproduction.

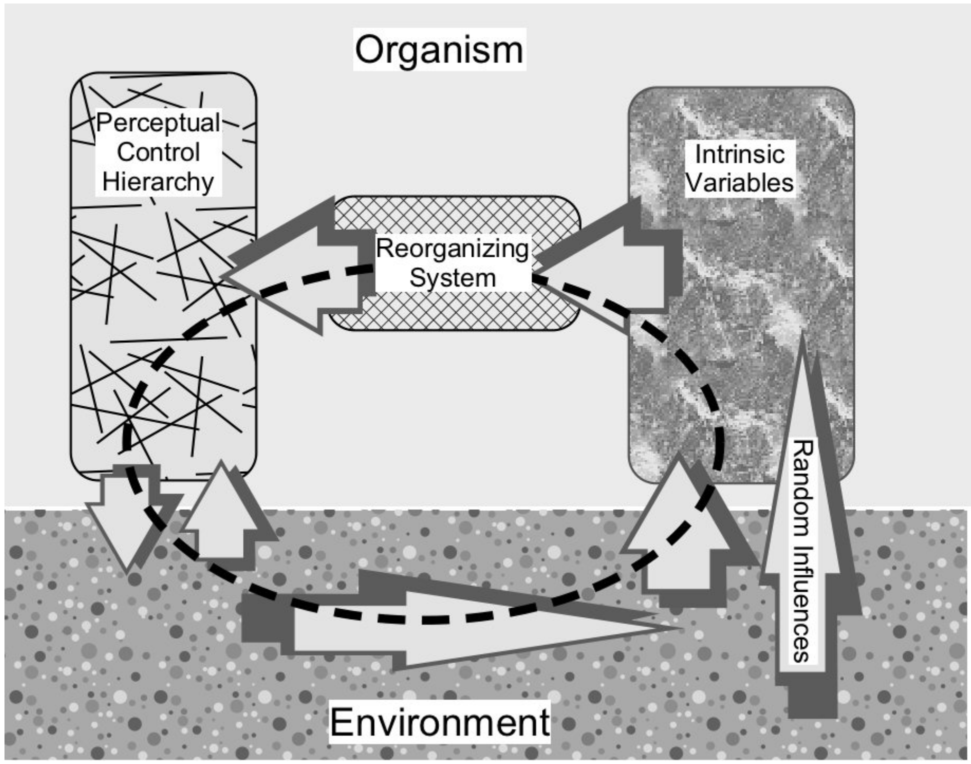

Figure II.3.3 (Figure 11.3 repeated) The relation between the intrinsic variable system and the control hierarchy seen as a control loop organized to stabilize the intrinsic variable structure. The reorganizing system is the output function of the loop. The output is the set of effects reorganization has on the perceptual control hierarchy.

This book is not about microbes or individual organisms. Soon we will begin to concentrate on the main topic of the book, how organisms work together socially if each controls only its own perceptions so as to get what it wants. Consideration, however, of some fundamental social processes at the micro-level may help in understanding how those same processes work when we as entire individuals are the elements in a "social soup" that allows us to combine and divide in various ways. In what follows, we may not explicitly refer back to the details of what has been discussed, but maybe in Volumes III and IV you will make better sense of the social world by imagining how these same processes apply.

The takeaways are:

- all organisms are structure (and so are networks of social relationships)
- to construct and maintain structure takes work, which requires energy
- the environment both is used to acquire energy to maintain and produce structure and is capable of destroying structure
- structure is most easily maintained by encapsulation that admits specialized use of the infrastructure
- encapsulation requires the encapsulated structure to use control through the external infrastructure while the form of the control is maintained or modified through the internal infrastructure

- in a sufficiently complex structure there will be substructures that reduce the entropic costs of the structure, such as classes of (templates for) cells within an organ that are all the same in some class, but differ among classes.
- any membrane-encapsulated entity at any level may take the place of an A or a B in a reaction A +B→AB, in which the AB has properties and potentials that are not available to A or B alone or acting independently without mutual interaction.

Enough about the microbes, although we should always remember that to our own microbiome we are merely infrastructure. Our microbiome will act to use and to influence its infrastructure (us) to its advantage, just as we act to use and influence ours (society).

Chapter II.3. Self Organization and Maintenance

In this chapter we begin to look at the whole form of a society — an organization — from a quite different viewpoint, whether the society is of control loops within an individual or one of which each unit is a whole living control systems. This different viewpoint is one in which specific applications of PCT to particular control systems are invisible, but the physical energy and work involved in building and maintaining structures is prominent. These points are well made by McClelland (LCS IV).

> *The kinds of activities described as work in everyday language are activities that create stable feedback paths in a shared environment for the benefit of other people. The word* ["work". MMT] *is also commonly used to refer to the kinds of activities that maintain these feedback paths in place. Thus, work activities produce some kind of environmental stabilization, the creation of some atenfel, molenfel, or molenex, which can then be used in controlling other perceptions. Manual workers create stable feedback paths by manipulating physical objects; they build things, make things, and clean things up. Agricultural workers produce fields of crops and confinements full of animals to be used as food. Transportation workers move truckloads of products from factories to stores, where sales workers make those products available to customers in exchanges with predictably structured protocols. Service workers manipulate and stabilize the immediate physical environments of individuals, including their dwellings and even their physical bodies, as barbers and hairdressers do. Healthcare workers attempt to stabilize the physiological functioning of people's bodies. Educators strive to turn out classes of graduates with predictable abilities and skills, people who can then be hired to put their skills to work creating various kinds of feedback paths for others. Government workers maintain stability and order for the community in a wide variety of ways, from removing trash to providing and enforcing laws designed to regulate commercial transactions and maintain public order, and thus preventing large disturbances that would make control of other perceptions difficult.*
>
> *The purposes of any given social structure are reflected in the work done by its members, that is, the ways they seek to stabilize some portion of their shared environment. Thus, we can classify social structures by the kind of work their members do: for example, families, ideally at least, stabilize a home environment for family members; schools aim provide stable flows of individuals with the tools to take action in predictable ways; businesses provide people with goods—objects that can be used as feedback paths—and services—routine actions that serve as feedback paths for controlling the perceptions of those who receive the services; and governmental structures are intended to prevent the kinds of disturbances to a shared environment that would make the work of other social structures more difficult.*
>
> *Even workers whose work seems somewhat abstract must produce physically perceptible stabilities, which can then be used as feedback paths for controlling lower- level perceptions essential for control of the higher-level, more abstract perceptions that provide the ostensible objectives of their work. Administrators and business executives create feedback paths by organizing the routine activities of others into predictable and efficient patterns for getting the work of an organization done. Knowledge workers put words on paper or images on electronic screens in order to send symbolic messages to others, thus facilitating their readers' control of higher-level perceptions. Entertainers offer their performances hoping to attract audiences, who will then use the performances as feedback paths for controlling perceptions of excitement or amusement. In every case, the creation of some perceptible product in the form of stabilized portions of the physical environment or stabilized patterns of human action—in other words, atenfels, molenfels, or some combination of them—*

provides the empirical evidence that work has been done. I argued above that these types of stabilities form the material and behavioral bases of social structures, and thus by producing these physical and behavioral stabilities people contribute to the overall stability of the social structures to which they belong.

In some kinds of work, people maintain feedback paths rather than creating them. People doing this work take the existence of certain feedback paths as perceptions to be controlled and then seek to protect them against the ongoing effects of disturbances. The janitor cleaning a building, the systems engineer fixing software bugs, the emergency responder driving an ambulance, or the baby's caretaker changing a diaper, all work to maintain feedback paths for others. Thus, the feedback paths in our shared environment depend on constant human attention and effort to do the work necessary to keep them stable. Without continual work, a humanly structured environment begins to crumble over time, like ghost towns or ancient ruins. The environments that most people live in are filled with feedback paths, both physical objects and routine actions, that have been shaped and maintained by human work.

We follow the ideas developed in Volume 1 on information and uncertainty, and in the recent Chapters on autocatalytic and homeostatic loops and networks. Whereas those Chapters concentrated mostly on individual control loops, now we begin to concentrate more on networks, often networks of control loops, after some preliminaries on the often misunderstood concept of entropy, and its relation to energy.

How does energy come to be an important consideration? We know as a historical fact that to sustain a "First World" lifestyle has required a large and rapid expenditure of the energy that plants slowly acquired from the sun over many millions of years, energy that has been stored in the earth in a form now manifest in coal, oil, and natural gas (mostly methane).

Is this rapid use of energy simply happenstance or is it inherent in the complexity of the urbanized world or of the internal structure of a living thing? How does such a complex system come to be in the first place, and how is it stabilized? What does "work" mean? No individual controller controlled for increasing the rate of energy usage needed to support this complex society, but to take advantage of the atenfels provided by society the average person needs more and more energy per time unit as the years go by. Why? We base much of this discussion on the insights of Ilya Prigogine and his colleagues about self-organized systems.

As early as 1947, Prigogine recognized that although the Second Law of Thermodynamics says that entropy always increases, this Law applies only to closed systems, systems that have no contact with anything outside some delimited space. In an open system, one that can accept energy and can output that energy in a different form to its environment, the Second Law does not apply. An open System with a through flow of energy can export entropy to its environment. Every refrigerator does this, as does every control loop (Section II.1.3).

Prigogine considered how complex systems might be self-organized and evolve. He had a precursor, as do most important insights. Lewis Fry Richardson (1922), studying the self-organization of complex patterns in the atmosphere, had observed in an often repeated (and often varied) ditty:

Big whirls have little whirls
that feed on their velocity,
and little whirls have lesser whirls
and so on to viscosity.

Richardson's whirls (often transcribed as "whorls") are initially created by the energy of a stream flow in any fluid. They occur when something such as a rock in a stream breaks the symmetry of the flow and creates a velocity shear in its neighbourhood. A velocity shear curves the trajectory of packets of the fluid, creating an eddy that has an opposed velocity shear at its edges.

Prigogine generalized and analyzed the phenomenon. The "big whirl" eddy might be the founding component of a self-organized structure, but its circulation in the stream would set up new shear stresses against the surrounding stream flow, inducing subordinate eddies which themselves had smaller eddies, all of which might maintain their stability for substantial lengths of time. The complexity of the stable eddy structure evolved from the dissipation through the smaller whirls of the energy that the stream flow continually feeds into the progenitor "big whirl". At the smallest scale, where small numbers of molecules are involved, their interactions become more and more random, dissipating the energy into waste heat. We will use Richardson's ditty when we later introduce "Rattling" (Chapter II.5).

II.3.1 Entropy and energy

Humpty Dumpty sat on a wall,
Humpty Dumpty had a great fall.
All the king's horses and all the king's men
Couldn't put Humpty together again.

(Old English Rhyme)

Whether or not Humpty Dumpty was a egg, as Lewis Carroll made him, the "great fall" obviously broke him into pieces that all the energy that the king's horses and men put into the effort of putting him back as he was failed to restore his structure. It didn't take much energy to break him (increasing his entropy), but to remake him (reduce his entropy to its former value) would have taken a lot, and that imbalance is the topic of this section.

Probably most people do not give entropy a second thought in their daily lives, if they have even heard the word. Maybe they should. We experience its effects every day. Drop a fine porcelain cup onto a tile floor. Was it easier to break it than it was to make it or than it will be to repair it so that one could ever tell it had been broken? Does it take more work to clean up all the pieces of the shattered cup than to collect the big, easily seen bits, leaving the hidden slivers and crumbs of porcelain until spring-cleaning time?

The examples illustrate one facet of entropy. It takes more work to build than to break, and more to repair than to break. We illustrated this numerically in the simple "markers on a checkerboard" Universe of Chapter 10. Things do break, and not always because of one's clumsiness. Things also decay. Even mountains erode away and become plains, given enough time. The skyscrapers of Manhattan are built on the roots of enormous long-gone mountains, and the skyscrapers will themselves collapse in a few generations if nobody works to maintain them.

In the introduction to this Chapter an extended quote by McClelland points out how much work it takes to maintain the social infrastructure. McClelland includes in this the abstract infrastructure of management and planning, laws and regulations. A law that is never enforced is no law, and it takes work by police, lawyers, judges, and prison authorities, among others, to enforce it. If that work is more than the result seems to warrant, the law will not be enforced, and eventually will decay into oblivion without being removed from statute books. Entropy is at the root of all these phenomena.

The meaning of the word "entropy" is covered by a multitude of everyday words, such as order versus disorder, structure versus decay, precision versus uncertainty, understanding versus bewilderment. In all of these oppositions, the underlying idea is that what is precise, ordered, neat and tidy and understandable is of low entropy, and their opposites are high entropy. Entropy, or at least the idea behind it, is perhaps not such a mystery, after all.[6]

If there is a mystery, it may be in the relationship between entropy and work and energy. That relationship underlies what much of this Chapter is about. I don't intend to go into the mathematical relationships much, if at all. It is adequate to keep in mind the notion that a tornado will take a house from a neat and tidy state to a pile of rubble, but will not do the reverse, that it takes a lot more work to become a professional pianist, golfer, or tennis player than to become a good amateur, that we need a lot of non-muscular energy to maintain a First World life-style, or that perfection is the enemy of the good because of the work involved in that last infinitely long step from excellence to perfection.

"Energy", however simple the concept seems, may actually not be quite so simple. It is an everyday term, but it is usually not described or defined. More commonly energy is discussed as something one can stockpile in the way one can stockpile carpets or bricks for later use. Oil companies may call themselves "energy companies" because that makes them sound benign, but they do not make energy. Nobody does. Energy cannot be created or destroyed, but it can change its form. It can be in heat, chemical structure, the height of an object above a floor. What power companies do is perform some such transformation, from oil in the ground to oil above ground for burning, from the energy in sunlight or from sunlight turned to the kinetic energy of wind and thence to electrical power.

Physicists find it convenient sometimes to treat the total energy of the Universe as zero, because that eases many computations, but the baseline number makes no difference to the effects. When physicists talk about energy (and entropy) they sometimes refer to an "isolated system", by which they mean a purely conceptual system into which no energy can come, nor can any leave. In an isolated system, energy can only change form, and the famous (or infamous) Second Law of Thermodynamics holds true. In an isolated system, and only in an isolated system, any change in the *form* of the total system energy is accompanied by an increase in the total entropy of the system[7].

Isolated systems, in theory, tend toward disorder. If any exist, we could not know about them, because they would change nothing in the Universe around them. Real systems that we can observe are never isolated, by definition. If they were, either we would be inside the box, in which case we could not observe it completely because our observation would have to include our own structure, which the observation would change, or we would be outside the box, in which case we could not observe it because no energy would leave the box to let us know what was inside. Even a black hole is not so secretive!

The idea that isolated systems tend toward disorder has a sound probabilistic foundation, but it can never be tested by experiment. All real systems can accept and expel energy. Whether they simultaneously increase or decrease their entropy depends on the form of the energy accepted and expelled. An untidy room can be made neat and tidy, but the person doing it has ingested low entropy chemical structures in the form of food, and has expelled waste as well as an appreciable amount of energy in the form of heat, which is the least structured, maximum entropy, form of energy. That person could not make the room neat and tidy when starving in a hot atmosphere — or in a cold atmosphere, for that matter.

---

6. The concept, using as an example the motions of the atoms of an ideal gas at a given temperature, is clearly illustrated at https://en.wikipedia.org/wiki/File:Translational_motion.gif. (Retrieved 21.09.04)

7. Technically, this is not strictly true, because there is always a possibility that completely by chance all the molecules in a small sub-region of the space might wind up moving in the same direction at the same speed, a very low entropy state locally at very low energy, and very cold. But for a region of more than a few tens or hundreds of molecules, the probability of it happening within the life of the Universe is so near zero as to make no practical difference. If you see an example in your lifetime and tell someone, they are likely to say you were hallucinating.

Much of Chapter I.10 was devoted to a discussion of uncertainty, structure, and “good form” from a psychological point of view. Now is the time to link that discussion to the concept of self-organization and energy. In the next Chapter we will introduce another concept that is a cousin to entropy and energy, “rattling”. In Chapter I.10 we made extensive use of a simplified Universe of Possibility, a checkerboard on which each square might or might not be occupied by a marker. The entire set of possible arrangements of markers was described by specifying for each of the N squares whether it contained a marker.

For this toy Universe of Possibility, maximum uncertainty is reached when the probability of any one square containing a marker is 0.5, independently of what markers might be on other squares. That uncertainty is N bits, because there are 2^N possible arrangements. If you were to describe a particular arrangement to someone else using a voice phone, you could do no better than list in order the squares and say “yes” or “no” according to whether the next square had a marker or not.

At the other extreme of order, all the squares might be filled or they might be all blank. All you would need to say is a single “filled” or “empty”. One bit of information specifies the entire pattern of N squares, but only if your listener knows those to be the only possibilities. If instead the listener believed all squares to be filled with the same colour, but the markers might all be black or all white, red or yellow, with equal probability, then the same entire pattern would need two bits for its description. If markers came in 16 different shapes as well as the four colours, four more bits would be needed. And so on.

Between these extremes lie patterns with different degrees of order, uncertainty to the listener you are instructing, or, as we will now say, entropy. For example, you might describe the checkerboard with the usual black and white squares by saying that diagonally neighbouring squares are all the same colour while laterally and vertically neighbouring squares are always in opposite colours, and square 4,7 is white. Saying that takes a lot less information than saying for each square whether it is black or white.

Suppose that on the template checkerboard that you are trying to describe as succinctly as you can there is a vertical cross of seven squares in each direction, intersecting at their central square, and a circle of seven squares diameter lies directly above and touching the cross. And suppose this checkerboard is big enough that you cannot see the edges. If the person you are talking to knows how to lay out crosses and circles, perhaps what I just wrote would be sufficient for the replica to be close enough to your template that you would say they were “the same” when you compared them.

But if the person did not know how to make circles and crosses, you would either have to say something like “Choose a starting square and put a marker on it, then put markers on the next six squares above that one. Go back to the middle square that has markers and put markers on the next three squares either side. Above the highest square with a marker, … .” Or, if you expected later to be describing other templates that included crosses and circles you might start with a description of how a cross or a circle might be laid out, and then give instructions like those with which I started, which use the label to point to the instructions already known to your friend.

This latter approach segregates two conceptually different parts of the construction description or program. One, the subroutine, says “Here’s how to do X anytime you want”, while the other, the main program, says “Do X at this location now, let me know whether it worked, and come back for your next instruction”. We are producing an informationally efficient way of describing how to create a structure according to a template, and this efficient way requires parts of the description to be encapsulated separately from each other.

The encapsulation cannot be complete. Information flows in both directions through the membrane when the main program is performed (which is when you pick up the phone and call your friend to start the description). Moreover, the “main program” itself might be a subroutine in a membrane that encapsulates it within a program at another level. One can make an analogy, perhaps, between the subroutine as one gene among many and the main program as an epigenetic regulatory process that is itself described by and evoked by genes.

This kind of information-permeable encapsulation occurs in many places that may seem very different on their surface. Using the language of entropy and energy, the "Universe of Possibility" becomes an "entropy space", within which everything about a pattern or structure can be described. In the physical Universe, the space we can observe extends out through time and distance to the beginning of the Universe, some 13.8 billion light-years away in space and years in time. Every year, it expands in all directions by one more light-year, providing that much more entropy space to the observer. Anything outside that region, the visible or accessible Universe, does not contribute to the entropy we could, in principle, measure by describing every property of every item and every relationship in the Universe.

The boundary between visible and invisible parts of the Universe is not a permeable membrane. It is defined by the speed of light. Everything that happens too far away for light to have had time to bring us information about it is unknowable to us. Betelgeuse, about 500 light years away, might have gone supernova while Shakespear was writing his plays, but we will not know if it did until the latter half of this century. If it went supernova at that time, it will be the brightest object in the sky after the sun for a considerable while. We or our descendants will certainly know about it once it passes the boundary between the visible and the invisible part of the Universe.

But we will not know everything about it in an instant. It takes time for information to be communicated through a channel of finite capacity, as we pointed out when considering perceptual delay in Section 12.4. The amount of information we could, in principle, get about the Betelgeuse supernova would start at zero when the event crosses this space-time boundary, and then would grow more or less linearly with time. We will pursue this issue again later.

II.3.2 Entropy and Energy Mapping

We have mentioned the food supply from time to time, but now we turn our attention to it. Malthus's point that an exponential increase will at some point exceed a linear increase applies here just as much as it does to the human population of the world. It is a universal truth, no matter whether Malthus, not knowing of future advances in agricultural technology, had a correct estimate of when the crossover would occur.

At first sight, the initial soup, which contains only a limited number of elementary units, seems much simpler to describe than does an elaborated multi-level autocatalytic superloop, with its myriads of products of different complexities. One might think it should be in a low entropy state, because all one has to do to describe it is enumerate the components and their concentrations. But this is not the case, as Boltzmann (1877) pointed out. To describe the initial soup requires more, much, *much* more. The location and momentum, and possibly some of the electrochemical properties or equivalent of the elementary units, must be described for every individual unit. Descriptions *en-masse* will not suffice. The soup is in a maximum entropy state.

To reduce the entropy of the soup requires a through flow of energy that can be used to export entropy to some other environment, otherwise known as an input of energy that is used to cool some parts of the soup while heating other parts. By "cool" here, we mean creating structure where there was none. In the complementary control view, "cooling" the structure means maintaining its properties stable by keeping perceptions of those properties near their reference levels. Creating structure at one place always requires exporting entropy within an energy flow, and that energy is often in the form of heat.

Processes in the initial lifeless soup have one source of material from which to build their complex products, and two possible sources for an energy flow that allows the elementary material components to be combined. One of these is energy stored in the heat and the radioactive elements of the Earth itself, the other being the sun. In both cases, the outgoing energy flow that allows for entropy reduction is completed by radiation to outer space. Ultimately, all life on the Earth can use no more energy per second for survival than that provided by the sun, though for eons the internal heat and the remaining radioactive elements may supplement it.

Autocatalysis increases a Universe of Possibility within the greater environment by creating an increasing variety of types of possible complexes from older complexes. At the same time it cools the Universe of Possibility into a Universe of Actuality by creating very few of these potentially plausible complexes and exporting entropy, often in the form of heat, into the larger Universe of Everything. Figure II.3.1 schematizes the process, which is quite general in connection with control.

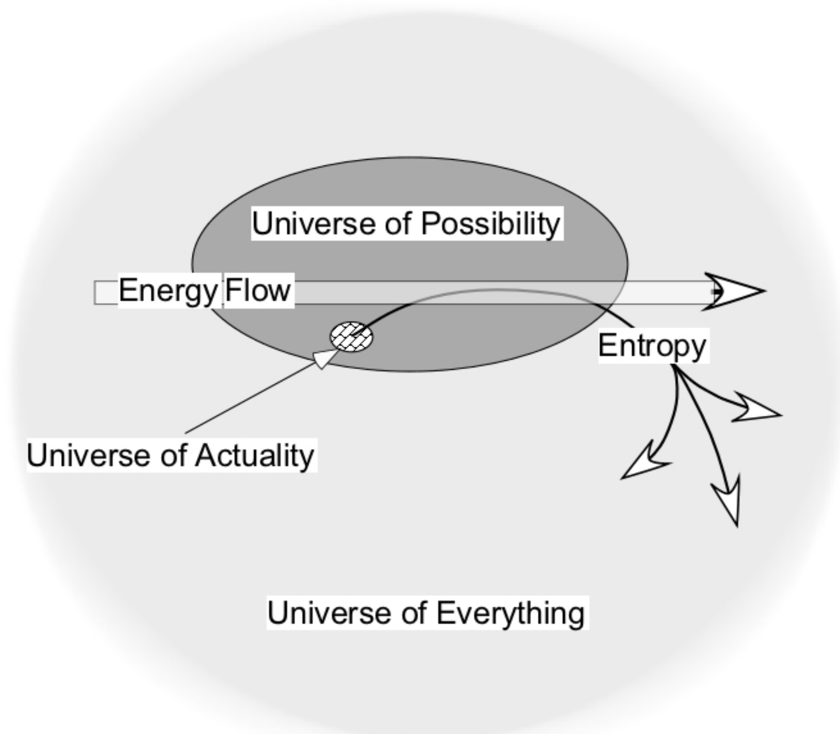

Figure II.3.1 The flow of entropy from the Universe of Actuality through the Universe of Possibility to the Universe of Everything. The three "Universes" are explained in the accompanying text.

To explain these three Universes, imagine that you have just moved into a new house, and everything is in boxes. You want to arrange your different sets of plates among the shelves and cupboards in your new kitchen. The Universe of Possibility is the set of different ways the crockery might be arranged, including leaving them in the boxes. But you use some energy in putting them into the places you think best (you control your perceptions of their locations). Wherever they are at any given moment is the current Universe of Actuality. The set of reference values for the locations of the different plates is the Universe of Actuality you want to produce, a very small subset of the University of Possibility, as we showed in Chapter 10, and it takes a fair amount of energy to refine the larger Universe into the smaller.

That energy, and with it the entropy that your perceptual control operations took from the Universe of Possibility, is dissipated into the larger Universe of Everything, which includes the air in the kitchen that might absorb the heat you create doing the work, not to mention everything you can see out of the kitchen window, the glasses, knives, and forks you have not yet placed, and the Sun and stars in the sky.

Control maintains the structured Universe of Actuality, preventing it from dissipating into the larger Universe of Possibility. People use the plates you carefully placed, and maybe don't put them back where you wanted them after they have been washed. Over time, if you don't do anything about it, your Universe of Actuality for the plate locations has expanded well beyond the refined state that cost you so much work to set up.

In a simple control loop, the Universe of Possibility is the range of values the perception might take without control, while the Universe of Actuality is the range it takes while being controlled. But with the creativity of Autocatalysis, the range is over the huge variety of complexes that might be brought into existence at appreciable concentrations, while the Universe of Actuality is the set of complexes that actually are built on the processes that participate in the loop. If the Autocatalytic loop is homeostatic, Universe of Actuality is further reduced because the range of values of their concentrations is also diminished from the possible to the actual in the same way as it is in the simple control loop.

With this in mind, we can address a third issue implicit in Figure II.3.2. Unlike the first two, this third issue is not an anomaly in the overlay of the control diagram on the autocatalysis diagram, but a question that arises in the mapping itself. Autocatalysis is important because it produces products — AB and CD in this example — that are more complex than the food elements — A, B, C, and D. These complex products are available as elementary units in further reactions that produce even more complex products (e.g Figure II.2.2). In control, what is the complex product analogous to the ABs and CDs?

The product of control is not a change in material. Controlling the value of some variable does not change the entity of which the variable is a property. Putting a glass on the table changes neither the glass nor the table. Nor does it create a single "table+glass" entity. What it does is create a stability, a reduction in the entropy of a portion of the local environment (Section 10.1 and Section 10.2). Because someone had a reference value for perceiving the glass to be on the table, it ceases to be equally probably on a shelf, in a cupboard, in the sink, but has become "on the table" with high probability and in any of those other places with low probability. The entropy of the Universe of Possibility that consists of relationships between places and the glass is much greater than the entropy of the Universe of Actuality if a perception of the glass location is controlled, but not if it is uncontrolled.

The effect of control, then, is the same as that of constructing a complex from simpler elements in that both reduce the entropy of the Universe of Possibility for the entities or properties concerned. That stability is available for use by other processes of the same kind. It may become an element of "infrastructure", which can be used in controlling other unspecified perceptions, as described in McClelland's quote. Just as AB can be a catalyst for a reaction C+D→CD, so a stability created by one control unit can be an enabler, a catalyst, for controlling a perception that could be controlled only with difficulty or not at all before that new stability existed.

Stability requires maintenance, the ongoing opposition to disturbances from the uncontrolled environment that is the essence of perceptual control. Just as the product AB of the reaction A+B→AB may decay away to its constituents A and B, so in the absence of continuous control or participation in a homeostatic loop will an environmental property drift away from a value it once had. And, as McClelland so rightly pointed out, it all needs ongoing work. Work needs an energy supply, a flow from a source to a sink in which the entropy extracted from the stabilized part of the environment can be deposited. It matters not whether the entropy is extracted by making a new chemical complex or by mechanically affecting a variable so that it takes on and maintains a particular value — or by simply thinking about maintaining a structure of thought in spite of disturbances.

Energy is used, transformed, in any localized reduction of entropy, in the environment or in the brain. But it is not lost, even locally. Some, but not all, is recovered if the stabilized property or the chemical complex decays. If you stop trying to hold onto a thought, you can more easily think about other things while the memory of the thought may decay and the thought be lost. In many biological processes, the energy recovery from the dissolution of structure is used to good effect, for example in the building and deconstruction of ATP to power control processes and homeostatic loops in various cells throughout the body. Mechanically, the release of energy is manifest in the controlled lifting and uncontrolled dropping of a ball that releases heat and possibly breaks things on impact. The effect is the same, whether the entropy was reduced by chemistry, by homeostatic processes of any scale, or by overt control. The functional mapping between autocatalysis, homeostasis and control is complete.

II.3.3 Is there a Latent Heat of Homeostasis?

When the kinetic energy of water molecules drops (we call that the water getting colder), there comes a point at which the water freezes. By this we mean that the kinetic energy of the relative motion of the water molecules drops to near zero as the molecules take up a regular crystalline structure. That kinetic energy (heat) must go somewhere outside the collection of water molecules. It goes into the environment, and until the water has entirely frozen the temperature cannot drop further. This dissipated heat is the latent heat of freezing (or of melting, in the reverse process). Maybe it is strange to think of the concept of "latent heat" as possibly applying to the formation and dissolution of homeostatic loops, but it should not be, after the preceding discussion of energy and entropy.

What does "latent heat" mean, other than as heat that must be supplied if a solid is to melt or a liquid to boil? Such changes of state are sometimes called "phase changes", and phase changes often involve the addition or removal of energy in order to pass from one phase to the other. The different states or phases have very different entropies although they differ very little in temperature. To oversimplify a little, but not too much, a gas consists of a bunch of elementary particle units, usually atoms or molecules, that move independently of each other, without correlations between the movements across pairs of the units. In a liquid, there are short-range and short-lived correlations among small clusters of units, while in a solid, those correlations are much wider ranging and long lasting. Starting from a frozen state and ending with a gas, each phase change implies a jump in the entropy, not a slow increase, as occurs when a solid, a liquid, or a gas gets warmer. Each phase change requires a finite amount of added energy, its latent heat.

In a liquid, small clusters of the units form and reform, with the whole cluster moving as a unit. These clusters are "structure" in the sense we have discussed. Their components move in a coordinated manner through their environment of other such clusters and independent elements, always with the possibility that elements or clusters will stick together momentarily before redissolving into independence. This tendency to "stickiness" gives the liquid its higher viscosity than the gas, in which the particles are moving too fast independently to have much chance of sticking, even momentarily. In the solid, the cluster structures join together so that they move in coordination as a unitary block. The structure encompasses all the units together.

Every bit of structure (using "bit" in the information theoretic sense or not, as you please) is a reduction in the entropy of the whole, which must be accomplished by exporting energy to a wider environment during the phase change of liquefaction or freezing. Conversely, when a solid melts or a liquid boils, it happens because heat has been applied, not only to raise the temperature, but also to increase the entropy of the mass by destroying the coherent structures.

These concepts translate almost directly into what we have described as happening in the elementary "soup" when an autocatalytic loop or network begins to form. Recall how the autocatalytic process begins with free-floating elementary units in the "soup", which occasionally react together to form "molecular" complexes. This level of organization or disorganization is like that of the gas just above the liquefaction temperature. The complexes move in a coordinated fashion until they decay into smaller parts or into their elementary constituents.

The next stage in building the autocatalytic loop is the beginning of catalytic clusters, in which some products of the reactions catalyze other reactions, increasing the concentration of novel complexes that might act as catalysts to other reactions. The more concentrated the "soup", the more likely it is that a potential catalyst will find itself in a position to actually catalyze a reaction before decaying, a fact that is directly analogous to the increase in the transition temperature of a gas to a liquid or vice-versa as the pressure increases.

These unstable complexes and their interacting reactions can act as nuclei for the next stage, solidification. At this stage, there are enough of them that the first autocatalytic loop comes into being, which locks the participating reactions into a steady production mode despite potential disturbances that would break up the individual clusters of reactions in the absence of the completed loop.

There is a significant difference between the lowering of entropy in the freezing of a solid and the lowering of entropy in the formation of an autocatalytic loop. In the freezing of the solid, heat is withdrawn and the motion of the atoms or molecules of the solid are reduced. This is done by means of a through energy flow that entrains energy and entropy withdrawn in a refrigerator-like action.

In the autocatalytic process, the entropy reduction goes into the creation of the structure of the reaction products, which may absorb energy rather than simply releasing energy to the environment. The same through flow of energy is required to export the entropy to the larger environment, as is the case for freezing a solid in an environment that is warmer than its freezing temperature, but some of that energy remains in the structure created. The structure is stressed, storing energy that will be released if and when the structure decays back into its components. Explosives work for this reason; the structure of the explosive molecule took up a lot of energy in its construction, and when the structure is destabilized sufficiently, it releases that energy when it falls apart. It is like the energy used to raise a tall building. If the building falls and breaks apart, it releases a lot of that energy back into the environment.

On the other hand, if two components tend to stick together without being pushed into fusing, like two magnets in which a north pole faces a south pole, or like an electron near a bare proton, creation of the complex sends energy into the environment rather than extracting it. For any structure that doesn't immediately fall apart once it is left unattended for an instant there is an energetic barrier against decay of the structure. If that barrier is overcome by energy introduced from the environment, then the structure may dissociate with the release of energy. Such a structure is produced in a reaction called "endothermic" and dissociates in an "exothermic" reaction. If an exothermic reaction releases enough energy to bring other molecules of the same structure over their energy barrier, we have a source of heat, possibly in the form of an explosion.

As we have described on several occasions, it takes a through flow of energy to maintain structure and keep the entropy of the structure low while increasing the entropy of the larger environment. We now mention that such process are necessarily endothermic, because if they were exothermic, they would form by simple association, like the magnet pair. Catalysis would not affect their production in any important way. When an autocatalytic loop begins to function, there is a sudden increase in the rate of production in endothermic processes of the complex structures formed by all the catalyzed reactions in the loop. To sustain such a loop takes a lot more steady energy flow to extract energy and entropy from the newly formed complexes than to keep producing the few complexes that form and decay in the free-floating soup.

The change from sporadic catalytic action that forms free-floating reaction clusters to an autocatalytic loop is sudden. The transition is a phase change in the chemical structure of the “soup”, just like the phase change that happens when a near-freezing liquid freezes. In both cases, the process extracts physical latent heat from the material. Since we are talking about a very similar structural “freezing” in both cases, it seems natural to call the energy required in forming the first loop the “latent heat of the formation of an autocatalytic loop”. To illustrate further the functional identity of the processes, we may note that when the autocatalytic loop is broken for some reason, the production rate of its products slows dramatically, and many of those that were created in the loop will decay into their components, increasing their entropy and releasing their stored energy into the environment just as do the molecules of an explosive that has been detonated.

If there is a latent heat of autocatalysis, it follows that there is a latent heat of homeostasis, a homeostatic loop being a form of autocatalytic loop. In the case of a frozen liquid, if the solid is simply placed in a warmer environment, it accepts energy and an increase in its entropy, and melts into a liquid. In the case of a homeostatic loop that breaks, what should we expect?

Whereas the frozen solid froze because energy was extracted from it, the homeostatic loop froze and stayed frozen because each reaction in it required a continuing energy flow to produce and maintain the low entropy structure of each individual product. That energy flow may well continue, but once the loop is broken, not only do the decaying products of the loop processes deliver their stored energy back into the environment as they decay, but also the energy flow is now not being used productively. There is indeed a latent heat of homeostatic loops, but that latent heat depends on the structure of each specific type of loop, just as the latent heat of freezing any liquid depends on the constituents of the liquid.

Clearly, this being a book about perceptual control and its application to various social issues generically labelled “Language, Culture, Power, and Politics”, the homeostasis of a loop of chemical reactions is not the true focus of our interest. Our analysis, however, provides examples that are more than just analogues of processes we expect to find operative in these areas of real interest; they are almost exact functional equivalents. We will carry this further in the next Chapter before we come full circle back to the social topics on which we focus in Volume 3.

II.3.4 The Nature of a Tensegrity Structure

Tensegrity is built on a dynamic balance between attractive and repulsive forces, attractive being manifest in the tensile “wires” of a physical structure, repulsive in its compressed rods. A tensegrity structure is an inherently low-rattling organization of its components because of the way it distributes imposed energy fluctuations among all its components. This would not happen if the components were isolated, each subject to all the varying variability of the stresses imposed on it from outside. Isolated components break more easily than they would as members of the tensegrity structure.

Tensegrity smooths the variation through the averaging effects of the Law of Large Numbers, the numbers being the number of rods or wires connected at a node, no matter where in the structure the stresses are imposed. The structure takes the strain; the individual components do not. The total energy absorbed by the structure is the same as would be absorbed by the affected component in isolation, but the variability of its variance as felt by that component is lower. The component is less rattled as a member of the tensegrity structure than it would be alone.

Is a tensegrity structure a solid, gas, or liquid? Its components are the pairwise attractive and repulsive forces between its elementary units, which we might as well call “atoms” while we address this question. In a gas, the atoms move independently, seldom encountering each other gently enough to form any kind of structure other than building individual atoms into individually moving molecules. In a liquid, the average encounter energy is lower, and structures tend to form, to be broken apart only by encounters more violent than average, and then reform. A tensegrity structure is analogous to neither a gas nor a liquid.

So is a tensegrity structure a solid? Earlier, we described two forms of solid, crystalline and glassy. They are distinguished by the relative mobility of their component atoms. In a crystal, each atom is held in place by its neighbours so that it cannot move through space. It can only vibrate in place. Under sufficient force, a crystal will break apart, no matter how slowly the force is applied, or in what direction. In a glassy material, Each atom is free to move, like a person trying to make way through a close-packed crowd. Glass is a liquid with extremely high viscosity. Under sustained force applied sufficiently slowly, a glass will flow like a liquid. It will break only if the force is applied too quickly for its atoms to flow.

Is a tensegrity structure crystalline? Apparently not. Is it glassy? Apparently not. Scarr (2014) asserted that the very molecules of which our physical structures such as bones and muscles are built are tensegrity structures, and I think we must say that control tensegrity structures such as the perceptual control hierarchy are equivalent to complex molecules each of which might be considered as something like a solid.

We might call the tensegrity-molecule a complex solid because its "atoms" (rod-ends or wire-ends) maintain a constant relationship with their neighbours, constrained by boundaries in their attractive or repulsive interactions. Perhaps a "soft particle" might also serve. Analogies of this kind, however, turn out not to be very useful. We have to treat a tensegrity structure as a form of matter requiring the new name, a "tensegrity". Its atoms are attracted to most of its neighbours, and repelled by only a few. Both are critical. The attractive forces enable the entire structure to cohere as a low-rattling form of organization despite the repulsive forces that would collapse if they did not exist within it. I suggest that a stable social structure or the structure of a stable perceptual control "hierarchy" will more probably be more closely analogous to a tensegrity than to any other phase of physical matter.

The transitions between solid and liquid, and between liquid and gas, involve the transfer of latent heat between the material and its environment while the material remains at a constant temperature. Is this true of a tensegrity disintegrating under the continuing application of heat? This question will be important if, as we will claim, social structures have tensegrity properties, because how societies and nations break under stress has a bearing on how to protect them from breakdown into what sometimes is called a "failed state".

The arguments above apply just as much to the development of bones and shells that offer structural support in a sometimes physically violent environment as they do to cells and tubes that structurally support the maintenance of high chemical concentrations in a low and variable concentration of the fluid environment. What, after all, is a bone but a tube that has an internal structure with the rigidity that permits it to do what is postulated above for a tube — harbour and protect the biochemical and cellular structures (e.g. stem cells, red blood cells) that permit the generation and regeneration of a variety of components that are essential to life?

Please do not believe for a moment that however tempting the ideas might be, the above is the way life actually began and evolved into the variety of vertebrate forms that have lived on this Earth. It is merely an example both of where thinking about the powers of PCT may lead, and as the kind of direction that might be worth pursuing when considering how life that depends on control of interior variables by acting on the environment could have "discovered" the value of control.

Nevertheless, the possible role of side-effect rattling in constructing organizations of proto-cells is a reasonable prelude to the more serious consideration of the importance of side-effects of control in creating organization in human societies. If you keep in the back of your mind a metaphor of the sometimes wave-washed lagoon when thinking of the irregular buffeting stresses ("rattling") that affect human and other societies from time to time, analogies to the tubes, cells, and bones may perhaps come to mind when we talk about social interactions in the last part of this book. According to Chvykov et al. (2021), these constructions are, and more often than not lead to, what they call "lower rattling" organizations of the possible ways components of any organization of components might have been arranged.

A further point to keep in mind is the concept of a permeable membrane, which also reduces the "rattling" of its interior. Membranes appear in different guises, not all of them physical. In the foregoing, we have treated a membrane in a rather abstract way, as providing on the one hand a shield that reduces the effect of some external variation on the states of variables inside the membrane, and on the other as a medium for desirable influences from outside to be represented inside and for states inside to influence states outside. Modular systems of any kind are similar, in that the cross-influences within the module are more dense within the module than from one module to another. Indeed, this is almost a definition of what constitutes a module.

These structures and energies are miniature models of much that we will see when we address the applications of PCT in social environments beginning in Volume IV, the side-effect interactions among individuals having much in common with the interactions among the components in the soup that lead to stable structures such as homeostatic loops and encapsulating them in selectively permeable membranes.

II.3.5 Entropy and Evolution

Cells are complex structures which must acquire energy and use it to maintain their structure. Then they must release some of it in the form of the heat that is inevitably generated in any natural physical or chemical process. The imaginary p-coli and di-coli need energy to drive their motors. They may also store some of their acquired energy as potential energy in new chemical structures, which they may use to grow their bodies or build their descendants (as a pregnant woman is said to be "eating for two").

The p-coli and di-coli family of Chapter II.3 get energy from the nutrition toward which they drive — their "food" — and they deposit most of that energy into their environmental soup by stirring it. Some is lost as heat because the organism heats up enough to allow for transfer to the surrounding soup, and the rest is ejected to their environment by producing more mass in the form of descendants.

If in total they deposited less energy than they got from their structured food, they would heat up indefinitely and eventually be incinerated. By expelling structured (low entropy) descendants they would not reduce their own entropy very much. They reduce their entropy by expelling high entropy heat and poorly structured waste material, increasing the entropy of the surrounding soup.

Lost in the first paragraph are the words "new chemical structures". Cells are at heart chemical structures, built, as Kauffman, following Prigogine, surmises, from autocatalytic processes that have created ever more complex homeostatic networks of chemical reactions. A homeostatic network keeps itself stable even though it may continually evolve to produce products of ever increasing complexity. By now, though, we can see why the ability for self-replication might exist as a byproduct of self-organization in a through energy flow (e.g. Prigogine, 1947; Prigogine, Nicolis and Babloyantz, 1972a, 1972b; Prigogine and Lefever, 1975).

The thermodynamic mechanics of the autocatalytic processes may be interesting, but more striking is an early observation by Prigogine (1947, quoted by Prigogine and Lefever 1975):

> *Ce n'est donc ni l'entropie, ni d'ailleurs aucun potential thermodynamique, qui permet de charatériser l'évolution d'un système, mais seulement la production d'entropie qui s'approche*

de sa valeur minimum.

Dans certain cas cette valeur minimum de la production de l'entropie ne peut être atteinte qu'en augmentant l'hétérogénéité, la complexité du système. Peut-être que cette évolution spontanée qui se manufeste alors vers les états à hétérogénéité plus grand pourra-t-elle donner une impulsion nouvelle à l'interprétation physico-chimique de l'évolution des êtres vivants.[8]

Prigogine and Lefever (1975) comment:

In this passage we have already the duality between creation of structure and maintenance of structure. The situation at present seems to be that these two aspects can be characterized by different behaviour of the entropy production. While the creation step leads to an increase in entropy, the steps of maintenance, on the contrary, seem to follow the theorem of minimum entropy production.

If we think back to the sections of Chapter I.11 on structure and uncertainty (entropy), we see that though this is intuitively correct, it is not necessarily self-consistent as written. If the added structure requires an increase in the number of components that might be related, the Universe of Possibility for the little local Universe is greatly increased, in the same way that it would increase by a factor of 10^{53} from 2^{49} to 2^{225} (176 bits) when a 7x7 checkerboard space is increased to a 15x15 space with squares identified only as occupied or empty. The actual information required to specify the new structure may at the same time increase very little, especially if the structural change involves the incorporation of motifs that existed in the ancestral structure. Depending on the degree of stability required, structural maintenance does indeed correspond to a minimum in the entropy generation rate, because any increase in entropy implies an unspecified variation in the structure.

But think of the implications of Prigogine’s 1947 statement in the context of our di-coli world (Section II.2.6). The An+Bn→AnBn reactions build structure, since one instance of the complex is clearly better structured than is one An and one Bn swimming alone somewhere in the soup. Compare the example of describing the cross and the circle. When your friend knows the circle to be above the cross, you would have to describe the location of only one of them, not both. The formation of the complex reduces the entropy of the pair, because the mutual relationship of the components is highly constrained within the complex as opposed to being completely uncertain while the circle and the cross, or the An and Bn components of the reaction were floating alone.

When a p-coli mutated into the successful di-coli of Figure II.3.4, it did so by adding a copy of an existing structure for which it already had the construction plans. It did not increase complexity by adding one element at a time such as an isolated motor unit or a comparison junction that connected nowhere. It increased the complexity of the resulting di-coli relative to its own complexity by making a trivial mistake, perhaps forgetting it has already made one sensor-motor array and then making another, perhaps misreading “one” as “two”, the same kind of error as we showed for the one-marker movements in Figure I.12.10.

What I am saying here is that the p-coli contains an *infrastructure* that is available for creating the new

---

8. *It is therefore neither entropy nor, for that matter, any other thermodynamic potential that allow the characterization of the evolution of a system, but only the rate at which entropy production approaches a minimum value.*

In certain situations this minimum entropy production rate can only be attained by increasing heterogeneity, the complexity of the system. Perhaps this spontaneous evolution, which shows itself in the states of greatest heterogeneity, may offer a new way of interpreting the physics and chemistry of the evolution of living beings. [*My translation*] complexity inherent in the change to di-coli. And that is what Prigogine implied in 1947. The entity must be sufficiently complex to contain the mechanisms — the infrastructure — that allow both the augmentation of its own complexity and prevent or greatly slow the natural decay not only of the newly complex organism but also of its own self. In the language we used in Section II.2.5, a template and regeneration process structure is separated by a permeable membrane from the working sensor-motor structure that it maintains, and perhaps rebuilds into descendant p-coli or di-coli.

Chapter II.4. Energy and Entropy in Control

In Chapter I.8 we introduced the emergent property called "tensegrity", a property of physical structures that gives the structure resilience, strength, and toughness. Mechanical tensegrity seems to be a characteristic feature of living organisms, and in Chapter I.8 we argued that it is likely to be a characteristic feature of the control structures of living organisms. In this Chapter the notion is extended using the tools of uncertainty analysis.

II.4.1 Control as cooling

Control is the process of reducing the uncertainty of an environmental variable, given a reference value for perception of that variable. As a practical matter, the environmental variable is not itself controlled because we have no telepathic means of determining its value at any moment, but instead a surrogate we call the perceptual signal is controlled. Good control therefore is a reduction in the uncertainty of the perception given the value of the reference. Nevertheless, reducing the uncertainty of the environmental variable given the reference value is what ultimately matters to the organism.

Why do I call this "cooling"? In physics, cooling represents a reduction in the entropy of a body, a reduction in the uncertainty about one component (e.g. atom or molecule) given knowledge of another. Entropy is a property of the body that depends on it having multiple degrees of freedom, as do many other properties, such as mass (dependent on the mass of its variety of constituent atoms), colour, dependent on the chemistry and arrangement of its molecules), roughness, and many more. Since we have mostly talked about control of a single scalar variable, the concept of entropy might seem not to apply. The temperature of a body, however, is determined by the variability in the energy of motion of its atoms, and this is a property that can be attributed even to a single atom, a scalar variable, as a function of time. The temperature of a body has no relationship to its global speed, the average directed velocity of its components. It depends on the internal relationships among the components.

When Boltzmann (1877) offered his statistical description of entropy, he determined that for an ideal gas, the component atoms in their collisions distributed the energy among them so that at any moment some had very low velocity, some very high, while most of them had moderate velocities. The distribution of energies would tend to a Gaussian form, and the entropy (which had previously been computed from gross measures such as pressure and temperature) would be calculable using a formula we have seen already: Entropy = $-\Sigma p\log(p)$, where p is the probability of an atom of the gas having a particular energy. The identity of this expression to the expression Shannon found to represent "uncertainty" induced Shannon to give the name "entropy" to what later writers have called "uncertainty".

In the case of a single scalar variable that is observed precisely, at any moment there is no uncertainty about its current state, but it changes over time, and has a distribution of states that it occupies, occasionally being high in its magnitude, occasionally low, but most probably in between if it is observed at well separated random moments. The distribution of values has an uncertainty, which would be the entropy of that variable. Reduction of the uncertainty of motion is cooling, so we can generalize the concept of cooling to the reduction of uncertainty of a single scalar variable, a single degree of freedom, which is exactly what a simple control loop does.

We are not really concerned with the workings of single control loops beyond noting that both reference and perceptual values vary because of inputs from outside the loop. The two inputs to a single control loop are unrelated except for possible connections somewhere outside the loop under consideration. These "outside connections", however, are exactly what the Black-Box and White-Box structures that develop the perceptual function of every control loop are about (Section I.11.4). The perceptual function defines what that control system controls — a perception (in Perceptual Reality) of some reasonably coherent pattern of effects that occur in the Real Reality of the environment of the control system when the control loop's action output changes.

The "outside connections" should not be ignored. Taking them into account, reducing the overall entropy of control error and "cooling" the local environment of the control hierarchy is exactly what reorganization does, by producing and refining the White Boxes that contain simpler White Boxes and connecting them in the form of new controllable perceptions that send their outputs to the inputs of the inner White Boxes. These inputs are the reference values for the perceptions controlled in those White Boxes — the lower levels of the perceptual control hierarchy.

The question, then, is not whether a single control loop reduces the temporal uncertainty of the single degree of freedom perception that it controls, but whether its operation within the hierarchy augments or reduces the mutual uncertainty among the error values in the hierarchy and between the reference values and the Real Reality environment. Here we are beginning to stray into the territory of four later Chapters, on autocatalysis, homeostasis, self-organization and maintenance, and "rattling".

One thing we will not ignore in this Chapter, however, is the cooling effect of tensegrity, if it were to be used in the control hierarchy, which we argue that it necessarily must be. Previously, we considered the distribution of the effect of a single energetic impulse among the component "rods and wires" of a tensegrity structure. Thermodynamically, a tensegrity doesn't cool overall. what it does do is tend to equalize the temperature over its nodes. In a control tensegrity, perceptual error variability is distributed across the perceptual levels, generally improving control quality within the tensegrity, which is another way of saying that the individual nodes perform their cooling (controlling) function better.

In a physical tensegrity structure, the shock wave of an event at one node takes time to propagate at the speed of sound within the individual wires and rods of the structure, faster along tense wires and slower along wires under less tension. An analogy in a control tensegrity is that wire tension represents perceptual error times loop gain, so that when one perception is suddenly disturbed, its effect is distributed at the speed of the control loops taking part in the structure. Control never takes full effect abruptly, so the sudden temperature rise of an abrupt disturbance of one controlled perception — a shock — takes time to be distributed by the cooling actions of the component control loops.

Here we are approaching the physical effects at the core of the "rattling" measure as applied to control, which is the topic of the next chapter. So let us move on.

II.4.2 The Use of Energy

The analogy between controlling and cooling is closer than a formal identity of the uncertainty and entropy formulae. Consider a home refrigerator. It has a mechanical mechanism that pumps a cold fluid through heat-exchanger coils in the interior[9]. If the interior air is warmer, some of its heat will be extracted from the interior and returned to the pump, whence by a process of evaporation and condensation the heat is redistributed out to the external environment of the refrigerator. At the same time there is a heat flow through the insulation of the refrigerator into the interior. As the interior temperature cools toward the temperature of the cooling coils, the rate of dissipation of that heat into the environment declines until it matches the heat influx through the insulation (plus the pulses of heating that occur when the refrigerator door is opened).

9. Some refrigerators use thermo-electric cooling, but the flow of energy is the same.

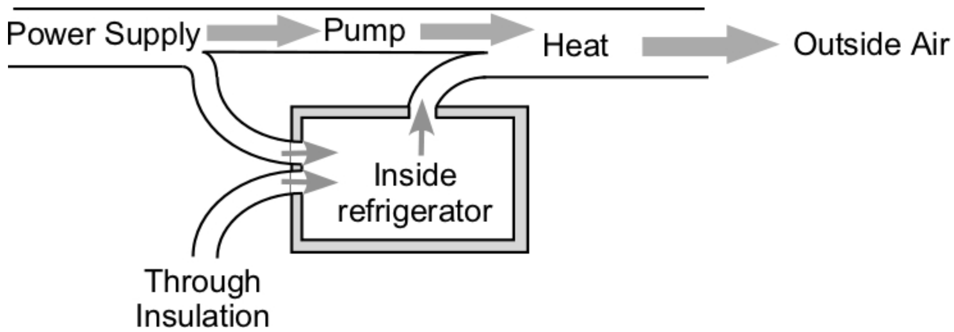

Figure II.4.1 Schematic of energy and entropy flow that keeps a refrigerator cold. The electric power supply runs the pump. Some of that energy is used for the cycle of evaporation and condensation of cooling fluid that extracts heat energy from inside the refrigerator at a rate that compensates for the flow of heat through the insulation. The rest is the waste heat from the pump operation, which together with the heat extracted from the refrigerator, is passed to the external environment.

This diagram of energy flow exactly models the energy and entropy flow in control. The external energy source, typically food for an animal, solar energy for a plant, provides energy to the control loop's functioning elements, such as nerve firings and muscular effort, and the actions of the effectors in opposing the disturbance to reduce its entropy is passed, typically as heat, to the external universe.

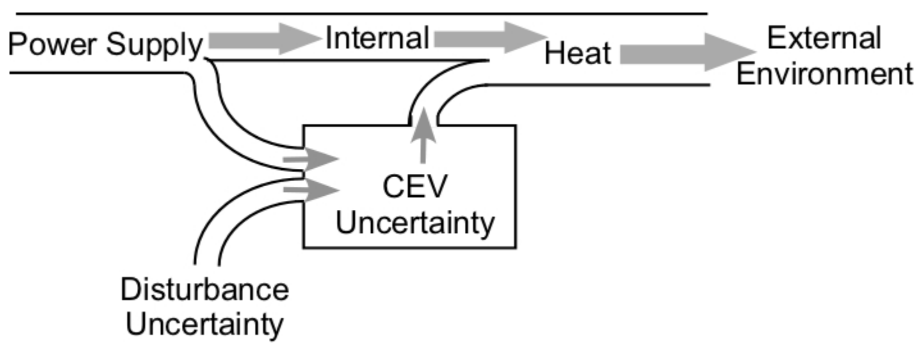

Figure II.4.2 Schematic of the energy and entropy (uncertainty) flows in a control loop. The power supply (perhaps food) drives the actions that oppose the disturbance, as well as producing waste heat. The countering actions reduce the entropy of the CEV in the environment, and the heat extracted by reducing the variations in the environmental variable is exported to the external environment along with the waste heat of the process, such as nerve firings and muscular internal friction .

Figure II.4.2 remains valid whether it describes a single control loop for controlling one scalar-valued perception or a whole lot of them, such as a vector of perceptions at the same level of the hierarchy or an entire complex hierarchic control system. From a thermodynamic point of view, control is simply a way of using a through flow of energy from a concentrated source to the diffuse external environment in order to cool specific parts of the local external part relative to the rest.

II.4.3 Control as Active Insulation

For the organism, what matters is not what happens in the external environment, but what happens inside the organism's skin. Above, we used the analogy of a refrigerator, which uses a flow of energy to cool its interior, and the analogy to the interior of the refrigerator was the internal perceptual environment of the organism. The organism uses a through energy flow, provided largely by its processing of food, to influence and stabilize the fluctuations of environmental variables against the varying influences of disturbances to the environmental correlates of the perception — the Real Reality correlate of the CEV. It does this because it is able to create an internal variable we call a perception that co-varies with the environmental variable, so that when the perception is stabilized, so is the environmental variable, and vice-versa.

If that were all there was to it, any environmental variable and its perception would be as useful to stabilize as would any other. But as we pointed out when dealing with reorganization in Section I.11.4, that is not all there is to it. The ability of the organism to act, to survive, and crucially to pass on its genes to descendants, depends on the effective maintenance of the interactions of internal variables that Powers called "intrinsic variables", which are largely biochemical in nature. These variables and their interactions are the reason why the control of perceptions is important. Control of the "right" perceptions reduces the kind of variability of the external environment that would adversely affect the intrinsic variables, while tending to maintain them at values that benefit the intrinsic variables.

The refrigerator does not rely entirely on the energy flows of Figure II.1.2 to maintain the thermodynamic stability of its interior. It has insulation that separates its inside from its environment. "Insulation" is for thermal isolation, the reduction of the influence of the impact of each molecule of the environment on any molecule of the interior (and vice-versa). Each molecule can be represented (as Boltzmann, 1877) did, by its own set of six degrees of freedom — three for location and three for velocity or momentum, each of which has a magnitude that varies very rapidly.

No organism that moves macroscopic effectors could possibly control against so many degrees of freedom changing so rapidly, not within a few tens of orders of magnitude. What they can do, however, is to control a few degrees of freedom that change slowly, and one of those degrees of freedom is what we label "temperature". If the temperature on one side of the insulation is subject to larger and more rapidly varying independent influences than is the other, the insulation protects the more stable side from the variable influences that would otherwise be transferred from the more variable side.

The equivalent of insulation for macroscopic mechanical influences (the effects of many of the external molecules moving coherently, as they may when they form part of an object) is a shell. Shells may be stiff or soft, but whatever their manner of reducing the transmission of mechanical effects across them, they reduce the influence of the external environment on the organism within. The same is true for changes in the external chemistry, since shells do not permit molecules individually to pass through. Some membranes do, but shells do not.

Although a shell may protect the organism within from macroscopic and chemical influences from outside, that protection cannot be complete, because, at the very least, the organism must accept food from outside and expel waste to the outside if it is to maintain its structures, its low entropy state relative to its environment. If all that could get in through gaps in the shell was food and all that could get out was waste, then the organism could survive in its environment if the environment happened to supply food without the organism having to act. But no organism is so lucky. Always other influences can get through the gaps in the shell, so the organism must do something to discriminate between the food it wants and the other influences, which it does not want. It must act.

Even plants act. Their food is largely sunlight and carbon dioxide, and they do not have to move in order to get either of those; their waste is largely oxygen, and they don't have to move to expel it. But they do change shape. to get more sunlight they may grow new branches, to avoid excess sunlight they may curl their leaves. Most plant actions, such as growing toward sunny gaps in overlying foliage, are very slow compared to our own movements, but not all are. The snap of a Venus Fly Trap, when it senses a touch on its sensitive part, is very quick, too fast for the fly to escape. Many flowers turn to face the sun as its orientation in the sky changes over the course of a day. For the most part, however, the actions of plants are too slow for our eyes to detect.

Other life forms from bacteria to fish and elephants must move if they are to acquire food. They must go where it is and ensure that it passes through a gap in whatever shell they may have, a gap that we usually call a mouth. To go where food is, they must be able to sense some aspect of their environment, perhaps a chemical concentration such as a smell, perhaps a sound, perhaps some pattern of light and shade. To enhance the likelihood that what passes though the mouth is food rather than some damaging toxic material, they need sensors we call taste.

In other words, simply to get food, mobile life forms must be able to form within themselves some kind of interior representation of some property of the environment. We call those interior representations perceptions. Some of those perceptions will be of things that would be food if they passed through the mouth-gap in the protective shell, others are of things that would be toxic. Organisms that acted to permit the entry of food while avoiding the entry of toxins would survive longer to produce descendants than would those that failed to discriminate.

In the language of PCT, the more successful organisms would control with a positive reference value the quantity of food they accept, and with a zero reference value the quantity of toxins. The sensor would not be the cell that transforms an influence from te environment into electrical impulses, but would be a complex of them, perhaps of taste, colour, and configuration (perceived shape), and their action output would be either ingestion or ejection of the material that entered the mouth.

The reference value for the quantity of food to be ingested would vary, depending on the current biochemical state of the organism — in common language, how hungry it is. The perception of "food" is opportunity, which does not automatically lead to ingestion. If the reference value for quantity of food is at that moment zero, anything that entered the mouth would be ejected, no matter its taste.

Fast forward thousands or millions of generations, and we may confidently assert that organisms that could sense at a distance whether some discriminable pattern of sensory input would represent opportunity or danger if it entered the mouth (or, at this later stage, if it were to otherwise affect the organism) would tend to survive and reproduce better if they could act to take advantage of opportunity or avoid the danger before it actually entered the mouth or struck the skin or shell than would those that were otherwise similar but lacked the capacity to act as effectively. The successful ones would control perceptions of distant events rather than of proximal ones.

Just as the early organisms could control for ingesting or ejecting things in the mouth by discriminative tastes, so these later organisms would discriminate using distance sensors, and would survive better if they could influence those distantly perceived patterns from their environment to admit the desirable influences and deter the undesirable. They learned to control the locations of things they could see, in effect cooling more and more individual degrees of freedom in their environment.

The effect is that of a thick, rather selective, membrane around the organism, within which the space is cooler — lower entropy — than is the average of the outer environment. This skin is equivalent to the refrigerator's insulator. The more perceptions an organism controls, the more discriminative the "membrane". In the same way, the membranes of cells have pores that open and close to admit or expel specific ions and molecules, but are closed to most.

The environment controlled thus by a multitude of controllers at many levels of the control hierarchy becomes ever more open to opportunities and ever less dangerous to the organism. The export of entropy from within the organism is implemented by perceptual control, but the gross effect is of export to a lower-entropy local environment, just as a hotter body can export some of its entropy to a cooler one.

II.4.5 Uncertainty and information around a control loop

From an information-theoretic viewpoint, the through flow of energy can be ignored, though the export of uncertainty cannot. The whole point of control is, as Friston points out for the "Free Energy Principle", to minimize the uncertainty of the CEV/perception given the reference value. In Figure II.1.3, the environmental variable CEV in question may be a single scalar-valued quantity or a whole vector of them. A disturbance (D) has a certain bit-rate (Section 10.2) when observed by the input processes of the loop. To balance the momentary value of the changing disturbance, the output must have the same uncertainty, observed from the same point. This must remain true even if the reference value remains constant.

Observed at the comparator, a well-controlled perception has a much lower bit rate than the disturbance. Indeed, the difference between these bit-rates, each given the momentary reference value, is a measure of control quality mathematically derivable from the variance ratio (and vice-versa), perfectly accurately if the variability is Gaussian, and closely if the variability is not too non-Gaussian. The bit-rate difference, however, is valid no matter what the form of variation, discrete (e.g. on-off or categorical) or continuous.

On the face of it, Figure II.1.3 presents us with a bit of a puzzle. Where does all the extra information at the output come from, when the error supplies so little? Think about that before reading below the Figure caption. A hint toward the answer is in Figure 10.1.

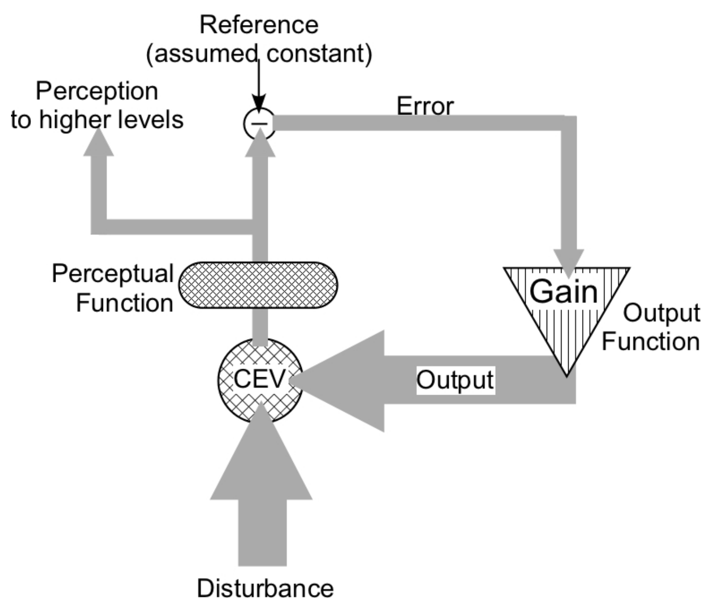

Figure II.4.3 Schematic of the uncertainty levels in and around a control loop. The arrow widths represent relative uncertainties. The uncertainty out of the CEV is that of the perception if noise in the path from the sensors through the perceptual function is ignored. The output function achieves its increase in uncertainty by averaging in some way (such as a leaky integrator) the error values in recent history.

The answer to the puzzle is, as Figure II.1.3 implies, in the Gain of the Output Function. Although the Environmental Variable, the CEV, has some precise value that could be specified to any desired numerical precision, the perception of it has a finite resolution. The output of the Perceptual Function specifies a macrostate for the CEV, not an infinitely precise value. Nevertheless, the perceptual value that is output is whatever the function produces from its exact input. The same is true for the error value. An exact value is presented to the Output Function.

What is the main property of the Output Function? Gain. If each stage in the loop between CEV and output has the same resolution, the Gain changes the overall size of the macrostate, allowing many more microstates within the same macrostate, and increasing the information in an observation of which microstate corresponds to the error input, thereby increasing the output bit-rate so that it can match the bit-rate of the disturbance.

The foregoing is a very crude and in some respects invalid answer to the puzzle, but it outlines a more complex correct answer, which starts by recognizing that a macrostate here is not a "resolution limited" box, but a distribution of probability for the value of the output of a function given the value of the input, both values being precise. A precise value for the CEV thus produces a precise value as input to the Output Function of the loop, but there is uncertainty in the relationship of one to the other. That uncertainty cannot be reduced by a simple Gain multiplier, but the uncertainty of a time-average is less than that of any single sample. That time-averaging can be done by the leaky integrator often used in control simulations. How well it works in producing an output that opposes the disturbance depends on two things: how fast the disturbance changes, and how the loop delay relates to the disturbance bit-rate.

In Figure II.1.3, the energy flows are ignored, but one aspect of the diagram needs explaining: the increase of uncertainty introduced by the output function. Since, if the output is to oppose the disturbance, its uncertainty must at least equal that of the disturbance if U(perception) = U(disturbance|Output) is to be minimized, the uncertainty increase through the output function cannot be due to random noise. It must be related to the disturbance, and the only source of information about the disturbance it can use is the fluctuations of the error value.

There is no reason why all past error values should have the same weight until they are ignored. The most recent values of the error are most nearly related to the current value of the CEV and hence of the disturbance, and what reason could be given for a not very old value of the error being fully weighted one moment and becoming valueless the next instant?

Both considerations lead to the conclusion that the best weighting would be one that drops smoothly over time at an equal weight reduction per second, namely a declining exponential. This is the weighting given to the error by an output function that is a leaky integrator. It is probably no coincidence that Powers found that simulations using leaky integrator output functions matched human performance very well in a variety of situations.

Since the transport lag together with the disturbance bit-rate determine how much the disturbance uncertainty increases before the output can oppose it, the added lag caused by the time-averaging must be included. But how much, given that we keep a sample of the input in the average forever?

The logical extra delay is to the centroid of the distribution, at which point the sample will have had half the effect it will ever have. This time turns out to be ln(2)/r≈1.44/r (seconds, samples or whatever unit time is measured in). If the weighting were equal over time this would add one bit of uncertainty to the calculation of (disturbance bit-rate)*(total lag), reducing the achievable quality of control by one bit. However, since the weighting is skewed to the more recent input data, and the centroid therefore has less lag than would be the case with equal weighting over time. The reduction from the "loop transport lag only" ideal is therefore somewhat less than one bit.

These possibly abstruse calculations are icing on the cake, the body of which is that the time averaging done by the output function is what produces the Gain and the enhancement of output uncertainty that is needed if the output is to oppose the disturbance optimally. The body of the cake is the same, no matter what the icing of any specific form of time-averaging. The exponential decay form of weighting generated by the leaky integrator output favoured by Powers seems to be optimum from several points of view, and provides good simulation matches to human tracking behaviour, besides being easy to implement. We have no reason to believe that it is very different from what biological control systems do.

II.2.9 What is Infrastructure?

Most objects can be used for many purposes; they provide atenfels for control of a wide variety of different perceptions. Objects are atenexes, because the control of any of the perceptions of their properties might be suitable for controlling some unrelated perception or other. Even a simple pebble can be used as a decorative object, a paperweight, the bob of a pendulum, an indicator of the depth of a well, a test for the strength of one's throwing arm, a surface against which to sharpen a blade, a source of edged blades (if its material is suitable), a creator of pretty ripples on a pond, and in many other ways. A bar of metal milled to a precision finer than a micron for use as a length standard can be used to prop open a window or as a murder weapon. We usually do not regard an object as an atenex if it is designed to have a single primary use as an atenfel, particularly if they are used by only one controller at a time.

Another way of thinking about the variety of atenfels provided by an object is to distinguish between designed "tools" (among which we may sometimes include "roles" people play — Chapter III.2ff.) versus natural objects that just happen to be at hand. A tool is designed and shaped to serve a purpose. "To serve a purpose" is an everyday way of saying "to enhance the control of a particular kind of perception". Although a hammer can be used as a pointer, a paperweight, or a bludgeon, in its role as a tool it is designed to make driving nails into wood easier than would be the case without the hammer. Controlled perceptions that include the relationship between a nail and some wood can often include the tool-atenfel of the hammer in their environmental feedback path.

An abstract concept, such as a role that happens to be performed by a person, may by design provide one or more atenfels. If you want a spill in an office cleaned up, you could do it yourself, or you could call a janitor. If you want to do a banking transaction, nowadays you may use an electronic device, but in earlier days you would use a bank teller. The persons of the janitor and the bank teller are individuals, but "janitor" and "teller" are roles that could be filled by anyone, or even a machine, who could competently play the role. The atenfel is associated with the role rather than with the person. If you decide not to call someone to clean up the spill and do it yourself, for the moment you are playing the role of janitor.

As McClelland points out in the passage quoted, much of our ability to control our perceptions is the result of people having worked together to produce and continuing to work together to maintain in good condition tools and larger-scale structures that provide atenfels for many particular perceptual controls — designed atanexes, or in a word, "infrastructure". For an atenex properly to be a component of infrastructure, though, at least one of its atenfels should function together with another atenex, forming a molenfel to ease control of some perception people might want to control. For example, think of pen and ink, cars and paved roads, or tall buildings and elevators, each of which is more useful with the other than it would be by itself.

When cars were first invented, roads were often no more than rutted and often muddy cart tracks. Controlling perceptions of where the new-fangled machines would go was much more difficult than it is now. Soon, however, roads began to be paved to replace the rutted mud tracks or dusty trails that were adequate for travel by horses and (uncomfortable) horse-drawn vehicles. Paving greatly eased your control of where cars could take you (as well as your comfort as a passenger).

The people who paved the roads may not have been controlling perceptions of the ease of future car travel, perceptions they can have had only in their imagination, if they had them at all. They were more likely to have been controlling their perceptions of cash in the wallet, by digging and tarring what and where someone asked them to dig and to tar.

The people who organized the paving, however, must have perceived at least that the paving would make it easier for unspecified others to travel to distant places connected by the newly paved roads, and that those others might want to do so. They might well have been controlling their perception of cash in their wallets rather than the ability of future travellers to move quickly and comfortably, but their means for controlling their cash perception was to design good paved roads and pay people to build them.

The existence of paved roads then provided atenfels for many other controlled perceptions, such as the ability to assemble complex objects from parts manufactured somewhere else. The existence of department stores made it much easier for people to buy what they wanted than was the case when purchases had to be made by visiting a large variety of small specialized stores. The fact that people had this new facility then provided the atenfels that allowed other entrepreneurs to create factories for producing a greater variety of products that people might want to buy. The factory products would be distributed to department stores by way of the paved roads and railways, which equally depended for their usefulness on the invention of trains powered by mobile energy sources.

We are talking here about homeostatic networks made possible by the creativity of older homeostatic loops. The people who pave the roads could get to the work site by using cars on roads already made. Their pay might be in the form of a paycheque that can be turned into cash in a bank building whose construction materials were brought to the site on those same roads. And so forth. There are myriads of connections that allow control of perceptions because of stabilities created by control of other perceptions. And as McClelland points our, it requires work to create those stabilities, and more work to maintain them. We will return to this work of maintenance when we deal with money and trade in Chapter III.7.

The word “stability” does not differentiate among all the perceptions that induce stability in the environment by being controlled and stabilized. What, then is being made “stable”? It is some property or set of properties firstly in perception, but also in the real world that provides the input from sensors to the perceiving apparatus and is influenced by outputs from the organism. The stability that matters is in real reality, however unknown to us it might be. What we know of it, however, is manifest in the CEV, which might be as simple as the location of a glass of water or as complex as a legal tax code. The CEV might be a configuration of material components or a purely abstract structure of relationships. It is as complex as the perception that is controlled, and that perception is a structure, a composition or function of simpler constituent perceptions — and of the corresponding CEVs and whatever is contributing in real reality to the perceptions.

The word “structure” is important. It represents patterns of relationships: “*This bit fits there, so that this other item can connect those parts*”. To put together a piece of furniture from IKEA, one either has in mind a planned sequence of controlled perceptions — a reference for a sequence perception — or one reads the instruction plan provided with the item. "Structure" is a reduction in uncertainty, which is the Shannon definition of information (Chapter 10). It is also, in the Kolmogorov approach to information, the information contained in what you did not already know beforehand about building the furniture from its components. Structure can be made quantitative, being the amount of information required to build the structure given that the component elements are known but their relationships are not.

What differs between "structure" and "infrastructure"? "*Infra*-" implies "below". In this case, "below" is conceptual, and it means "supporting" or "foundational". Just as the foundation of a building allows the part in which people live and work to persist rather than to collapse into soft ground, so the "structure" aspect of "infrastructure" is in the uses to which the structure may be put. Any structure may serve as an atenfel for controlling different perceptions, but infrastructure is used more as an atenfel for perceptions involved in building further structure.

The boundary is fuzzy. The less the structure is for immediate use, and the better it serves as atenfels for control of a variety of other perceptions, the more the structure deserves to be called "infrastructure". Education and experience build infrastructure for the individual organism to control perceptions of many kinds, and as we shall see starting around Chapter III.6, organizations and networks of social practices and legal systems provide infrastructure that enable their individual members to control perceptions they could not control by themselves.

There is a stigmergic[10] aspect to infrastructure. In the context of inventive homeostatic loops, the stable infrastructure loop products are catalysts for many different reactions of a wide variety of kinds. They are atenexes that "react" with other atenexes to provide the molenfels that are actually used for controlling perception. Just as a paved road offers no benefit over a path or track to anyone without a wheeled vehicle, and as a telephone network is no use to anyone without a telephone, or an electricity supply to anyone that has no electrical appliance, a structure that is infrastructure is ordinarily used as a molenfel with some other structure.

Much civic infrastructure is the structure of communication, but the products whose production is enhanced by improved communication also can be infrastructure. Tools, roads, houses, phones, construction machinery, and so on and so forth, all serve to ease perceptual control or to enable control of certain perceptions that were not anticipated by their constructors. To create these atenfels requires a variety of different skills, not least among which is the ability to organize.

Perceptions that leave their tolerance bounds require action to return them to the reference macrostate. In other words, all structure needs maintenance. When the macrostates, the specifications of the structures, are complex, such maintenance may require even more organization than does the initial creation of the structure. It is often easier to build an entirely new house than to renovate one that is in a bad state of repair. As McClelland points out, maintenance requires work. Work uses energy, but what the energy does is reduce or keep low the entropy of the structure's components considered together, as we discussed in the first part of Chapter xII.1.

McClelland uses the term "work" several times in the quoted passage. Maintenance, whether of a conceptual or a physical structure, requires work, though not necessarily the work of direct perceptual control, as we have already seen in our discussion of homeostatic loops. In this context, "work" is used in the sense of physics. It is measurable, and is no metaphor. Doing work generates heat, and that, too, is no metaphor. Every action in the real world generates heat, whether the action is physical (including electrical) or chemical. It is even true of perceptual controlling and of conceptual work, imagining and thinking, since conceptions are built by neural impulses that are electrochemical events in the brain. The removal of heat from the brain is one of the major evolutionary constraints on its structure and that of the

10. "Stigmergy" is the study of the abiding effect of past actions on the current environment.

skull that contains it.

In everyday parlance, “work” is the energy used in deliberate behaviour, which according to PCT is always the control of perception. As I argued, back in Section 6.2, the ultimate measure of “worth” is the availability of energy to control arbitrary perceptions. Maintenance and construction work is action to control some perception in the worker, but the apparent worth of infrastructure work is distributed across a community of people who can use the results of the work to reduce the energy they use in controlling their independent perceptions.

Infrastructure as a rule has this effect of allowing the distribution of structure created in a concentrated locality to a wide-ranging set of possible uses, another way of saying that infrastructure catalyzes a wide range of possibilities, or of saying that infrastructure provides atenfels that combine easily to form useful molenfels.

For the worker, the work ordinarily augments her individual worth by more than her cost of performing the work, either because of the money earned or because performing the work reduces error in some controlled perception such as her perception of the perceptions of herself by others. Typically, her augmented worth is in her perception of the amount of money available to her, but it need not be. It could be improvement in control of, say, self-image if the work is done as a volunteer for a charity.

Whatever the worker controls, it seen by others as part of a job. The job defines a role the worker is playing at that moment. The worker’s higher-level controlled perception may be the perceived value of the work to the community, but the effect it has on the ability of an anonymous other person in the community is always a side-effect of the worker’s control actions when controlling for playing a particular role.

Infrastructure is a statement about levels of organization. There would be nobody doing the work described by McClelland if nobody was concerned with easing the interactions among people they would never know. “Infrastructure”, as a word, means “the structure below”, but what is below and what is above depends on your viewpoint. The infrastructure of water and electricity supply, of laws and administrative bureaucracy, are all there to facilitate the ability of any person to control by means of actions performed by others. The multi-person structure incorporates all those other people, so “infrastructure” can equally be seen as being at a higher level than the people who use it. From the viewpoint of any one such person, though, the infrastructure is “just there” making things easy.

Neither “above” nor “below” is any more appropriate to the relationship between people and infrastructure than it would be in comparing the roles of nodes and links in a generic network. The network comprises both on an equal footing. We will use this duality extensively in Chapter II.6. In a multi-level control hierarchy, the level below consists both of active controllers and the linkages between the levels that may occur through even lower levels and through the environment, all of which constitutes the infrastructure for the higher-level controllers.

II.4.6 A Two-level Hierarchy

For the purposes of the following discussion, we can reduce the control loop to a single construct, a reducer of uncertainty between the disturbance signal and the perceptual signal fed to the next level of the control hierarchy (as suggested by the different thickness of the disturbance and the perception arrows in Figure II.1.3). The uncertainty that is fed to higher levels is what remains of the uncertainty of the disturbance after what can be corrected has been corrected. At that higher level, several lower level control loops provide their perceptual signals to create the perceptual equivalent of the higher level CEV in the environment. If they are independent, their uncertainties simply add (Figure II.1.4).

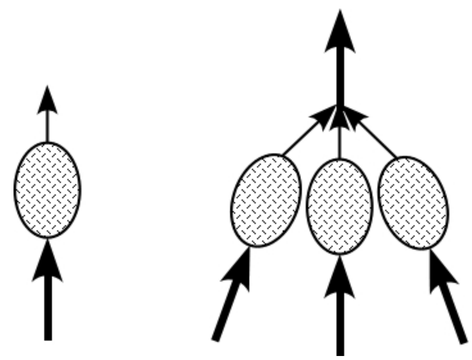

Figure II.4.4 (Left) A schematic suggesting the uncertainty reduction achieved by a single control loop, and (Right) the combination of the uncertainties of three perceptual signals that combine to provide the input to a higher-level perceptual function. The signals themselves do not combine, but their uncertainties do, additively if they are independent.

Figure II.1.4 shows nothing of the actual workings of a control loop. It illustrates only that the combination of signals that serve as inputs to the perceptual function of the next level would add, providing a greater uncertainty than any one alone. This uncertainty is not, however, that of the perception in the higher-level loop to which they contribute. It is the disturbance uncertainty that remains at the higher level.

Figure II.1.5 (Figure 8.17 repeated) illustrates a two-level structure slightly differently from the way such structures are usually shown in the PCT literature.

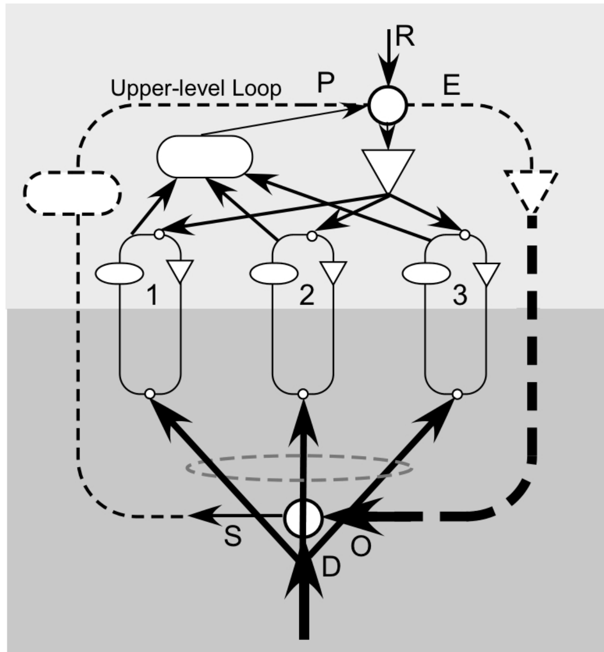

Figure II.1.5 (Figure I.8.17 repeated) A two-level control loop. Three simple control loops send their perceptual signals to the second-level perceptual input function and receive their reference values from the second-level output function. The CEV of the second-level loop (lower white disc) is disturbed and the disturbance opposed by the combined outputs of the simple loops, shown here as the output of a virtual output function. Likewise, the second-level loop can be seen as having a virtual perceptual function that incorporates the actual perceptual functions at both levels. Line weights suggest uncertainty bit-rates.

As Figure I.8.17, Figure II.1.5 was used in illustrating tensegrity of control. It is no different when the arrow widths are interpreted as bit-rates rather than as variances. The higher the disturbance bit-rate, the more difficult it is to control the perception, and this configuration distributes the higher-level disturbance bit-rate across disturbances to the three lower-level CEVs. Each of the lower-level loops (loops 1, 2, and 3 in the Figure) has an easier job than if all the disturbance bit-rate to the upper-level were caused by disturbance to just one of them.

This is where the tensegrity effect comes in. Suppose only loop 1 were initially exposed to a varying disturbance. It would control away most of the uncertainty that acted as a second-level disturbance. The second-level output opposing the second-level disturbance would, however, be transmitted to all three first-level reference inputs, causing the uncertainties of their errors to increase as though they had experienced their own individual disturbances. The total still will match the uncertainty imposed on the higher level by the loop 1 disturbance, but the uncertainties of loop 2 and 3 tend to oppose that sent to the higher level perception (because of the sign reversal of error=reference–perception in each loop in which the uncertainty is entered at the reference input rather than at the disturbance input — the CEV).

In Figure II.1.5, a virtual loop is shown dashed. It has a virtual perceptual function and a virtual output function corresponding to the real ones shown connected to the lower level loops, but a real comparator (as real as are any of these hypothetical structures) and a real CEV in the external environment. The real CEV is the variable that is important. Perhaps, as in Section 5.3, it is the location of a chair, and the lower level CEVs (the small dots at the bottoms of the simple loops) are the locations of a leg, the seat, and the back of the chair.

The virtual perceptual function of the outer loop is the result of perceptual processing by the lower level loops and by the second-level (real) perceptual function that has the lower-level perceptions as its inputs. The same is true of the virtual output function, which produces the output that is actually created by the simple output functions as they are passed through their individual environmental feedback paths to be recombined at the CEV.

I used the example of a chair and its components to represent environmental structures that have become the objects of stable perceptual functions because controlling them as unitary entities reduces conflicts in the control hierarchy, improving control by the lower-level loops. How is this improvement achieved? It is not achieved in the way Powers reduced conflict in his demonstration of reorganization in a 14-degree of freedom arm model in LCS III. In that demonstration, the output functions reorganized to become orthogonal and thus avoid conflict. Here, however, we assert that the perceptual and output functions have become more interdependent, not less, and yet the potential conflict among them is reduced. This seems to contradict Powers's approach, but it does not.

The reason that there is no contradiction is that Powers's model arm acted in an environment in which its own structure was the only structure in its Universe, which meant that if perceptual signals were pre-tuned to report only a single variable that could be changed independently, such as wrist angle, independent control automatically eliminated the disturbance effect on one perception of the side-effect influences of the actions in controlling another. With the chair in the environment and its components being acted on separately, the side effects of controlling the location of one component would disturb perceptions of the others except under very special conditions. Those conditions involve making the control actions less independent rather than more so.

The "special condition" was described in Section 11.1 as the development of a higher-level perception of "chairness" that would coexist with and would use the perceptions of the chair's legs. seat an back as its elements. The incoming uncertainties would not be uncorrelated in the way Figure II.1.5 presumes, but would be highly, though not perfectly, correlated. The 6-D perception of the chair's location and orientation would have a certain total uncertainty, but this uncertainty would be much less than the sum of the uncertainties of the 36 component variables taken individually.

The same, then, is true of the reference values sent down to the component reference inputs at the lower level. They relate to each other so that the chair moves without the "leg X position" output trying to move it in an incompatible way. Of course, it is always possible that an independent higher-level control unit wants the chair-leg X position to be in another place for a different reason, which would generate a resource-limitation conflict (which we discuss in Chapter II.7), but that is a different matter entirely.

II.4.7 Distribution of effects

Suppose in the chair example something, such as a large dog, is applying varying forces to the chair leg. If those forces were unopposed, even by friction on the floor, the chair leg would move with accelerations always proportional to the the applied force — F=Ma, where M is the mass. So what is the mass being moved? It is not only the chair leg, it is the chair, which includes all its components.

If the chair location is a controlled (vector) variable, then some opposing force might be applied to some other part of the chair, but could still result in the force on the leg being properly opposed without the side-effects affecting other perceptions of the components. In practice, however, there probably would be side-effects caused by torques, so reorganization would have been likely to adjust the inter-level weighting so that the opposing force was preferentially applied to the component being disturbed.

This example may be a bit extreme, because the chair is a rigid object, but the point that it illustrates is that if a perceptual function has been generated by evolution or reorganization, and has been stable because it corresponds to some structure in external reality, then the effect of an influence on one component of the structure will be likely to be felt on other components directly, and if that influence disturbs a controlled perception of the component and therefore of the structure, those influences will be back-propagated down through the distribution of reference values. The propagation of effects through the control hierarchy (so far only a two-level hierarchy) is reminiscent of the propagation of effects through a physical tensegrity structure (Chapter 7), an analogy we will soon pursue further.

From the standpoint of uncertainty, we have a similar distribution issue. We might look at the uncertainties of the disturbances on the components of a structure, or of the disturbance to the structure as a whole. The latter implies that we consider the dashed outer loop in Figure II.1.5, while the former implies that we ignore the outer loop and consider only the simple loops and the ECU (the triplet of perceptual function, comparator and output function) of the higher-level controller.

It would be better to do neither at this point. Instead, we consider the outer loop of Figure II.1.5 together with the component loops and the upper ECU. To this assemblage of "real" and virtual components, we have four places where disturbances can introduce uncertainty from outside, the structure itself as a CEV, and the three components shown in the Figure. The uncertainty due to variation of disturbance to the structure is necessarily passed to the CEVs represented by the components, but they could, in principle, be subject to additional disturbances, so the total disturbance uncertainty passed into the two-level hierarchy might exceed that of the disturbance to the structure. This would occur if the structure were not hard and rigid, like the chair, but softer and more flexible, like a doll in which the arm can be manipulated without affecting the legs, but the doll as a whole can at the same time be picked up and carried.

The uncertainty passed into the individual component CEVs can be corrected by those individual control loops, whether it comes from the uncertainty of the disturbance to the structure, from an independent source, or from both together. But if the reference values for the three components remain constant, either the structure will not have its disturbance countered or there will be conflict. For each component, we have two *independent* uncertainties, that distributed from the disturbance to the structure and that from its own distinct disturbance.

But behold! From the structure of a control loop, the component loop has two different independent source of uncertainty: the disturbance and the reference. From the argument above, the uncertainty in the reference is distributed from the output function of the higher-level controller ECU, and is strictly concerned with countering influences on the structure. From that point of view, a control loop is simply a pass-through filter that sets its CEV to the value of its reference signal (Figure II.1.7).

The first thing to notice is that if the tensegrity structure as a whole is not to heat or cool continuously until it burns up or freezes to death, it must, on average, export as much energy and entropy as it absorbs. A single control loop does this by using a through energy flow that comes in at low entropy and departs with increased entropy. This is literal, physical, energy and entropy, because it is the literal, physical, environment that is affected by the control loop. Internally, however, the physical energy that must be exported is in the form of nerve firings. In PCT, a "neural current" is the average firing rate over a bundle of nerves, so the "neural current" that is the output from each processor around the loop to the next is the source of the energy that must be exported as a consequence of control.

Each node of a normal tensegrity structure consists either entirely of wires connected at a point where the tensions in the wires balance in all dimensions, or of a collection of wires that together apply pressure to a rod. Tensegrity-like structures may include two or more rods connected at a flexible junction where wires may also connect. Whatever the actual structure, at every node the sum of the compressive and tensile forces in every direction must balance asymptotically in the absence of external disturbance.

When force is applied at a node, it unbalances the forces at that node, increasing or decreasing the tension on a wire and the compressive force on a rod. Each of those forces, multiplied by the amount the wire or rod changes from its resting length, is the energy stored within the tensegrity member. This stored energy is manifest in the changing "stiffness" of a wire or a tug-of-war rope (Section 8.1) that has no intrinsic stiffness, as well as in changes in the length of a wire in tension or of a rod in compression.

Every wire or rod connected at that node distributes its changing tension or compression force to the node to which it connects, and those nodes likewise distribute the forces they receive further around the structure, always changing the energy stored in the components to which they link. Eventually, these distributed forces arrive back at the node where the original force was applied, resisting it. The tensegrity structure as a whole acts like an elastic body, the elasticity depending on the tensions in the wires, manifest in their stiffness, which is itself distributed over the structure.

Uncertainty is the central concept behind a principle called "Low Rattling" (Chvykov et al. 2021) that applies to organized (and initially, unorganized) structures. We will discuss the rattling measure starting in Chapter II.5. Suffice it to say here that tensegrity structures are, by their nature, organizations with low rattling measures compared to the rattling that would be experienced by the components if they were to be connected differently, so as not to exhibit tensegrity properties. Accordingly, the low-rattling principle suggests that sufficiently complex structures will tend, over time, to change into structures with tensegrity properties. Low rattling is a property, but not an emergent property, of tensegrity.

In a control analogue, the tension in a "wire" is manifest in the error values in the controlled perceptions throughout the hierarchy, modulated by the loop gain. We will discuss this further when we come to discuss "rattling" in Chapter II.5, where "Low Rattling" is connected with reorganization of the perceptual control hierarchy. At this point, it is sufficient to say that a tensegrity structure is a low-rattling structure, and hence a structure toward which both individual internal organization and social organizations should be expected to tend. Before we even look carefully at a perceptual control hierarchy, we should therefore expect to find that any control system whose structure evolved rather designed would be a tensegrity structure. This expectation applies to the perceptual control hierarchy, which should on average reorganize as a tensegrity structure (in coordination, of course, with the biochemical organization of hormones and enzymes that produces Powers's "intrinsic variables").

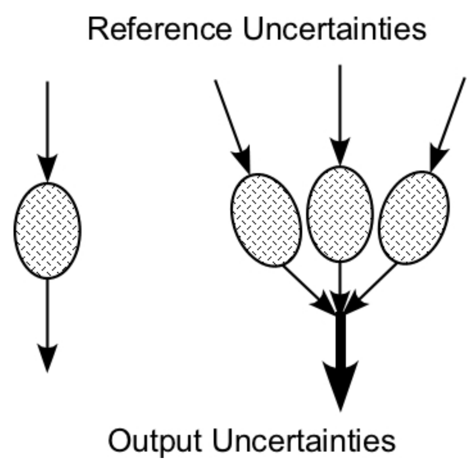

Figure II.1.7 The output side equivalent of Figure II.1.5. The reference uncertainty of a control loop is the same as the uncertainty of its reference if the disturbance is constant. However, the uncertainty of their combined effect on a higher-level variable can be anything up to the sum of their individual uncertainties.

How much effect the output has on the combined structure depends on the effective gain of the higher-level ECU (Figure II.1.3), rather than on the behaviour of the component control loops, at least if the control quality of the supporting loops is good. The important point is that the uncertainty passed through from the reference to the CEV is independent of any uncertainty associated with the component loop due to its individual disturbance. Hence, the upper-level ECU can treat the lower-level components as though they were simply passive (though possibly noisy) parts of its environmental feedback function.

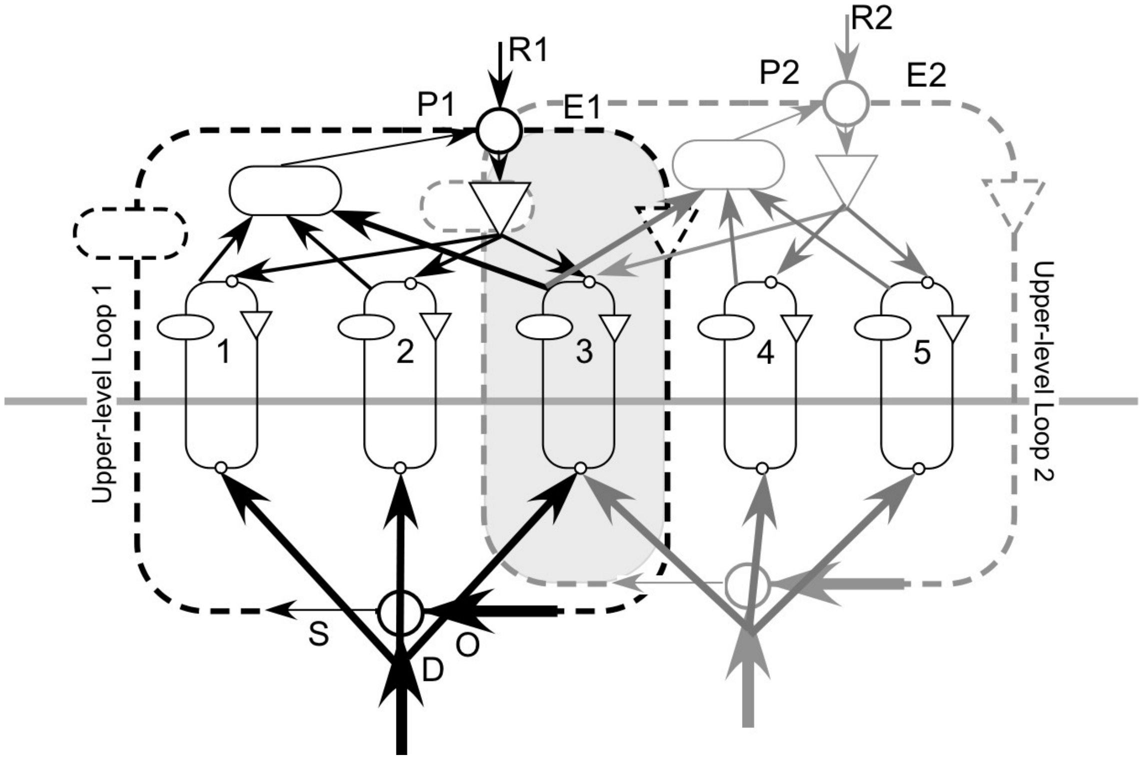

Figure II.1.8 (Figure 8.18 repeated) A second upper level control loop may share some of the same lower level perceptions and/or outputs as part of its control of its structured perception.

Since the error uncertainty due to the disturbance to the CEV of loop 1 of Figure II.1.8 is much greater than the error uncertainty induced at the reference input to loop 3, the distribution of uncertainty "stress" to loops 4 and 5 is less than to loops 2 and 3. There is no demarcation line between parts of the structure that help in control against the disturbance introduced at one point. Rather, there is a declining effect as more traversals up and down between levels are needed. What happens in this regard as more levels are implicated in the distribution of uncertainty?

It would be unusual for only one CEV to be disturbed over a long enough period for the distribution of uncertainty stress to arrive at a stable state. More probably, the uncertainty of the disturbance will be non-zero at the CEV of every control loop at every level of the hierarchy. How do the uncertainties of these differently applied disturbances generate their effects on the structure?

The key question that must be asked is about their informational independence. If all the disturbances were independent, the total uncertainty to be distributed through the structure would be their sum. But since we are dealing with a hierarchical structure built by evolution and reorganization, they cannot be independent. The disturbance to a CEV at one level is seen by the corresponding perception only in the form of variation of the perceptual inputs of the components.

Thinking of it in the other direction, if the CEV that is an environmental structure is disturbed, that disturbance will affect the component CEVs in an informationally related way, the exact relationships being determined by the structure. If you move a chair, its legs, back, and seat will all be disturbed in a coordinated way. They are informationally related (technically "redundant"), though not necessarily identical, as we saw in Chapter 7. The way they differ depends on the actual structure of the chair.

When a set of variables is redundant, none is master. Any or all of them may be treated as being a bit less uncertain than would be the case if they were independent. Just the total uncertainty is reduced because of the structure-imposed redundancy. The component loops, however, are unaffected by the relationship between the uncertainty of their particular disturbance and that of the disturbances to any other CEV. They have to deal with what they are given, and each will deal with controlling against its own disturbance, distributing the remaining uncertainty through the rest of the control hierarchy by way of the reference values influenced by the higher-level "structural" perceptual control units. The lower-level unit has no part in the way the remanent uncertainty is distributed, but it benefits because the uncertainty of its own disturbance is reduced by the operations of higher-level control.

II.4.8 Control Tensegrity III: Hypernodes and Reorganization

In a two-level structure, the tensegrity effect is produced by the up-and-down cycling of effects between the levels. A disturbance to one lower-level control loop result in a disturbance of lower magnitude to all the higher level loops to which this one contributes. They, in their turn, alter the magnitude of their output signals which affect the reference values of the lower-level loops that they use to produce their actions in controlling their perceptual inputs. All of these effects feed back to the control loop initially disturbed, usually in the sense of reducing the stress on that initial loop.

The effect of this one initial lower-level disturbance is thus distributed over a whole set of other loops at both levels (and slightly also to the adjacent levels above and below). The control collectively performed throughout the structure results in a reduction in the uncertainty that is passed upward by any control loop being appreciably less than would be the case for the same loop in isolation. A control loop embedded in the perceptual control hierarchy is thus stronger and better able to control accurately than the same loop would be in isolation.

As we saw in Figure 8.18 (repeated as Figure II.1.8, above) when two structures overlapped, both using lower loop number 3, the uncertainty reduction is not passed uniformly through the levels. It is modularized according to the overlaps of support to the upper-level structures. The effect is analogous to using wires of different elasticity in a physical tensegrity structure. If, in Figure 8.18, we call the whole structure of “outer loop 1” (higher-level loop 1) and of “outer loop 2” local tensegrity nodes, we can see that simple loop 3 provides a link between them. That link passes some information between them so that when outer loop 1 counters a disturbance in any of its supporting lower-level loops, it passes some of that uncertainty through loop 3 for outer loop 2 to deal with.

How much? That depends on how well outer loop 1 controls, because the uncertainty in its CEV is effectively the uncertainty of its perception. Keep in mind that the structure of Figure II.1.8 is a small element within the large hierarchy, just as at the beginning of this section an isolated control loop is weaker than the same loop embedded in a perceptual control hierarchy. Relatively speaking, though, how much uncertainty is passed through (and dealt with) the two local tensegrity nodes of the Figure depends on how well the linking inner loop 3 controls its own perception, and how much uncertainty (or variation) in its reference value is created by the effects of different qualities of control in its partner inner loops, because the uncertainty of the outer loop 1 CEV is the only contribution to the uncertainty of the CEV of simple loop 3 other than its own quality of control and the uncertainty production rate of its disturbance.

Since we assume nothing about the actual structure perceived by outer loop 1, that contribution will average out so that if the simple loop 3 were undisturbed, its CEV would have an apparent disturbance with an uncertainty equal to the uncertainty of the controlled perception of outer loop 1. The residual uncertainty is the “reference” contribution, since the effect is passed through the reference input of simple loop 3, and shows up in the uncertainty of the controlled perception in outer loop 2.

If outer loop 2 has some support in common with other outer loops (upper level controllers of structured perceptions), that uncertainty will be partially passed on even further. The relationships among local tensegrity nodes can be seen as a network in which the common supporting loops serve as limited capacity links between the nodes (Figure II.1.9)

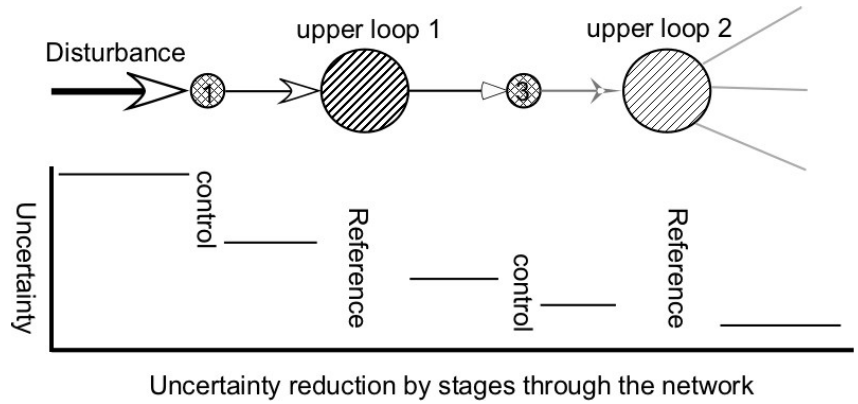

Figure II.1.9 Stages of uncertainty reduction in the two-level hierarchy of Figure I.8.18 when the only disturbance is to the CEV of lower-loop number 1. (Upper) The hierarchy as a network in which the nodes are the upper-level perceptions and the links labelled 1 and 3 are the lower-level loops. T (Lower) Reduction of uncertainty by stage of the network. The links through the lower-level control loops are uncertain due to variation in the reference values passed to them from the upper level, whereas the uncertainty passed to the upper level perceptions is due to imperfect lower-level control.

Similar diagrams could be drawn to represent the distribution of uncertainty through the hierarchy for the insertion of uncertainty at any CEV. Introduction of uncertainty to the structured CEV of an upper-level loop appears directly as informationally related disturbances to its component CEVs, and is reduced by the lower-level loops, apart from any uncertainty in the structure that remains when the lower-level perceptual values are precisely specified.

What we see is a gradual smoothing with distance through the tensegrity structure. This reduction in variation of the uncertainty due to a disturbance is due to control by intervening loops. The wire "pull" on the perception and thus the CEV results in a reduction of uncertainty each conceptual cycle between the levels.

The stress in a physical tensegrity structure is also reduced as a function of distance from the point where an external influence is imposed, but the reason for this reduction is different. If a force is applied to a point in a tensegrity structure, it is opposed by the sum of forces applied by the wires and rod that meet at that point. Each individual component changes its equilibrium force by less than the applied force. At the other end of the rod or wire, the same thing happens, and the effect of the external stressor is further reduced.

This is different only in detail from what happens in control, because when the output of the higher-level structure control unit is passed to the reference values of several lower-level supporting loops, each of them acts to reduce that uncertainty, which together they do better than they could if they alone were the only loop supporting the upper-level controller. The summation effect works in conjunction with the limited overlap to enhance the reduction of the effect of a disturbance laterally within a level through the control structure.

The chair used as an example of a two-level structure in Chapter 5 is rigidly constructed, so there is no residual uncertainty about location due to imagination or memory if the relationships among all of the components have been precisely specified and are sensed. In that case, the uncertainty left after control of the lower levels is all that the upper level must counter by its independent control action. But not all higher-level perceptions in the full hierarchy are so precisely determined by their lower-level constituents based on sensory data. The higher in the hierarchy of levels, the more likely it is that imagined or remembered values form part of the perception, and in these cases the uncertainty of the memory or of the relation between the imagined and the real world act as places where uncertainty is introduced. Figure II.1.10 shows a slightly larger two-level structure than the one shown in Figure II.1.8.

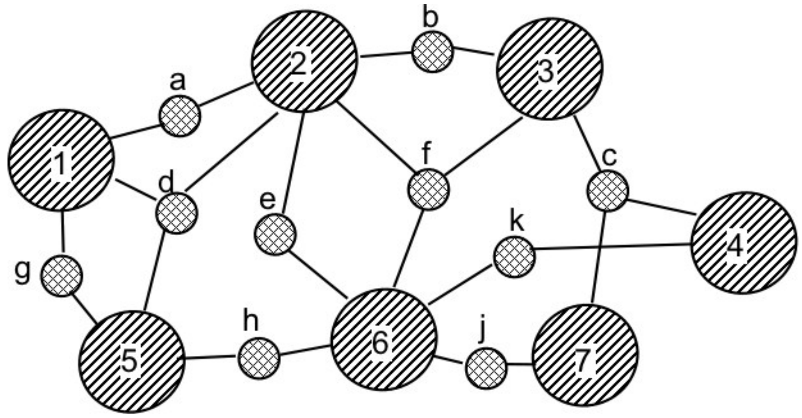

Figure II.1.10 A network of uncertainty transmission relationships between two levels of a hierarchy. Large circles represent upper-level control loops, small circles represent supporting lower-level loops. A link through a lower-level loop between two upper-level loops implies that the lower-level loop supports control by both upper-level loops, and thus transmits uncertainty between them in both directions.

As discussed above, a characteristic of a two-level network such as this is that each down-up traversal reduces the transmitted uncertainty to be dealt with by the next loop in the chain. In the absence of any perceptual structures at levels above or below these two, we can use some of the tools of Social Network Analysis (SNA) to assess the relationships among these control loops.

One tool that allows us to treat the network as a hierarchy of modules of increasing scope has been called "hypernode analysis" (Bjørke, 1996, 2003, 2006; Bjørke & Myklebust., 2001; Bjørke, Nilson and Varga, 2010). In its simplest form, hypernode analysis is a repetitive series of operations. Each single stage groups at least two nodes together according to the similarity among their connections. The grouped nodes are treated as a single "hypernode" at the next stage, and the derived network can be compacted similarly, depending on the relative similarities among the hypernodes produced at the previous stage. If we say that two upper-level nodes in Figure II.1.11 are connected when there is at least one path between them that goes through only one lower level node, then upper level loop 1 is connected to upper level loops 2 and 5 through lower-level loops a, d, and g, while 6 is connected to all the others except 1. The interconnection list before constructing hypernodes is shown in Table 13.1.

Table 13.1 Shortest path length of connections between upper level nodes in Figure II.1.10

	II.1	**II.2**	**II.3**	**II.4**	**II.5**	**II.6**	**II.7**
II.1	II.0	II.1	II.2	II.3	II.1	II.2	II.3
II.2	II.1	II.0	II.1	II.2	II.1	II.1	II.2
II.3	II.2	II.1	II.0	II.1	II.2	II.1	II.1
II.4	II.3	II.2	II.1	II.0	II.2	II.1	II.1
II.5	II.1	II.1	II.2	II.2	II.0	II.1	II.2
II.6	II.2	II.1	II.1	II.1	II.1	II.0	II.1

II.7	II.3	II.2	II.1	II.1	II.2	II.1	II.0

The corresponding similarity or difference measures may be measured in a variety of ways Bjørke used mutual uncertainty between connection patters as the primary measure, but as this is simply an illustrating example, we will use a much simpler measure that can be performed by eye. In this crude analysis, the measures in Table 13.2 are the summed absolute differences between the values in the rows corresponding to the upper-level loops being compared. Smaller totals imply more similar connection patterns, so the total of the absolute differences in number patterns represent degrees of dissimilarity between two nodes in their pattern of connectivity.

Table 13.2 Dissimilarities of connection pattern among the different upper-level nodes.

	II.1	**II.2**	**II.3**	**II.4**	**II.5**	**II.6**	**II.7**
II.1		II.6	II.13	II.13	II.5	II.9	II.13
II.2	II.6		II.6	II.8	II.3	II.5	II.8
II.3	II.13	II.6		II.4	II.7	II.3	II.5
II.4	II.13	II.8	II.4		II.9	II.5	II.2
II.5	II.5	II.3	II.7	II.9		II.6	II.9
II.6	II.9	II.5	II.3	II.5	II.6		II.5
II.7	II.13	II.8	II.5	II.2	II.9	II.5	

At this point, we choose a threshold for considering nodes to be part of the same hypernode. There are several dissimilarities of 5, but few below 4, so 3 and below seems a reasonable threshold[11]. This places 4 and 7 together as a first hypernode, then 2 and 5 plus 3 and 6. The corresponding hypernode structure is shown in Figure II.1.11.

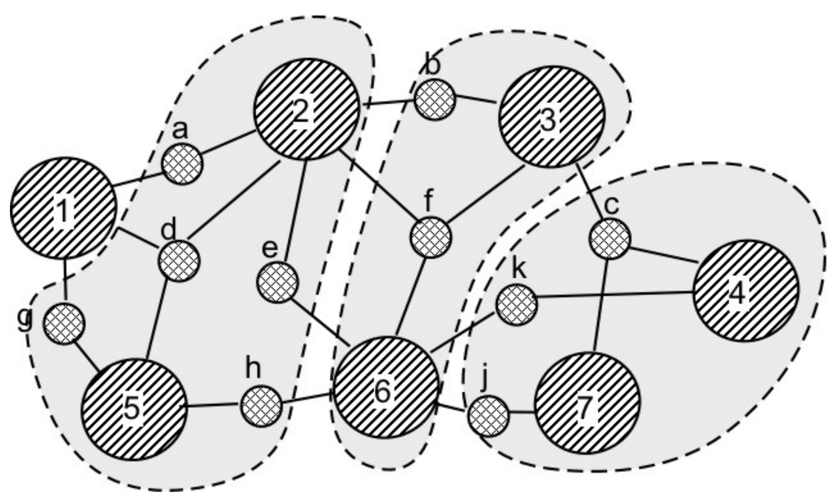

Figure II.1.11 The first hypernode compression for the network of Figure II.1.10. There are four hypernodes, as the upper-level control loop "1" forms a hypernode on its own.

Why should we be interested in Bjørke's hypernodes? There are two reasons. Firstly, loops that are

11. In a large network, there would be more numbers, and since in a proper analysis the similarities would be computed using mutual information, they would be distributed over more of a continuum. The hypernode structure corresponding to a selected threshold can be displayed and the threshold changed interactively to suggest values useful for some purpose.

encapsulated into hypernodes tend to support each other's control better than they do loops outside the hypernode. The reason is that they have more supporting lower-level loops in common than they do with upper-level loops outside the hypernode; supporting lower-level loops are the medium by which the uncertainty "stress" is distributed through the structure from any specific source of disturbance.

Secondly, because the higher-level loops within the hypernode have more supporting loops in common, their perceptions considered together are likely to correspond to as yet unperceived structures in the environment, which makes it likely that reorganization will pull the perceptions within a hypernode together into a controllable perception at a yet higher level. The hypernodes themselves are not that structure, but the commonality among their supporting lower-level loops enhances the likelihood that their variations will be correlated. If that informational interdependence carries through to other high-level control loops not shown in the diagram, the arguments in Chapter 10 increase the likelihood that the combination will be perceived first as a syndrome, and later reorganized to become a structured perception within the control hierarchy.

How would this affect the hypernode structure, if we carry it through several levels of the control hierarchy? Figure II.1.12 suggests a multi-level projection of three high-level structured perceptions that were built by reorganization because their components turned out to be informationally related.

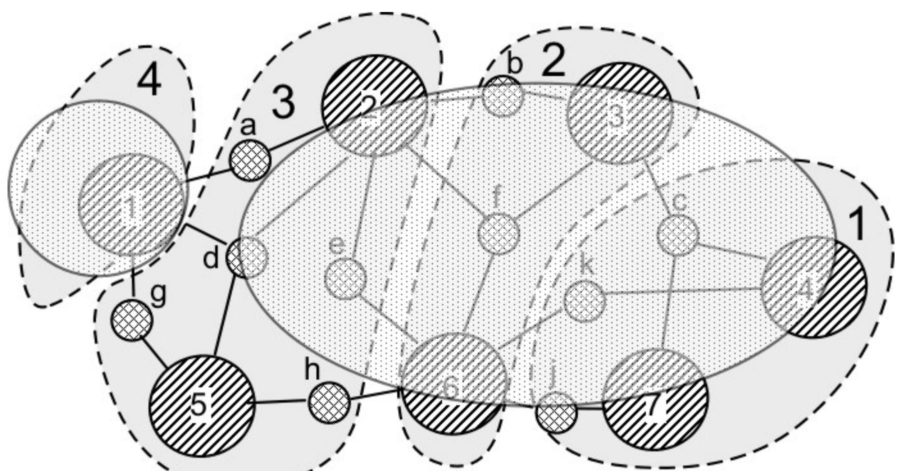

Figure II.1.12 Two higher-level control loops built on the basis of the hypernode structure of the network of Figure II.1.11. The leftmost top-level structure (large translucent circle) controls only through second-level node 1, which is supported by lowest-level loops a and g, whereas the other (the ellipse) uses second-level loops contained in three of the first set of hypernodes, and is ultimately supported by all the bottom-level loops except g. The leftmost top-level control is performed by habit, in that it uses only second level loop 1, which on its own is hypernode 4.

In an organism, the low-level controlled perceptions have been tuned by billions of years of evolutions, and almost certainly represent properties of the real world environment. The same is true of object perception[12], if "almost certainly" is replaced by "very probably". As we go higher and higher in the control hierarchy, the probability that the real world works the way we perceive it to do becomes smaller and smaller, until we reach the domain of scientific enquiry, a domain that is often presented to

12. Later, we will argue that there is no object perception in the non-conscious perceptual control hierarchy, and that the perception of objects as such is purely conscious.

the public as describing facts of nature, but which depends for its very existence on ignorance of the actual facts of nature.

We can see this change of certainty as we go up the levels in the hypernode structure. Unlike the neural synaptic connections that collectively provide the circuitry of control, hypernodes are a purely analytical convenience, helping us, taking either the Analyst's or the Controller's viewpoint, to distinguish clusters of controlled perceptions that seem to behave in an informationally interdependent way more closely within a hypernode than between hypernodes.

The divisions between hypernodes are far from distinct, but nevertheless, we may expect them to illustrate regions of Hebbian reorganization within a hypernode and anti-Hebbian reorganization between hypernodes, which, in coordination with e-coli reorganization would enhance the modularity of higher-level perceptual structures. The effect would be similar to the lateral inhibition effect that was the topic of Chapter 9.

We will see it again between independently initiated autocatalytic loops that do not share reactions, which might merge or might repel each other. When we consider rattling in Chapter II.5ff, we will see it yet another context, the relative calming effect of separation between locally stable parts of the organization, such as distinct autocatalytic networks, which we discuss in the next Chapter. These tendencies toward separation are akin to tensegrity rods in the organizational structure, which oppose the variety of attractive tendencies inherent in perceptual control, not least in the tendency of creative autocatalytic loops to shorten that we discuss in the next two Chapters.

Chapter II.5 "Rattling": Process and Measure

What is "rattling"? In everyday language we know what it means. Something intolerable happens unexpectedly and we have no equally quick means of correcting the resulting problem. We feel "rattled". This "intolerable" event might originate from our internal thoughts, such as remembering that we forgot to turn off the stove when we are on a plane out of town, but more probably we are rattled by some external event that causes uncorrected error in one or more controlled perceptions.

Perhaps we observed an accident on a roller-coaster when we were waiting our turn to ride it. Perhaps we intended to prepare a special dish but find an essential ingredient to be out of stock at the grocery. Most people would agree that they were "rattled" under such circumstances, though some placid people might not. These "placid" people tend to have wide tolerance bounds, and control perceptions above the "rattled" level simply by doing something else. I suspect most of us have been rattled at some time or other. What Chvykov et al. (2021) did was to propose a computable measure of rattling that applies to organizations of elementary entities that can act in ways that influence each other,

The measure of "rattling" proposed by Chvykov et al. is in the same domain as a measure of variance or of uncertainty (Chapter 10), both of which apply to the variability of a set of measurements. The measurements used in computing variance might be of the same variable over time, of a collection of different variables at the same moment, or of a collection of different variables over time. The same is true of the the variables used in the rattling measure. In this case, the variable whose variance or uncertainty is computed is not the values of the variables, but the velocities with which those values change.

Chvykov et al. did not propose variance over those change velocities, because the variance measure is valid only when the variable in question has a "Normal" (Gaussian) distribution, and they wanted to include the sudden velocity changes cause by impacts. The uncertainty measure is appropriate where the variance measure is not, although they are identical if a distribution happens to be Gaussian. So their "Rattling" measure is the uncertainty of velocity changes of the values over a collection of variables.

In their experiments and simulations the variable of interest to Chvykov et al. was velocity of change in the location of an active unit they called a "smarticle", a number of which were in a region bounded by an impassable boundary. Any one smarticle could hit another and cause it to move, but it could not move autonomously. However, the region was small enough that not all the smarticles could be out of range of all the others at any one time. The organization in question was the pattern of relative locations of the set of smarticles.

We will be interested in changes in perceptual or reference variable of a control loop, and control loops will affect those variables by the mutual disturbances caused by side-effects. Since only error variation affects the action output of a control loop, we will not worry about either perception or reference individually, though we could, but will deal only with the velocity of change in the error variable for the loop. We consider the error variable to be not the difference between reference and perception, but the degree to which that difference exceeds the tolerance bounds for the loop. Its contribution to the rattling measure is the uncertainty over time of the velocity of change in the error variable.

What Chvykov et al. noted was that over a whole collection of interacting entities, the inter-relationships among the entities — their organization — would change as a result of their interactions. The greater the rattling, the greater the average change rate. They used a thermodynamic analogy, that the faster the molecules move on average, the higher the temperature, and heat moves from hotter to cooler, not the reverse. The metaphor suggested, and Chvykov er al. showed mathematically, by simulation, and by physical small-scale illustration that organizations of self-energized entities with no energy input from outside tended to move from high-rattling to low-rattling states.

You might well ask what the rather esoteric rattling measure has to do with Perceptual Control. Rattling describes a process and a measure that will apply to any organization of active entities that can influence each other, which control loops do by signals that can alter each others' disturbance values or the reference values. Disturbance values in the control of one perception can be increased by side-effects of the output actions of other controlled perceptions. Reference values are altered by the outputs of higher-level perceptual controls that are acting to stabilize against their own disturbance changes. The rattling measure could be applied to both.

Rattling, both as a process and as a measure, applies directly to the organization and the reorganization of the perceptual control hierarchy. It provides another look at reorganization that augments what we have discussed in various ways and will discuss further. But it does more, since it applies to any organization of active elements, including social organizations of active organisms (living control systems), from bacteria to humans, and complex ecologies of living things. We will treat the rattling of social organizations in Volume 2.

In the next Chapter we will investigate a closely related process, "crumpling". We mention it here because it has a tensegrity relation with rattling that affects the reorganization process in the control hierarchy. Crumpling is an active process when we crumple up a sheet of paper and throw it away, but we will not be discarding crumpled control loops. We will not deal with the process so much as with the results of the process found when we examine the crumpled paper.

We will use crumpling largely as a metaphor for perceptual processes that belong to the same rattling-arranged organization of control loops. Although one must never take a metaphor as exactly representing the analogue — as Korzybski (1933) said: "The map is not the territory" — nevertheless, a good metaphor is often helpful in understanding the analogue. Crumpling turns out to be in many ways a good metaphor for the development of novel category perceptions and their relationships within the control hierarchy.

In the rest of the book, especially in Volume 3 and Volume 4, the rattling and crumpling concepts and measures over an organization will be increasingly important, though they will in may cases remain implicit. This Chapter introduces and explains the rattling concept, and uses it as a different approach to several of the topics we have previously discussed, particularly those related to individual reorganization of the perceptual control hierarchy. The same is true of Crumpling, which we ignore until the next Chapter.

II.5.1 Quantifying Organizational Rattling

The entities treated by Chvykov et al. were not control systems, but they can move if they are affected by external forces and are not physically constrained. How they move is irrelevant except that their internally generated energy may or may not influence the movements of neighbouring entities, just as the movements of underground rocks in an earthquake influence buildings at the surface. The authors call their particular examples of interacting entities "smarticles". A smarticle moves only when touched or hit by the waving arms of another smarticle. In the cited paper, smarticles can move freely over a bounded domain, bounded so that they cannot all disperse out of range of being hit by the actions of other smarticles. Smarticles are not influenced by external influences from outside their restricted domain. Our control loops and thus their disturbance and error variables are influenced by outside influences as well as by the side-effects of controlling other perceptions.

To illustrate the principles of their largely mathematical paper, Chvykov et al. built a small tangible set of actual smarticles and recorded their motions over the floor of a circular "cage". They measured the total rattling change over time of the set of smarticles, which approximated a declining exponential. They also simulated a much larger set of interacting smarticles, and found a smoother decline over time of the same form. The decline was smoother because of the averaging over a larger number of rattled smarticles.

To measure a change in a variable such as a amount of rattling, the variable needs to be defined. As we discussed in the introduction to this Chapter, Chvykov at al. define the quantity of rattling of some variable as the uncertainty (Chapter 10; they use the word "entropy") of the velocity of change of that variable. The uncertainty could be of the velocity of the variable over time for a single entity, or over a collection (an organization) of similar entities at a given moment, or over a collection over time.

As pointed out by Shannon (1949), uncertainty is additive. If variables X and Y have values {x1, x2, x3, …} and {y1, y2, y3….} independent of each other, then the uncertainty over many measures of the collection {X, Y} is the sum of the two uncertainties of the xn and yn sets of values.

If the distributions are Gaussian, then their uncertainties are directly proportional to their variances. If they are not mutually independent, then the variance or the uncertainty of the combination of X and Y is reduced below the simple sum of variances or uncertainties by their covariance or "mutual information". Rattling of a large collection, an organization, of entities is similarly additive, with reductions related to their mutual influences. (Chapter 10, Shannon, 1949, or Garner, 1962).

Within an individual, the variable being rattled changes with the analyst's viewpoint, but often it is a perceptual variable or the environmental variable corresponding to a perception. Sometimes it is the error variable in a control loop or the output variable, other times the output at the muscular interface to the environment. In other words, any quantifiable variable can have a rattling measure just as it can have a variance measure or an uncertainty measure.

In what follows, we will often talk about variability of variance as a guide to whether rattling is intense or weak. The rattling measure offered by Chvykov et al. (2021) is the uncertainty of velocity of change in whatever measure or measures are under consideration. As we have seen earlier, uncertainty and variance in a continuum measure are interchangeable conceptually and possibly numerically. Hence it is quite legitimate to use variability of variance (or equivalently variance of variability) as a conceptual measure of rattling intensity, as we shall often do.

What does this have to do with the Quality of Control (QoC), which is usually measured by the variance or uncertainty of the error variable? Rattling is measured by the uncertainty of change velocity of the error, the derivative of the error. The derivative of a waveform is basically uncorrelated with the waveform itself, so we should at first sight assume that the rattling measure would be uncorrelated with the QoC. This, however, is not so, because the non-correlation of the waveform with its derivative depends on the Fourier analysis that decomposes the waveform into sine and cosine components, and their derivatives are cosines and sines, orthogonal to (uncorrelated with) their original sines and cosines.

The Fourier analysis, however, also incorporates amplitude measures, and if the original waveform is amplitude-modulated, the derivative's amplitude is modulated in lock-step with the modulation of the waveform. The error variance is a measure of the modulation of the error over some period of time, and the uncertainty in the velocity of that amplitude variation is the rattling measure.

When the error value has rapid excursions, control is unable to keep up. Putting it another way, the negative loop gain of a control loop cannot exceed unity if the loop transport lag exceeds half a cycle of the highest frequency in the disturbance. High rattling corresponds to poor control because the uncertainty of the velocity of change in the error depends on the QoC, which is inherently poor when the disturbance changes more rapidly than the control loop can handle.

What constitutes a rapid excursion in the disturbance depends on the level of concern in the perceptual hierarchy, since the upper levels inherently have longer loop transport lags than do the lower levels on which they depend. Accordingly, it could happen that the rattling experienced at some low levels might be more than at higher levels, or it could be lower. There is no necessary correlation between rattling and the regions of the hierarchy experiencing high rattling. As we shall see, the implication is that reorganization is more likely to change parts of the hierarchy experiencing high rattling than low, which corresponds to Powers's intuition that locally poor QoC leads to increased reorganization rates.

One last part of the scenery needs to be set, namely the stage itself, a high-dimensional descriptive space. Imagine a minimal organization that consists of two unit entities, each able to affect the other by a link to which we can assign a weight in the form of a number between plus and minus infinity. There are two one-way links, and their weights completely describe the entire organization. The two numbers, the link weights, could be labelled Wxy and Wyx, and the total description of their organization could be specified by a single point in the X-Y two-dimensional plane. If the effects are mutual, each having as much effect on the other as vice-versa, then the organization lies somewhere on a diagonal represented by Wxy=Wyx.

A different descriptive space could be built from these same two variables, X and Y. At a given moment, the variables have values x and y. Their joint value is a point at (x,y) in the X-Y plane. Whereas Wxy and Wyx describe the quasi-permanent structure of the organization, x and y describe its behaviour from moment to moment. Even if Wxy=Wyx, x and y could vary in any way at all, even chaotically following "strange attractor" orbits, depending on the processing being done within the nodes designated X and Y. If either or both of X and Y were subjected to external influences, x and y could vary in any way at all, while Wxy and Wyx remained invariant. The plane (X, Y) is a descriptive space, but there are choices about what to describe in it.

Consider a minimal control loop consisting of a perceptual input function that works as a simple pass-through connector, a comparator that produces an error value that ranges between plus and minus infinity depending on the difference between the instantaneous values of the reference input and the perception, and an output function that is a leaky integrator with a gain rate G and a leak rate L. The loop also has a transport lag T. The structural parameters of this trivial control loop are completely described by the values of g, l, t, of G, L, and T respectively, provided that the form "control loop" is assumed *a priori*.

The organization of this particular control loop can be described as a point in a three-dimensional space of G, L, and T onto which g, l and t, the values of the three parameters G, L, and T respectively, can be mapped. The Quality of Control (QoC) is sometimes measured as the ratio between the error variance and the disturbance variance. The QoC of a control loop depends on the values of g, l and t, so for every point in the 3-D GLT space there is a value of QoC.

This very basic control loop has another 3-D descriptive space, that of the signal values at any moment. Because of its basic nature, specifically that the perceptual value equals the CEV value, the loop has three different signals at any moment, Perception, Error, and Output with their corresponding values p, e, and o. These are related through the location of the structure in GLT space, and are affected by momentary and historic values of the reference and disturbance inputs. In any particular organizational context, these latter would rattle p, e, and o.

If we add another parameter to the loop description, such as a tolerance bound B, the space of description of the loop becomes four-dimensional. If we have a trivial hierarchy with two such loops, one receiving perceptual values from and sending output to the other, the entire organization has four dimensions to describe each of the two loops, plus two dimensions to represent the weights of the perceptual and reference links, making the descriptive space of the organization ten-dimensional. As we add more control loops in a working hierarchy, the number of descriptive parameters quickly gets very large. The descriptive space is of correspondingly high-dimensionality.

The whole organization is completely described by a single point in the space of description of the structural variables within it, however high its dimensionality, and associated with that point is a QoC value for the organization under external rattling influences. The actual measure of rattling depends on the movement velocities in the different dimensions of the descriptive point in the equally high descriptive space of signal values, but those velocities do depend on the structural descriptive point, and changes in the structural description will change the rattling values for any external influences.

Chvykov et al. (2021) show that the descriptive point of any organization will tend, over time, to move downslope towards regions of the space where the rattling is less, in preference to moving upslope to where rattling is higher. We will show that the QoC value of the structural point is the most relevant measure. We may call lower-rattling regions "calmer" and regions with high rattling "tempestuous". For an organization to move of its own accord toward calmer regions of the description space is self-organization. When applied to a PCT control hierarchy we call it "reorganization".

In this Volume 1 the collection or organization of primary interest will consist of a perceptual hierarchy, with the interconnected ECUs as the entities. In Volume 2 (beginning with Chapter III.9) the entities will usually be individuals or groups of humans or other social organisms, rather than structures within the individuals. The rattling measure in both volumes is usually a sum over a collection of entities or a collection of variables, or both.

II.5.2 Leaves and Smarticles

To help the reader to visualize rattling, we use organizations of entities even simpler than the "smarticles" of Chvykov et al. (2021). The entities we consider in this Section are purely passive, including leaves, snowflakes, sand grains and the like that can be moved by the wind. How they are rattled is one component of how the more complex self-energized entities such as smarticles or those of our central interest are rattled. From this viewpoint, rattling theory can be seen as a generalization of what we called in Section 11.1 the "winter-leaf" phenomenon, in which dry autumn leaves blown around by variably gusty winds from all directions move preferentially from places open to the winds into relatively sheltered areas where they may settle.

As we described above, the rattling measure depends not on the variance in direction and intensity of the wind, but also on the variability over time of these variances. Calm days give the leaves time to settle down into fairly stable configurations, while stormy days move all the leaves except those near the bottom of the most sheltered leaf piles. Variation of wind variance over time is as critical as are the simple variances of wind direction and intensity. This temporal variability is an essential component of the rattling measure.

Steady laminar winds can move airborne leaves without changing their spatial organization (upper sequence of Figure II.5.1), while non-steady winds change it (lower sequence). The entropy of the spatial organization does not change in the steady flow, whereas it continually increases in the turbulent flow. After long enough in the turbulent flow, the organization of the leaves will be even more widely distributed and uncertain than it was when the trees dropped them. How rapidly the entropy of the spatial organization increases will be variable, fast during gusty periods with winds from changing directions, slow when the air is in one of its calmer moods, with momentarily gentle breezes from a constant direction.

In measuring the rattling forces applied to the organization of winter leaves, what matters is the *variability* in the amount and direction of supplied power. How that affects the distribution of leaves into piles at long-term equilibrium depends on the configuration of wind-shadows. Some of these wind-shadows are created by static environmental irregularities, some by existing leaf piles. It is not the total externally applied energy alone that matters, but its variability for a given average power level (power is the energy per unit time — the rate of energy application, not the amount of energy applied).

Tying the winter leaf metaphor to the velocity uncertainty measure of rattling presented by Chvykov et.al, we can see areas of wind-shadow as regions of relatively low rattling, whereas in open spaces they experience high rattling powered by variable winds. In the Chvykov et al. (2021) paper, the elementary members of their group, which they call "smarticles", are internally powered and "rattle" each other, whereas in our winter-leaf metaphor the variable wind power, an external energy source, is what rattles the leaf.

Rattling totalled over an entire organization is a generalization of the winter-leaf phenomenon. The leaves are passive, and don't directly influence each other except than one leaf may shelter another from the wind or knock it from a calm place into a windier place. Their individual rattling values for location are essentially independent and therefore additive. If one leaf shelters another, they tend to move in a correlated fashion together more than they would if they moved separately. A whole pile of leaves shadows an incoming loose leaf more than does any one other single leaf. Leaves at the bottom of a pile are seldom rattled at all.

The "rattling" environment of self-energized but non-controlling entities like the smarticles of Chvykov et al. has no such static regularities (except in their experiment, the boundary fence that limits the ability of the smarticles to escape each other) to influence the motions of the different entities. They can reduce their overall rattling only by changing the ways and the degrees to which they influence each other. In the case of the smarticles, only the location of a smarticle is changed by an interaction with another smarticle.

Chvykov et al. (2021) generalize this basic idea from the three-smarticle physical demonstration to higher-dimensional spaces representing more smarticles. The energy of a “smarticle” at any moment consists of energy from its internal source and energy supplied by its recent interactions with other smarticles. Compare this with Boltzmann's analysis of the energy distribution among the particles (atoms) of an ideal gas caused by their bumping against one another, leading to the principle that an enclosed gas would tend toward an energy distribution of maximum probability — maximum entropy. In this distribution, most atoms move slowly, while there always a few moving fast, in an exponential distribution.

Chvykov et al. prove mathematically and demonstrate both physically and in simulation that in the case of mutually influencing active “smarticles”, the distribution of the velocity uncertainty of “rattling” trends in the same way. To make sense of this statement, we must consider the velocity uncertainty as being measured not only across dimensions of the description space for a rattled unit, but also over time within each dimension. The rattling experienced by a unit is the total dynamic entropy over some time span combined over all descriptive dimensions for the unit. We will describe this differently in Figure q17.3 below.

In the asymptotic steady state, for a given average rattling value, most experience low total rattling while there will always be a few that are strongly rattled. In the words of a common aphorism, you can please some of the people all of the time, and all of the people some of the time, but you can’t please all of the people all of the time. Chvykov et al. would suggest that you can’t even please all of the people any of the time.

For a given total energy in Boltzmann’s ideal gas, entropy is maximized by a declining exponential distribution of kinetic energy over the atoms; in the case of the leaves and the smarticles, total rattling is minimized by reorganization. Boltzmann’s gas has no organizational structure. What Chvykov et al. suggest is that even when the collection of elements has structure, within that structure the distribution of velocity variation is of the same character as is the distribution of kinetic energy in Boltzmann’s bumping ideal atoms.

As Chvykov et al. point out, these asymptotic distributions are achieved only in the infinitely long run. They, like us, however, are interested in the dynamics of how this asymptotic equilibrium distribution may be approached. For their smarticles, the approach of total rattling over their spatial organization to an equilibrium steady state is similar to a declining exponential.

We will return to the rattling approach to social self-organization in Volume 2, Chapter III.10. In this Chapter, we will concentrate on the perceptual control hierarchy within an individual.

II.5.3 Turbulence, and Group Self-Organization

As we have seen, the entropy of a group of bodies in an isolated system tends toward maximum entropy as the individual molecules jostle against one another. Entropy, however, tends to be minimized by self-organization if some can be carried away by a through energy flow to a wider environment. Now we see self-organization also as a reduction in rattling, a dynamic cousin of entropy. Self-organization, according to Chvykov et el. is a result of reducing the total rattling of one entity by another averaged over the group of interest.

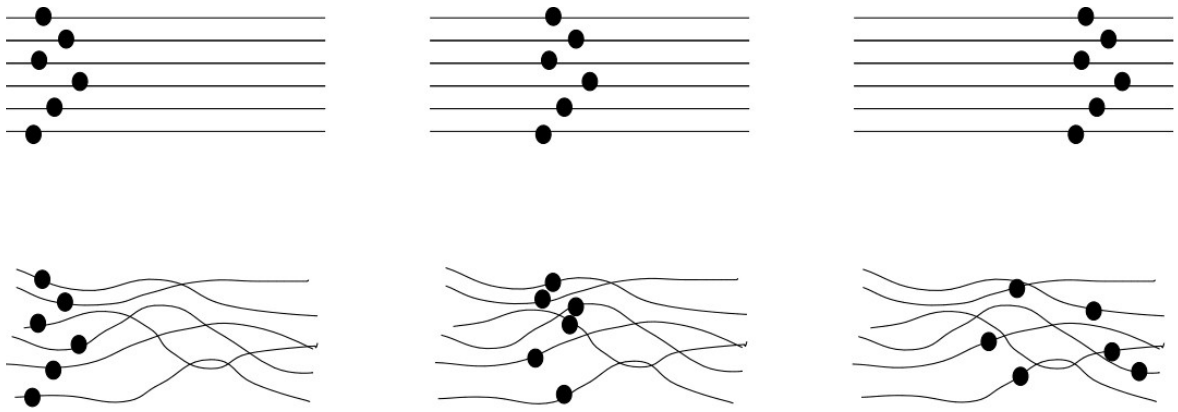

Figure II.5.1 The movement of a set of leaves in (upper) a smooth, regular, wind in what is called a "laminar flow", and (lower) the same set of leaves in a turbulent wind. In the laminar flow, the pattern of relationships among the locations of the individual leaves does not change. The organization of the pattern is calm and unrattled, and thermodynamically cold. In the turbulent flow the pattern changes from moment to moment. Considered by itself, the pattern n a turbulent flow is highly rattled and thermodynamically hot.

Figure II.5.1 illustrates that the rattling of the organization of the leaves does not depend on the strength of the wind, but on its character. If as a child you ever played "pooh sticks", you have observed the phenomenon. To play pooh sticks a group of you stand on a bridge over a stream or river and simultaneously drop small sticks onto the flow. You then race to the downstream side of the bridge and see whose stick emerges first. However carefully you synchronize the fall of the sticks, there are always some that move fast and some that move slow, depending on what parts of the stream flow they happen to hit.

The depiction of turbulent flow in the lower panels is far from exact. Remember Richardson's little 1922 ditty, quoted in the introduction to this Chapter:

> *Big whirls have little whirls*
> *that feed on their velocity,*
> *and little whirls have lesser whirls*
> *and so on to viscosity.*

When Richardson says that "Big whirls have little whirls", those little whirls are generated by shear stresses at the boundary of the big whirl, but that's not all that can be said about them. The big whirl carries them along with its circular flow as it itself is carried by the main stream flow. Figure II.5.2 may suggest this process, again without being at all literally accurate. Each little whirl is generated by the shear of the big whirl's flow against the stream flow in which it is embedded, at all scales. The littler whirls may spin off into the local general flow or may be captured as part of the larger generating whirl. The metaphor here is that of organizations and sub-organizations that may spin off into independent entities or retain membership in the larger organization.

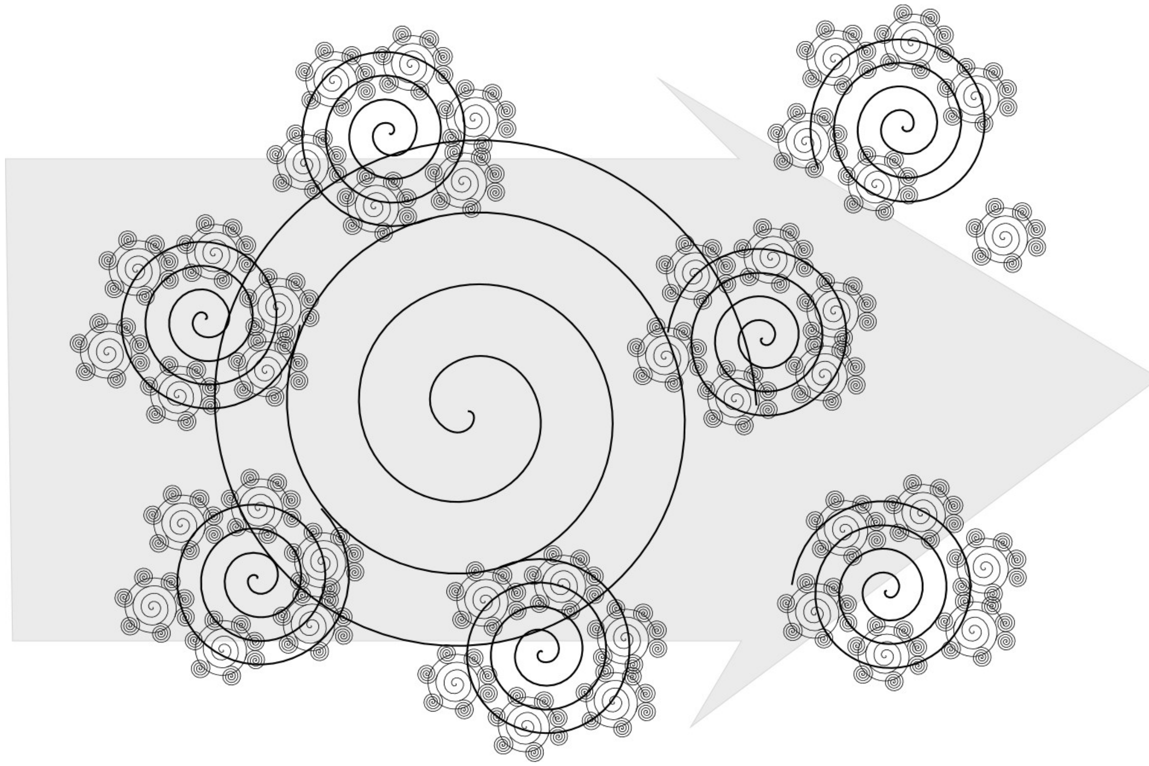

Figure II.5.2 "Big whirls have little whirls/that feed on their velocity". Most are carried along with the swirling movement of their generating bigger whirl, while at all scales others may break off and independently whirl away in the general flow.

In Chapter xII.1, we noted Boltzmann's 1877 proof that the temperature of an ideal gas confined in an isolating container was proportional to the entropy, where entropy was determined by the uncertainty associated with the movements of the atoms of the gas. The more precisely the relative locations and kinetic energies of the gas molecules could be known from one moment to the next, the colder the gas. Figure II.5.2 again suggests a temperature difference between turbulent and laminar flow, while at the same time showing how the low entropy energy of the main flow is transformed into high entropy energy in the form of waste heat.

Figure II.5.2 also suggests how the turbulent flow will contain pockets of high-velocity flow, where the local velocities of several scales of whirl add together in the same direction, together with pockets of relatively calm, where the local velocities add in opposition to each other to sum to near zero. Again, this collaboration and opposition between whirl velocities recurs at different scales. Just as "Big whirls have little whirls", so too may big relatively calm regions have within them even calmer smaller regions. Everywhere, the local variances change as you move from one region to another or as you wait for a while in one place. The space-and-time shape of the "whirly world" is variably variable everywhere and everywhen. Turbulence rattles.

If we apply Boltzmann's insight to the organization of the leaves, we note that in a laminar flow, if at any moment you know the location of one leaf and the relations among the leaves in the original pattern, you know where all the other leaves will be at that moment. The organization is thermodynamically cold (upper panels of Figure II.5.1) and is not being rattled even though the linear flow may be very fast. In the language we have been using and will continue to use, the leaves in a laminar flow are "calm" as a group, no matter how strong the wind.

The situation is quite different in a turbulent flow. In a turbulent flow, the entropy of the organization grows quickly as time passes. The leaf pattern is thermodynamically hot (lower panels of Figure II.5.1) and the organization is experiencing rattling. How "hot" it is depends on the strength of the flow, which drives the turbulence and therefore the rattling of their organization.

A phase change will occur in a flow as the flow gets faster. A measure called the "Reynolds Number" increases past a threshold that is like the freezing or boiling point of a liquid, in that if the number is passed, a laminar flow will become turbulent. We might ask whether the explosively increasing rate of autocatalytic innovation might have some equivalent phase change threshold in rattling theory. We have already seen the phase change described by Kauffman (1995) that occurs when catalytic processes link up in a loop, but that phase change has no time scale.

At the moment, I have no answer as to whether there might be a rattling-related phase change with increasing velocity of innovation, but I suspect that if there is, the effects will be localized, just as are the stormy and calm regions in a turbulent flow of wind. Some parts may for a while sustain laminar flow, while others break off in the general rattling turbulence[13]. In "rattling language", some parts may remain calm while others are strongly rattled by rapid innovation.

We will use this "calm within turbulence" analogy in Volume 2 to show how a revolutionary organization that collectively acts to rattle greatly the organization of the culture to which its members belong can at the same time be internally calm and very little rattled, despite acting very violently to rattle other parts of the larger organization, as was demonstrated at the US Capitol during the certification of the election of Biden to the Presidency of the USA on January 6, 2021. The same may happen among different parts of an individual perceptual control hierarchy, being manifest not only in the form of locked conflict that is addressed by the Method of Levels (MoL), but perhaps in mental aberrations we might call paranoia, schizophrenia, or multiple personality disorder. We will go no farther with this kind or speculation, but the possibility should be kept in mind.

As discussed above, in a turbulent wind, periods of extremely strong gusts mingle with periods of relative calm. This would be true of the wind direction as well as of its strength. Calm periods in wind direction (near-laminar flow) may occur even when the wind strength is high, and the wind direction can be highly variable even when the wind strength is negligible. The two aspects of wind variability are not highly correlated.

No matter which wind property you consider, if you compute its variance over moderately short time intervals, that variance measure will be variable and the rate at which it changes will have some uncertainty. To a certain degree that uncertainty of the velocity of variance change is a statistical consequence of failing to average over long enough intervals, but in the turbulent flow there is an excessive variation in the variance beyond the statistical expectation. It is this excess that rattles the organization of the leaves carried in the wind.

Figure 11.6, which describes a possible control loop that reorganizes the parameters of a perceptual control loop, uses exactly the kind of short term variance averaging I mean. The reorganization controller in that Figure controls its perception of the ratio of two varying variances to optimize the control loop parameters. The winter-leaf effect relies on the variation of the external power vector (direction and magnitude) of the winds relative to the static irregularities of the environment, such as walls, gullies, and hedges. We can say that the wind variation "rattles" the leaf, moving it hither and yon until it arrives at a

13. I recognise that such metaphors can be carried too far, and this metaphor will be reined in as we discuss rattling both in this Chapter and in social conditions in Chapter III.9 and beyond.

calmer place.

Rattling is directly related to another measure, diffusivity. Diffusivity is normally applied to molecules and particles moving in a medium. A particle in a region of high diffusivity is more likely to be found a short time later far from where it started than the same particle would be if it were in a region of low diffusivity. Over an entire space, a diffusivity value is associated at any time with every point in the space, no matter the dimensionality of the space. Figure II.5.1 showed the effects of organizational diffusion as the individual particles go their own ways in the turbulent flow, but remain in a static pattern in the laminar flow.

The whole space can therefore be mapped (in a two dimensional analogy) as a diffusivity landscape with highs and lows and slopes in different directions. In higher dimensions, the same applies but is harder to visualize. However many the dimensions, the particle is more likely to move downslope than upslope, and most likely to move in a direction nearer to the steepest downslope than in other directions.

Before we go too much further, we should be clear about what is meant by "a high-dimensional space of description" for an organization.

II.5.4 Descriptive Spaces: Interior and Exterior Rattling

We saw self-organization in Chapter II.4, in the work of Prigogine and collaborators. There, a through flow of energy carried away entropy and dissipated it as heat. Organizational structure, as we discussed in Chapter 10, is the reduction of uncertainty about one part of an organization by knowledge of another part. The more and stronger the structure, the lower its entropy, the mutual uncertainties among its parts.

The space of GLTB plus parameter values for interconnections among the loops describes the current organization of simple control loops as a whole. But this is not the only description relevant to the organization. Rattling is a moment-by-moment dynamic quantity that depends not only on the current slowly-changing organization, but also on the "strength and direction" of the current environmental winds, to use the winter-leaf analogy once more. If there is no disturbance to a control structure, it is not being rattled, no matter what its organizational structure.

To follow the variations in signal values that will provide the required measure of organizational rattling we need a different space of description, in which the dimensional values are not values like gain rate, leak rate, transport lag, and tolerance bound, but values like disturbance value, perceptual value, error value, and output value. These values map one-to-one directly onto the connections in the organizational structure description.

It is probably easier, however, not to imagine them being incorporated as added dimensions in the structure description, but to consider a separate dynamic description space in which each control loop has a time-varying four-dimensional description consisting of signal values in the four links of the basic loop.

Of these four different signal values, we need take the rattling measure from only one. But which one? When we consider the mutual influences of control loops through direct effects or side-effects, the environmental values in Real Reality are what matters. Especially for side-effects, when we consider the cross-effects that occur within an organism, we need to consider output values.

The output values of control loops and the resulting direct and side-effects through the environment on other control loops correspond directly to the hits by the flailing arms of one smarticle that move other smarticles. Error values and perceptual values, in contrast, have no direct influences on other control loops. They are interior to the control loop. The perceptual value corresponds to the value of the CEV in Real Reality (and is exactly the value of the CEV in Perceptual Reality), so its values are not a very good measure of how rattled the loop is, inside the organism.

The key rattling measure of all the interior possibilities is likely to be based on the variation in the error values of different perceptions. Each perception defines a dimension in a high-dimensional space, and all the values in the control loop related to its control can be plotted along this axis including not only their values but also their rates of change of value — their velocities.

We have two different measures related to rattling associated with each control loop that perceives an external environmental variable, *experienced* or *internal* rattling based on error values (differences between perception and reference that exceed tolerance bounds) and *active* or *external* rattling based on output values that might affect other control loops.

Most people find it difficult to visualize a space of more than three dimensions, let alone a space of hundreds or thousands of dimensions that might be needed to incorporate all the variable perceptual dimensions with their error values in all the perceptions at all the levels of a perceptual control hierarchy. To imagine the entire organization as a single point in this high dimensional space, one does need to think of this space, though perhaps only in principle, not in detail.

An alternative way to represent them is as a vector, each element of which represents the value of the variable on one of the dimensions. To see the relationship between the vector and the space, consider a two dimensional description space — a graph — in which a vector of two values (x, y) represents a single point in the space. Likewise, a point in three-dimensional space can be represented by three values, the vector (x, y, z). We can continue this into a vector representation of a point is a space of any number of dimensions by a vector (x1, x2,,x3,...xn). When we want to measure the rattling over the high-dimensional space at any moment, we can find the uncertainty of the distribution of vector element values, if the vector represents velocities of change of the elements. Figure q17.3 suggests such measures.

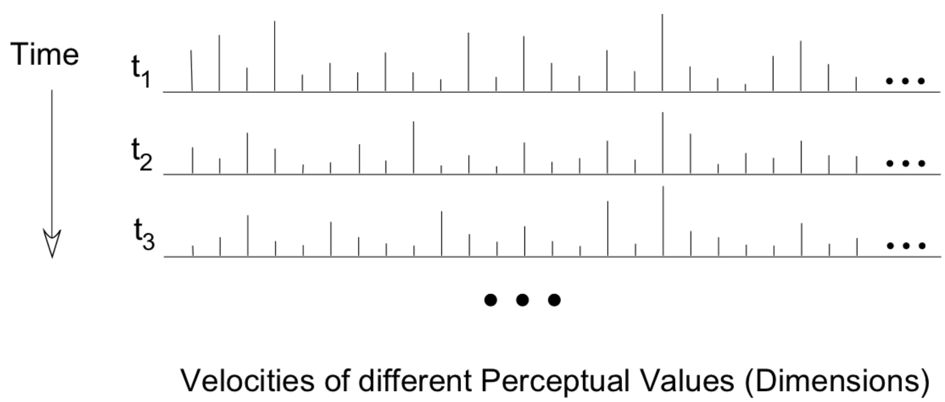

Figure q17.3 Velocities of change of variables on the same set of dimensions (perceptions in this case) at three moments in time. The uncertainty of each distribution (its rattling measure) can be computed.

Not only can the uncertainty of the velocity distribution over the descriptive vector of velocity values be determined at any moment, but in the same way so can the uncertainty of the distribution over time be computed for each element of the vector — each perceptual dimension. The elements of this "time vector" are samples of the value of the element at times just far enough apart that their values are effectively independent. These time-based uncertainties measure the rattling experienced by the variable on that perceptual dimension, which we have taken to be the rate of change of its error value. This measure is the contribution of that variable to the total rattling.

Finally, it is possible to compute the uncertainty of the entire matrix of error velocity variation over perceptual dimensions over time[14]. This overall uncertainty is a measure of total interior rattling experienced by the part of the perceptual control hierarchy under consideration. This interior rattling total is the value that rattling theory says would tend to diminish rather that grow when reorganization changes the connectivity or parameter values of the hierarchy structure.

The same measurement procedure can be used when the basic variable per dimension is the output of the control loop for that perception rather than the error value. The measure based on output values would be of the exterior rattling. Since in most control loops the output value changes slower and more smoothly than the error value, the exterior rattling is likely to be less than the interior rattling for most perceptual control hierarchies.

Distinctions between interior and exterior rattling values do not occur among the "smarticles" considered by Chvykov et al. (2021), but they are important among control loops and hierarchies. Smarticles do not "experience" the interactions that rattle them, in the sense of having an internal value that is related to but not identical to the physical force applied in an interaction. Control loops do. They "experience" changes in their perceptual values and error values, and they act on the environment through the output from their Output Functions.

Experienced and active rattling are intimately linked through the QoC (Quality of Control). For a given value of QoC, the error variance is proportionate to the disturbance variance, and the output is directly correlated with the disturbance. The disturbance to the CEV of a control loop in a complex of other control loops consists to two independent components whose variances add to form the total disturbance variance, an external component analogous to the winter-leaf wind, and the summed side-effects and direct effects of outputs from other control loops. For the next little while, we will ignore any "wind" contribution to the disturbance, and consider only the effects that occur among the control loops in the organization.

There are four ways in which the actions of one control loop can affect the ability of another to control. We discuss them in detail in Section III.1.2. For now, suffice it to say that side-effects of the action of one loop on another are more likely to add to the disturbance of the other than they are to shield the other from some sources of disturbance or to enhance the precision of the other's perception of or ability to act on its CEV.

There is, however, always a chance that the side-effects of action by one control loop will be beneficial to some other control loop. Although in an unstructured organization, most of the interactions of the outputs of the various member control loops on the others will augment their rattling, at least some will damp it. And according to Chvykov et al. (2021) the tendency of the organization as a whole will be to move from the higher rattling structure to the lower more readily than the reverse.

Prigogine noted the reduction in the entropy of successively smaller whirls following Richardson's dissipation rhyme of big whirls feeding lesser whirls. In effect, Prigogine noted that the littlest whirls were localized depositories of high entropy and pockets of high energy. If their entropy were carried away into the larger environment, the remaining organization would have relatively low energy per degree of freedom, and reduced remaining entropy.

In this way Prigogine arrived at a self-organized modular structure created by a through energy flow

14. To be accurate, we must recognize that in a control hierarchy, perceptions do not vary independently. Their relationships mean that within the time-by-dimension velocity of change matrix a lot mutual information (Section I.10.2) must be subtracted from the sum of the individual vector rattling measures to find the total rattling over the entire matrix. We need not worry about this nicety for the purposes of the present discussion, but one should keep it in mind.

from (in Richardson's case) the driving stream flow to the uncountable "littlest whirls" that consisted of strongly interacting molecules for which their motions are seen macroscopically as heat. In the case of the perceptual control hierarchy, we called such progressive self-organization "reorganization", in which each control loop uses its own through energy flow to dissipate the entropy that would otherwise be added from the impinging disturbance. Now we are seeing it again in rattling. A low-rattling organization is calm. Calm means low disturbance from other parts of the organization, and low energy variations on average — low entropy or low uncertainty.

Shannon (1949) showed that uncertainty is additive, in the sense that if you subdivide a vector of values such as a sequence over time into arbitrary divisions, the total uncertainty is the sum of the individual uncertainties. Rattling, therefore, is an additive measure in the same sense. Parts of an organization have total internal rattling values that sum to the organizational total rattling. Both of these bald statements again ignore the possibility that the arbitrary subdivisions provide some mutual information (Section 10.2). In the case of rattling, this mutual uncertainty has a physical correlate, the strength of the influence of one arbitrary division of the organization on another.

Whereas Prigogine's process of entropy reduction was dynamic, the entropy measured was that of a moment of interest in a volume of interest. Over time, that "static" entropy changes, as some is carried away from the region of interest to its environment by the through energy flow. Chvykov et al. (2021) describe a different entropy or uncertainty, that of a distribution of velocities, which we might call "dynamic", since it is a measure of the irregularity over time of the dimension by dimension velocities of movement of the point that represents the organization of the signal values through the description space. This dynamic entropy, the rattling measure will drift hither and yon as the units of the organization are changed by reorganization.

High rattling means much irregularity across the dimensions, as different units are moved this way and that way, changing over time, some faster, some slower. It is this dynamic entropy of velocity measure that is likely to be additive except for the mutual information of pairwise interactions. The greater the mutual information — the lower the velocity uncertainty of one component if you know the velocity of another — the more coherent the organization and the calmer it is overall. It is the structure that matters.

In Volume 2, the self-organized structures of interest will be organizations of social beings, whether they be people or sparrows, herring or elephants, ants or wolves, dolphins or honeybees. Here we are interested in self-organization — reorganization — of perceptual control within the individual. Even at this early point, however, we should note that a social organization's description space includes the organizational descriptive spaces of each individual that belongs to the social group, along with parameters associated with all the ways one individual might influence another. Nevertheless, the social organization is still described by a single point with an associated rattling value in that unimaginably high-dimensional space, and it still tends preferentially to move downslope rather than upslope.

II.5.5 Rattling and Individual Reorganization

How does rattling apply to a group of Elementary Control Units (ECUs, consisting of perceptual function, reference input function, comparator with tolerance, and output function) that have not necessarily yet formed into a single perceptual control hierarchy? Surely such a complicated structure is unlikely to be a "minimum rattling" structure? I think that the findings of Chvykov et al. (2021) about the spatial organization of their Smarticles apply equally to at least the external rattling interconnections of ECUs. The result will be that even if control loops start out having random interconnections they eventually will be grouped into a a structure that is largely or entirely a perceptual control hierarchy.

Back in Section 5.3 we argued that a single second-level perception that replaced correlated changes among several simpler perceptual variations, was easier to control than controlling the individual first-level perceptions independently. The total "rattling" of the first-level perceptual errors would be subsumed into the rattling of the more complex second-level perception. In Chapter 8 and again in Chapter xII.1 we discussed, firstly how a tensegrity structure distributed stresses over all its members, and then how a two-level hierarchy with enough controlled perceptions at both levels would probably be a tensegrity structure.

Chvykov et al. (2021) start with a fixed number of smarticles that trend toward low-rattling regions of their available fenced-in organizational space. Their organization is a set of pairwise spatial distances among the smarticles. We address a different problem: whether the construction of new control loops at higher levels, as in Figure 11.1, should calm the rattling experienced by existing ones enough to reduce the total rattling experienced by the expanded organization of control loops. The answer seems to be that it can, and will if the higher levels built onto the hierarchy correspond to perceptions of properties of Real Reality.

Not just for smarticles, but more generally for organizations, we should consider not only total rattling over the entire structure, but also some measure such as "rattling density" and "rattling density variation" across the structure. Let us consider why.

High rattling density in small regions of the perceptual control hierarchy lead to an (Analyst's) expectation that this part of the hierarchy will reorganize more rapidly than will a calmer region. Without knowing of rattling theory, Powers imposed this same requirement on reorganization because it simply made sense to him. Extrapolating the results of Chvykov et al. (2021) in the way we did puts his intuition on a firmer footing.

It all comes down to the distribution of energy, or rather of rattling power, across the set of entities being self-organized. Referring back to the wind-blown leaf metaphor, each leaf is equally subject to the dispersing energy supplied by the wind, no matter how many leaves there may be, apart from the energy-distributing effects of wind-shadowing in the environment. Adding a leaf does not (appreciably) affect the wind power that moves the existing leaves. Nor, apart from shadowing effects, does the rattling experienced by the new leaf influence the average rattling experienced by the totality of the old leaves. The rattling measure is additive for non-interacting entities, as we discussed above.

The immediate consequence of this additivity per leaf to the total rattling is the physically self-evident fact that a calm place is a calm place however many leaves may be in that place. The addition of a leaf should not change the rattling measure of that place as a whole. It is not the total rattling that counts, but the rattling per leaf, or the local rattling density. The variation of this rattling density across the space available for organization determines the strength of any trend for the leaves to stay longer in low-rattling environments than in regions where they are more strongly rattled. The net effect is as if there were some force moving the leaf organization as a whole more in the direction of low rattling than toward higher rattling.

Always omitted from the leaf metaphor is any active rattling that might be imposed by one self-energized entity on another, something leaves do not do. This interaction effect need not be positive, increasing the local rattling density, but if the connections are random, statistically it is more likely to be positive than negative. If it is positive, the organization as a whole will trend toward a place in which the newly added entity adds less to the rattling density, and perhaps even to a place in which the addition reduces the local rattling density in the neighbourhood of the newly added leaf.

How do these considerations apply to structures of control loops? In Chapter xII.15 we present an evolutionary argument for the development of homeostatic loops of beneficial side-effects of control actions. Here we take a different viewpoint based on rattling, which we more or less ignore in Chapter xII.15, and consider energy flows. Each control loop has its own through energy flow, which serves as a local energy source. The loop uses that energy to entrain and export the entropy that is separately introduced by the disturbance. We look a little more closely than before at what this means.

If control were perfect, the energy used by the output function exactly and at all times would oppose and balance the energy supplied by the disturbance, so that none would remain to affect the power of the signal flow within the loop. That signal flow power drives the rattling of the error signal, which in the ideal of perfect control would be zero. Such an ideal can never be reached in practice, and the balance between disturbance power and output power through the environmental feedback function to the CEV is never exact except for isolated instants when the output-disturbance difference crosses between positive and negative in either direction.

Variation in the imperfect balance between output and disturbance introduces variation into the error signal and a variable energy flow rate (power) around the loop. This energy (not power if there is an integrating output function) is magnified by the output function to rebalance the opposition to the disturbance. Rebalancing is never perfect, and would be imperfect even if the disturbance value were to remain at the value it had a whole loop transport lag earlier. Imperfect control means that the perceptual value changes, if only slightly, as the disturbance value changes.

The perceptual and the disturbance values will probably not be highly correlated because of the stronger effects of the history of changes in the disturbance. Nevertheless, the velocity of perceptual changes will over time have a recent probability distribution. This distribution has some entropy (uncertainty) that varies with the entropy of the recent disturbance variation. In a linear system, the entropy of the distribution is equivalent to its variance, and uncertainty in the velocity over time defines the “rattling” measure for that control loop. How rattled the interior signals chasing around the control loop are depends on the Quality of Control (QoC) of the loop.

The worse the control, the greater the average energy in the variation of the perceptual value. Some of this energy is absorbed and dissipated in the tolerance band, but when the instantaneous variation of the perception differs from the reference by more than the tolerance bound, energy is supplied to the error signal and thence to the output function and to the signal flow around the loop. Energy lost in the tolerance zone is dissipated, unused for control, but perhaps contributing to the “temperature” of the ensemble as a whole.

In the language of Chvykov et al. (2021), PCT tells us that so long as the reference value for the perception remains constant, variation in the error value is “driven” only by fluctuations in the disturbance. Their paper, however, deals with smarticles that supply their own energy without friction. Each smarticle is a “driver”, and each smarticle is also “driven” by whatever influence the other “drivers” have on it when their arms hit each other.

The control loop likewise can be considered to be a “driver”, even while the reference value is kept constant, because its application of energy to the environment when opposing disturbance variation has side effects on other control loops. These side-effects influence the disturbances to the affected other loops in one way or another. Furthermore, if the reference value changes, so does the reference-perception difference, which may then exceed the tolerance bound and cause energy to be distributed to the environment and perhaps cause side-effects elsewhere in the hierarchy.

Not all side-effects contribute toward the disturbance (rattling) of another control loop, either within a single individual or in a different control hierarchy. Figure II.5.3 illustrates four ways in which a side-effect may influence control, of which only Type 1 adds to the disturbance to the affected loop. The other three may either enhance or damp the effect of the disturbance, but they will not add to it.

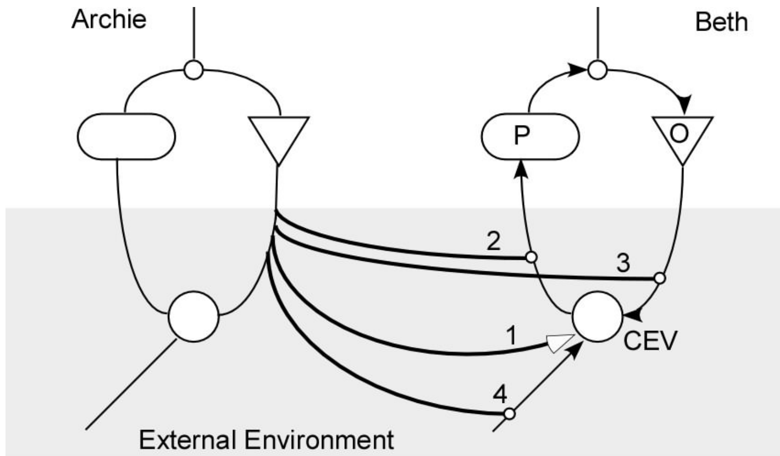

Figure II.5.3. Different ways Archie's actions might influence Beth's ability to control some perception. (1) Directly influence Beth's CEV, (2) Change Beth's ability to perceive her CEV accurately, (3) Change Beth's ability to influence her CEV, (4) Change the ability of external disturbances to influence Beth's CEV. The diagram shows Archie's influences as side effects of his controlling some perception, but the same possibilities exist if Archie controls his perceptions of those influences.

If a side effect of Type 4 damps the effect of the disturbance, it does the same as "shadowing" among the winter-leaves. If a Type 4 disturbance enhances it, that increases the rattling of the disturbance on the affected loop, increasing the effect of variation in the disturbance on the average energy supplied to its error signal. Damping places the affected loop in a relatively calm (low rattling) place, whereas enhancement places it in a high-rattling (leaf exposed to the wind) place from which reorganization will more readily move it.

Interactions of Type 2 increase or decrease the precision with which the perception in the affected loop represents the CEV, while interactions of Type 3 change the precision or strength with which the output can counter the disturbance. Each of the three "non-Type-1" interaction types might either increase or decrease the rattling experienced by the affected loop in the form of error uncertain variation velocity of change.

We finally come to reorganization and its relationship to the minimization of the rattling measure proposed by Chvykov et al. (2021). To summarize how we got here,

- we concentrated on the variation in the error signal, as measured by the uncertainty of the velocity with which it changes.
- For a given disturbance variation, in the absence of tolerance, the energy implied by the error signal is a proportion of the energy introduced to the loop by the disturbance.
- That proportion is a measure of the Quality of Control, still in the absence of tolerance.
- With tolerance, it is the variation in the value of reference minus perception that determined QoC, and
- some of the energy is dissipated by the tolerance, reducing the effective loop gain.
- The rate of variation in the error signal affects the rattling being experienced by this particular

loop and determines the energy applied by the output to oppose the disturbance and reduce the entropy of the CEV.

- Some of this energy is dispersed into side-effects that influence other control loops, either by adding to the disturbance or by enhancing or damping the effect of the disturbance on the other CEV, affecting the magnitude of its rattling.
- Over the whole set of control loops involved, Chvykov et al. argue mathematically and show by simulation and demonstration, configuration change tends to minimize total rattling. Within a single hierarchy we call this reconfiguration reorganization,
- with a quasi-exponential decline over time toward an eventual minimum.

II.5.6 Building a Control Hierarchy through Rattling

In the foregoing summary, we ignored the effect of variation in the reference value. A sudden change in the reference value suddenly changes the instantaneous error value, just as does a change in the disturbance value. Change in the reference value likewise affects the force applied to the CEV to bring its perceptual value nearer to the reference. The force and the distance over which it is applied together determine the energy used in this change, a simple multiplication if the CEV is the physical location of an object with mass, but energy nevertheless in whatever form the CEV might be, if only the neural energy expended in making all the changes involved in the "neural currents" around the loop.

The control structure we have so far described is a randomly interconnected hodgepodge of control loops and partial loops, not a hierarchy. Side-effects connect hither and yon, to many other elements of the "control soup", not unlike the autocatalytic "soup" we discussed in Chapter II.2. What we are looking for is the equivalence of the catalysis that turns the elemental soup into a series of interacting structures. Catalysis enhances (or reduces) the ratio of elements combined into molecules compared to elements separated.

The construction of a molecule from its elements reduces the entropy of the region of the soup that contained the original independently moving elements. Only when the catalyzed reactions support one another to form a homeostatic loop will the density of the molecules relative to their elemental components continue to be high and the soup as a whole to be at a lower entropy.

We are asking a parallel question to the one Kauffman (1995) asked. Under what conditions will stable and productive loops be built and persist against the global tendency for entropy to increase (variety to decrease). Kauffman's answer (Section II.2.2) was that if there is enough variety among the components, then loops of serially catalyzed reactions will occur and be both stable and productive, creating an ever increasing complexity of molecular species and therefore a widening variety of self-supporting loops.

Now, if we look at it from the point of view of rattling, the components that form part of a homeostatic loop are all in a relatively low-rattling organizational state. The mathematics of Chvykov et al. (2021) would lead to the same conclusion, but we can add a third view of the same. The components of a homeostatic loop are, together, in a lower and less variable entropy state than they would be in the absence of the loop. Less variable local entropy implies lower rattling.

All three ways of looking at the formation of homeostatic loops lead to the same thing: homeostatic loops are self-restoring negative feedback loops, and are therefore in organizational "relatively calm" places. There they will tend to stay maintaining the dynamics of the variations of their (mostly intrinsic) variables despite changes in their local contexts. "Relatively calm" does not mean unrattled. It means less rattled than in neighbouring parts of organizational space. Control is never perfect, and neither is the stability of the variable values in a homeostatic loop. Nevertheless, the stronger the negative homeostasis feedback, the less rattled are the components of the loop (with the same caveats about loop delay as affect the stability of a basic control loop).

One way to enhance the stability of a long loop is to shorten it, bypassing a sequential set of catalytic reactions with a novel reaction that also produces an effective catalyst, but by a path that is shorter and thus less subject to disturbance. The shortest homeostatic loop is a control loop (Section II.3.1, Figure II.3.2).

The next consideration is that some of these homeostatic loops may pass through the external environment, which contributes both to the varying "wind" that disturbs them, and to the functional structure of the loop that reduces their entropic variation. This varying "wind" consists of the set of external disturbances to the parts of the homeostatic loop that pass through the environment. The fewer of them and the better controlled they are, the less rattled are all the loop components, both internal and external. "Fewer" implies loop shortening.

The argument of Chvykov et al. (2021) thus implies that part of the organizational structure of the "soup" will tend to be in the form of control loops. They persist because they cannot get simpler and because as whole loops they act directly to reduce the interior rattling to which they are exposed by disturbances acting on the CEV, their controlled environmental variable.

This argument ignores the side-effects of actions by homeostatic loops and control loops on separate structures that also are part of the entire organization. Nor have we suggested why loops should be built that use pre-existing loops in their construction, or even how the controlled variables we call perceptions will be built from patterns of other perceptions. In all of this we accept the proposition that an organization is more likely to move toward than away from a place where it would be less rattled than where it is now.

The individual's reorganization moves the control hierarchy to a structure of greater control competence (lower interior rattling). The side-effects of one control loop on another within the hierarchy act like the arms of the smarticles, slowly and erratically reorganizing the hierarchy and moving the side-effects of control actions preferentially away from disturbing other control loops toward benefitting them, improving their ability to control.

Finding such a low-rattling place is an extraordinarily complex problem in a high-dimensional space, but not in a local space of few dimensions. Let us consider a very simple case, in which there are four distinct inputs (1, 2, 3, 4) that pairwise create perceptual values "A" and "B", and as a complex of all four, these same four inputs create a perception "C". (Figure II.5.4). In reading this Figure from left to right, imagine the successive development of White Boxes to emulate the behaviour of a Black Box (Section 11.4).

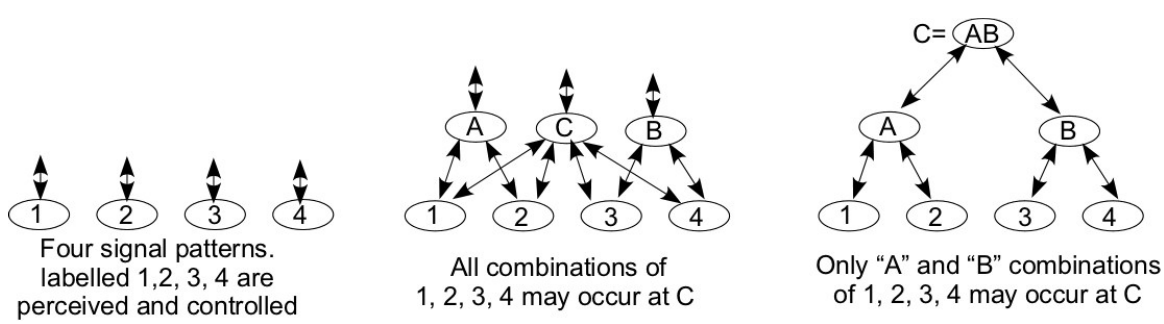

Figure II.5.4 Evolution of a hierarchy from a flat distribution of controllers. Double-headed arrows represent perceptual values up-bound and output values down-bound. Rattling density decreases left-to-right from one diagram to the next if the disturbances as 1, 2, 3, and 4 are statistically unchanged.

We argued above that the tendency toward rattling reduction will lead to the construction of control loops that each control some perceptual variable. Using Figure II.5.4 we now argue that these control loops will not function independently, but will find organizational regions of lower rattling density if they reconfigure to create hierarchies of control.

The initial "soup" of homeostatic loops incorporates some control loops, but the soup also contains separated component parts of control loops, neurons that connect to other neurons rather densely but haphazardly. Most of these synaptic connections will be destroyed by the Hebbian-anti-Hebbian (HaH) process (Section 9.3), but some will survive because of consistencies in the relative timings of the connected neurons. The double-headed arrows in Figure II.5.4 represent pairs of such longer-lasting connections.

In the left panel of Figure II.5.4, each of the four ECUs (Elementary Control Units,1, 2, 3, and 4) is part of a control loop that extends (not shown) into the exterior environment that causes the four variable disturbances that contribute to the rattling that we measure at the four error signals. The question we are addressing is whether the connections of the four ECUs in the left panel into the seven ECUs connected as the three-level hierarchy of the right panel will be more likely to reduce than to enhance the total rattling as compared to a flatter structure such as the two-level seven-ECU middle panel or the original single-level structure, which has fewer ECUs in all.

Compare the middle panel. The environmental disturbances rattling 1, 2, 3, and 4 have not changed, but ECUs A, B, and C are rattled by, respectively, the results of perceptual control by specific groups of lower-level controllers {1, 2}, {3, 4} and {1, 2, 3, 4}. Random synaptic connections would have included several other groupings, but most would have been pruned away because of inconsistency in the resulting controlled perceptions. Whereas 1 and 2 have a regular combinatorial pattern such as a function, as do 3 and 4, the combination of 1 and 3 might be quite random, and pruned away because they proved useless.

The rattling-inducing disturbances are not uniformly distributed in direction and magnitude, and Chvykov et al. (2021) would argue that the interconnections between {1, 2} and A would tend toward relative parameter values that define low-rattling regions of the {1, 2} disturbance space, and similarly for ECU B and the {3, 4} disturbance space.

The situation for C is a little different because C has four connections to 1, 2, 3, and 4. Whereas the rattling space in which A and B drift has only four dimensions (two directions and two magnitudes) within which to seek low-rattling regions, C has 16 dimensions to search, a vastly bigger space (compare the Universes of Possibility on different sizes of checkerboard discussed in Section 10.9). Low rattling regions do exist in that space. By extreme good fortune the parameters of C's connections to 1, 2, 3, and 4 might be close to one such region, but it is much more likely that they will not be.

You might ask whether the rattling measure could be a criterion for an e-coli approach to a lowest-rattling optimum. But if you do, you should also ask about mechanism. What would restrict the e-coli's advance to following a straight line within the lowest-level 16-dimensional search space before it tumbled? Moreover, if the relationships among the four basic perceptions are more complex than simple linear correlations, why would the e-coli procedure not find some minor local optimum rather than a substantially lower-rattling region elsewhere in the space?

There may be answers to these questions, but they are not necessary. Remember that A, B, and C have haphazard synaptic connections before pruning. At least some instances of C will have connected to A and B as well as to 1, 2, 3, and 4. If the disturbances to A and B are related in a way that means instances of A and B tend to occupy a relatively small part of their available space, the argument in favour of C having its connections to 1, 2, 3, and 4 pruned while its connections to A and B are strengthened by the HaH process is exactly the argument that led to A be connected preferentially to 1 and 2, while any connections of an instance of A to 3 or 4 would have been pruned.

Both local changes of the organizational structure, from flatter into more hierarchical, reduce the interior rattling experienced because of the ability of the higher-level components to control their inputs. Control, as we have pointed out many times, is calming of the interior rattling because it reduces rattling experienced at the error signal. However, the route to local reduction of rattling density is an exclusionary one. It works by assuming that the rattling reduced by the creation of a new top-level controller is all that can be reduced.

Even in the "1,2,3,4,A,B,C" complex, 1 and 2 may usefully be controlled in a different way as well as the optimum way found by the re-creation of C at a higher level. Thinking of it this way, the parameter set that connects A to 1 and 2 is akin to the first principal component of a principal component analysis (PCA). More variance, more rattling, might be inherent in the remaining principle components (in this example, there is only one remaining component, but in general there may be many).

In Taylor (1973a) I argued that the HaH process would tend to create a principal component representation of everyday scenes. Now I am arguing the same again, that minimizing the rattling density measure of a collection of controllers will also lead to the construction sets of controllers all at the same higher level that perform the same function of partitioning the variance in non-interacting ways, though if the organization of the Real Reality Black Box does not result in repetitive patterns that are linear functions of their constituents, the linear algebra of PCA will not apply literally. Only the variance partitioning principle will survive, while the non-linearities will be incorporated in the perceptual input functions of the novel ECUs. It is not guaranteed, though it is likely, that high average variance will be coupled with high velocity uncertainty (rattling).

In Section 9.3 on the HaH process, several photographs of different kinds of scenes suggest that the equivalent of PCA variance partitioning will have different results in different kinds of environments (Figure 9.7), and even across different parts of the same sensory scene (Figure 9.8). In Taylor (1973a), I suggested that entire groups of analyzing neurons might be built together if an entity encountered such a variety of scenery, with the different groups being differentially activated depending on the local current scene.

A situation analogous to that described as categorical labelling (Section 9.6, especially Figure 9.16) should be expected to occur. In Figure 9.16, different symbols that represent the same letter described differently ("A, B, C, D, and E" in Caps, Lower Case, sounded out, in Greek, or as school grades) cluster together both as letters and as a descriptive character such as those listed earlier in this sentence. Likewise the two photographs in Figure 9.7 might be labelled "Natural Foliage" and "Human Construction", while individual regions of Figure 9.8 might be similarly distinguished. We will approach this a little differently when we come to "crumpling", later in this Chapter.

In an auditory context, such partitioning might distinguish patterns called "Natural sounds", "Human speech", "Instrumental Music", "Singing", and so forth, down to lower variance partitioning of "Poetry" versus "Prose" speech or "Style of Mozart" versus "Style of Beethoven" or "Rock Music" versus "Classical Music". Such a listing illustrates that variance can be partitioned and sub-partitioned in many ways in both vision and audition, and presumably the other senses as well. We should expect rattling to be similarly open to partitioning, presumably along directions that could be similarly labelled.

How does this affect the trend toward low-rattling? Indeed, do aesthetic preferences even perhaps contradict the "approach to low-rattling" theory of Chvykov at al. (2021)? The immediate intuitive answer to both question is (in my case) "Maybe". A more considered answer is more definitive. Different people may have different aesthetic preferences, but that contradicts nothing we have said.

When one considers the distribution of rattling that exists at an asymptotic steady state, however, there is no contradiction at all. Most people will share aesthetic preferences in any one of the dimensions we may investigate. Few, if any will share them across all dimensions of aesthetic preference, and in any one dimensions only a rare few will have extreme taste in any dimension.

Low rattling is not the same as lack of variety. Chvykov et al. demonstrate that equilibrium rattling of an entire organization implies a quasi-exponential distribution of rattling, out from a minimum toward the extreme in every dimension. The total magnitude of the rattling affects the mean rattling within this long-tailed distribution, but does not truncate the distribution.

The long tailed quasi-exponential distribution applies as a long-term equilibrium, not only within the population of perceptions controlled by an individual, but also across a population. This point follows from treating whole populations of people and their control hierarchies as one self-organizing structure that trends over time toward a low-rattling equilibrium, and that equilibrium exists only when the distribution of rattling is the exponential.

Within an individual's control hierarchy, this implies good control of most perceptions, for most people, but very poor control of some other perceptions. The exponential distribution presumably applies also across people. Most will be quite good at controlling most of their perceptions. The mean of their rattling distribution across the perceptions they control will be low. For a few, however, the mean of their distribution across the perceptions they try to control is high. Accordingly, it seems that aesthetic taste actually conforms, at least qualitatively, to the proposition that organizations trend toward low-rattling configurations.

In general, organisms that control well tend to be calm without seeming to do very much to remain calm, though they actually are actively using energy in "cooling" their controlled perceptions. But, as Chvykov et al. demonstrate, there are likely to be places in the organization of an individual's perceptual control hierarchy where control is not very good. An expert motor mechanic may or may not be a great pianist or politician, a Nobel prize winner in Physics may or may not be a great golfer or acrobat. The perfect all-rounder probably does not exist even in minimal rattling organizations.

As a generic structure the hierarchy is a low-rattling organization, or would be, if all that mattered were the ability to control against the vagaries of disturbances to the sense organs. But we know it isn't. The organism containing the perceptual control hierarchy must survive in the environment in which it lives, which means its biochemical homeostatic dynamics must also be protected against rattling. The interactions among the homeostatic loops, and between them and the growing perceptual control hierarchy must also trend toward low-rattling regions. The right perceptions must tend to be controlled by useful actions. This is a major topic of the Chapter II.7, in which we introduce an abstraction called "The Mechanic".

Whatever the mechanism, reorganization tends to lead more strongly and more often toward good control or low rattling throughout the hierarchy than it trends away from good control. Chvykov et al. (2021) showed that in their simulation a 47-dimensional system showed a smoother exponential decline toward a steady-state value than did the 6-dimensional system of three physical smarticles. I have not attempted to determine the rattling measure for Powers's Arm2 demo (Powers 2008), but one would expect its 196-dimensional structure to show a yet smoother decline in rattling.

Many times through this book we have assumed or asserted that hierarchic perceptual control structures will tend to be tensegrity structures. Let us see how and whether this works with the tendency toward low rattling demonstrated for simple “smarticles” by Chvykov et al. (2021).

II.5.7 Rattling and Modular Tensegrity

Tensegrity (Chapter 8) is a property of a structure or organization. Rattling is a measure of a property of a structure, but that property is not the tensegrity property. It is in a different domain. Tensegrity deals with balances of tensile and compressive forces. Rattling deals with velocities of change of dynamic variances. The don’t sound related at all, but I believe they are, in two ways. One is in the distribution of rattling among the nodes of a tensegrity structure, while the other is in that the interactions among active entities that produce organizational rattling can describe an abstract structure of rattling that has the form of a tensegrity in its tendencies toward less rattled structures.

Tensegrity is a balance between compressive and tensile forces that counter each other at different points (nodes) in a structure, but that operate in different directions in a multidimensional space at the nodes where they interact. When we described a minimal 3-D perceptual control structure that displayed tensegrity properties (Section 8.10, Figure 8.20), that structure was necessarily a two-level hierarchy with three ECUs at each level.

How would this apply to a rattled control hierarchy that displays tensegrity properties? In the minimal 3-D tensegrity control structure of Figure 8.20 six ECUs are interconnected by the control equivalents of “rods and wires”. The environmental structures that correspond to the relevant perceptions are all accessible to rattling influences, variable in force and “direction” (where “direction” here is defined by the relationship among the magnitudes of the influences on the three sensor-level perceptions). The structure of Figure 8.20 being a minimal structure, the “Law of Large Numbers” has little force, since overstress at any one component breaks the entire structure if it ever occurs, even momentarily.

What we omit in the transfer between the physical and the control versions of tensegrity is that the control “wires” that “pull” the CEV and thus the perception toward the momentary reference value do not suddenly snap as physical wires may do. If a control wire cannot “pull” more forcefully, it just doesn’t, but continues to pull with maximal strength, which lessens as the muscle “gets fatigued” — uses up its energy supplies and fails to dispose of its waste material and heat.

The limit is a property of the ECU’s output support, the environmental feedback function to and in the Real World of the environment. Any passive components of the environmental feedback function may break if they are overstressed, but the processing components of the control loop, such as the output function, are very unlikely to break. After the counter-force that overwhelmed it, the output function may require some physiological recovery time, but when it recovers, in most cases it can act just as it had done earlier, perhaps even better.

Apart from that difference, the argument remains the same for control tensegrity and for physical tensegrity. Rattling all the nodes with independent “noises” will with high probability lead to low variability being observed at most of the nodes most of the time. Only occasionally will the forces at any node converge in such a way that they combine to create an influence that momentarily exceeds the ability of an ECU to oppose effectively. Again, the tensegrity structure creates a region of calm in the space of reorganization of the hierarchy — if the disturbances to the various nodes (ECUs) are unrelated statistically.

How might rattling interact with tensegrity? Does a tensegrity structure embody a “calm” region of the space within which reorganization takes place? If so, would the resulting hierarchies tend toward any particular form? A minimal tensegrity structure fails totally if even one of its components breaks. To enhance the survival probability of the structure requires redundant interacting “rods and wires” at each level of control. The “1,2,3,4,A,B,C” structure of Figure II.5.4 is not a tensegrity structure of any kind. Would a tendency to low average rattling density tend to build non-minimal hierarchies that do exhibit tensegrity properties?

If we consider the basic properties of a structure to be its resilience, fragility, and ability to absorb imposed energy, structures that display mechanical tensegrity have high resilience, low fragility, and are able to distribute imposed energy throughout their structure by stretching and compressing their “wires and rods”, and thereby moving the nodes where wires meet the ends of rods. In a perceptual control tensegrity structure, some of these nodes might be perceptual values, but most would be intermediate level reference values.

A tensegrity structure inherently calms a region of any structure that incorporates it, by incorporating the smoothing effect of the “Law of Large Numbers” into its very construction. In a randomly organized physical tensegrity structure, the total effect of N varying values of tension and compressing components (mainly wires) meeting at a node varies on average as $\sqrt{N}$, so the variation of the load at each individual wire varies as $1/\sqrt{N}$ throughout the entire structure. But neither evolved nor designed tensegrity structures are randomly connected, so we should not expect a uniform smoothing to occur at every node or on every rod or wire, no matter what or where external influences are applied to the structure.

The local rattling measure at a node is not the load on a rod or wire, but the uncertainty of the resulting change velocities of the perceptions or the intermediate-level reference values caused by the variation in stretching and compressing of the rods and wires over time. We have argued that the tensegrity structure itself distributes the load on the various components in a way that smooths — reduces the velocity uncertainty node movement due to wire-stretching and rod compression, and hence the average rattling — throughout the structure.

As always with a steady-state rattling distribution, if a sufficiently large tensegrity structure has evolved rather than being designed so as to equalize loads, there will be a distribution of local rattling measures that approximates an exponential with many low-rattling nodes or components and a few where relatively high rattling persists. Accordingly, an evolved structure will not have such a uniform $1/\sqrt{N}$ distribution of effect, but will have a distribution with some lesser uncertainty because of the correlations of effects around the different nodes of the structure. Once again, we come up with the same kind of distribution noted by Chvykov et al. (2021), in which most effects are at the low end of the distribution while some few are at the high end.

A tensegrity “region of calm” in the reorganization space is a region within which reorganization, no matter the mechanism, is relatively slow compared with other regions that are not “sheltered” by the emergence of the tensegrity structure within the hierarchy. But how might the relatively complicated tensegrity structure be built in the first place?

We actually provided much of the answer to this question in Chapter 5, when we discussed the development of perceptual functions for complexes of simple perceptions. We pointed out that six perceptions of the translational and rotational states of a chair are much simpler to control than would be 36 independent perceptions of those states for the four separate legs, one seat and one back of the chair, all of which are constrained by the construction of the chair to move in coordinated ways. The perceived environment of relations among the chair parts is very calm, all the rattling being to the chair as a whole.

Over the 42 potential degrees of freedom (6 for the chair and 36 for the parts) only six at a time can be treated as independent. The effects of action on one component, say the angle of one leg of the chair with respect to a light-source, affect the translational as well as the rotational properties of all the other components, though possibly not of that one.

This spread of effects from control of one property of one component to different properties of other components is not unique to structures that display tensegrity properties, as the chair example demonstrates. But the example also demonstrates that the same reorganization that produces the second-level perception of "chair" also causes distribution of effect across other apparently unrelated lower-level variables. If we consider the chair-leg orientation control, rotating that chair component does not require that any specific leg be moved. Indeed, a different leg might at the same time be controlled to both rotate and remain fixed in location. The action of rotation thus can be combined with the action of controlling a translation perception that is related only because the two entities are related by way of a third, the whole chair.

What happens if the 42 apparent degrees of freedom of the chair are simultaneously and independently rattled? The chair and its components together have only 6 degrees of freedom, but the rattling appears to have 42, all independent. As seen by the chair and its parts those 42 rattling degrees of freedom somehow must cluster into only six, but which six? There is no way to tell, since the six are distributed through the chair. The Law of Large Numbers, however, argues that usually the combined effects of the six independent variations is on average proportionately smaller by $1/\sqrt{6}$ than the possible peak total. Nevertheless, occasionally there will be moments in which all these effects do coincide in phase, creating a total effect that is indeed the sum of the magnitudes of the individual effects. Once again, we arrive at the same kind of distribution, a very few outliers with large magnitudes, while most sum to much smaller totals[15].

We just showed how control against an influence in one property of one part can act as a disturbance to a different property of another part, with considerable freedom as to which "other part" it might be. As an example we chose a different leg, but we could as readily have chosen the seat or the chair-back. The effect of controlling the rotation of that leg was distributed among the other degrees of freedom, but might be made manifest in a variety of places. Reorganization produced the "chair" complex as a perception with six degrees of freedom, but it did not create a tensegrity structure.

The chair, with one "chair" perception represents only one set of six degrees of freedom, but a tensegrity structure requires more flexibility at the upper of at least two perceptual levels. In the minimal structure of Figure 8.20, each upper level ECU contributed to the reference values of two lower level ECUs, and each lower level ECU contributed to the perceptual inputs of two upper level ECUs (not the same pair). That was a minimum requirement to generate a low-rattling "calm place" in the reorganization space. Depending on the inter-related properties of PR and RR, adding further interconnections between hierarchical levels on the perceptual side or the output side would be likely to reduce the fragility of the minimal tensegrity structure and increase its calming effect.

Think of rattling not as an abstract mathematical construct, but as a problem for control. If you know only "one way to skin a cat" and that way is blocked, you are stuck with an unsolved problem — an unskinned cat. But if you have another way to do it, and preferably more than one, environmental circumstances are unlikely to prevent you from controlling that perception. Looking up the hierarchy rather than down, if you have the ability to use a particular kind of knife for cat-skinning, you are likely to

15. During World War II, Churchill is said to have been interested in several very unconventional weapons, one of which used a similar concentration of quasi-random effects. If a stone is thrown into a pool it creates a big splash at the point where the stone falls, and ripples that have a deterministic arrival pattern at every point on the shore. The weapon would have inverted the process, using flapping boards on the shore to create a great wave, apparently erupting out of a relatively calm sea at the expected location of an enemy ship. I have no idea whether an experiments were ever carried out to test the idea.

be able to use it not to "skin a cat", but perhaps to skin a rabbit of a different colour.

The Powers hierarchy is often shown with a complete set of interconnections upward and downward between all perceptions at one level and all perceptions at the neighbouring level upward and downward. The implication is that every kind of tool, whether a knife or not, is equally useful for every task, whether to "skin a cat", to warm the baby's bathwater, or to build a barn. That implication clearly is inappropriate to the world in which we live, whether PR or RR.

We reorganize to connect to any next-higher level ECU tools that have some chance of helping control that ECU's controlled perception. We do this not necessarily by building new connections, but mostly by pruning away connections that do not enhance control, but are more likely than not to increase rattling at the receiving neurons. In Chapter II.6 we will introduce another metaphor, "crumpling", that bears on why and how this happens.

Knives do one kind of job, hammers another, and clocks and hand-warmers yet others. All of them may be used for different purposes, and most of those purposes can be served by more than one tool. The set of up-and-down interconnections us built by the same kind of rattling-reduction reorganization as is the rigid chair. Nevertheless, the resulting modular structures will, at least sometimes, turn out to be tensegrity structures. As such, they will be "calm places" in the space of reorganization, structure descriptions of low local rattling.

We seem to have arrived at the conclusion that the rattling measure described by Chvykov et al. (2021) applies to the reorganization of complex perceptual control systems into localized (within the reorganization space of functions and interconnection parameter values) but mutually interacting *modules* each of which is largely or entirely a tensegrity structure. We may treat each module as a complicated but potentially useful atenfel for control of higher perceptual levels, while the same kind of tensegrity-displaying structural modularity is likely to develop by evolution or individual reorganization at levels all the way to the surface of interaction with the environment (sense-organs and muscle fibres, for example).

We should therefore expect the resilience property of tensegrity systems to be expressed by finding that many, if not most, perceptual functions incorporate inputs from more than one sense organ type. Many blind people, for example, can and do use bat-like echolocation to obtain an internal picture of their location; the "taste" of food incorporates not only sensations from taste-buds in the palate, but also odours interpreted in the nose, the apparent colour of the food (e.g. Spence, 2015; Velasco et al. 2015), its shape (Velasco et. al., 2016) and even the tools used to eat it (Harrar & Spence, 2013).

From the tensegrity viewpoint, it is not at all surprising that food taste incorporates so many disparate modalities, given the importance food has in the maintenance not only of the available energy supply, but also in the health of the dynamic biochemical homeostatic loops. Nodes of the perceptual control hierarchy related to the perception of food taste would be expected to be particularly calm and unrattled[16].

We should again expect the entire organism to contain a distribution of places, including perceptual control "nodes", most of which contain well controlled perceptions. Some perceptions, however, statistically should be expected to be badly controlled and to distribute rattling throughout their local connection neighbourhood, including hierarchies below them in the output side of the hierarchic structure. Under "normal" situational variation, these incompletely learned perceptual controls would rarely be invoked to propagate their energetic rattling, but if the perceived situation changes significantly, the output to the environment might occasionally include a period of substantial rattling despite lower-level intervening tensegrity modules.

16. From my own experience, I can attest that this does not apply universally. A few years ago, sugar suddenly began to taste as bitter instead of sweet. That phenomenon has gradually become weaker, but still persists, and has substantially changed my food preferences.

The foregoing development has many obvious logical or mathematical jumps or gaps, but if it is anywhere near valid, it shows how and why organic structures tend to be both vulnerable to strong attacks, and resilient and able to recover almost completely from lesser disturbances, reorganizing to meet such lesser attacks should they recur in a way analogous to immunization against attack by diseases. Chapter III.9 addresses parallel issues when control hierarchies are distributed across numbers of interacting individuals that form societies, and cultures.

The Analyst's viewpoint that we have implicitly been using is simply not available to the perceiving systems in a controller/observer. All the observer can perceive depends on the information that had been available to its sensory systems up to the moment called "Now". Powers often advised and cajoled students of PCT to think from the point of view of the controller. That is not always easy to do when we really want to understand how the whole system, including any one controller, works. Nevertheless, it is important advice.

II.5.8 Rattling and the Growing Hierarchy

We assume that the perceptual control hierarchy is built on a physical structure of neurons chemically interconnected by synapses that tend to excite or inhibit the electrical firing of the receiving neuron following a firing of the neuron of which they are a part. For a long time, it was recognized that a newborn baby had a very tangled set of synaptic neural connections, which over time became simpler by the deletion of many of those connections. We have implicitly or explicitly used this image several times in earlier portions of the book, and will do so again.

Okabe (2020) examined this simple view of the development of the maturing brain by directly observing the birth, growth, and death of individual synapses in a maturing mouse brain, finding that it was not wrong as a gross approximation, but that it was incomplete. New synapses were being built and old ones reduced or eliminated throughout life. The animal's synaptic network remained dynamic, rather than becoming settled as the mouse aged.

However, it is also over-simple to say that the synaptic network of neural interconnections remained dynamic through the life of the mouse, because the regions of the brain where dynamic change occurred shifted, some places remaining relatively stable after a period of change, others beginning to or continuing to change. In particular, the hippocampus, where memories apparently are built, continued to change at a rate unabated as the mouse matured and aged. Although Okabe did not mention this, the impression I get from his conference presentation is of a brain learning, leaving stable those regions that have learned well, and varying regions that support material still being learned.

In terms of the perceptual control hierarchy, new learning is in part developing new perceptual functions on top of (or as components of) already learned perceptions whose control helps stabilize the dynamics of the intrinsic variables, and in part learning to improve the quality of control of existing and newly created perceptions. The creation of new perceptual functions using as a start an existing tangle of excitatory and inhibitory interconnections would involve considerable reduction of inhibitory synapses with perhaps some construction of new excitatory connections. Both should be subject to the principles of Hebbian-antiHebbian (HaH) learning (Chapter 9), depending on consistencies of the effects on the Perceived Reality (PR) environment of actions created by the output components of perceptual control.

The complex network of synaptic connections is an "Organization" subject to the general rule of Chvykov et al. (2021) that it is more likely for a dynamic organization to move toward a configuration of lower rattling than toward a configuration of higher rattling. If we add to this the observations by Okabe, we find another useful metaphor, the fire front in a forest fire (Figure II.5.5).

Figure II.5.5 A forest fire front. Note the still burning spots behind the front that is advancing leftward in the image (from https://www.firefightingincanada.com/wp-content/uploads/2019/05/c7cb0544a5148ffb234155ff250e5922.jpg)

Consider the organization of the forest, and rattling that before the fire had created a fairly stable ecology of living control systems — trees, fungi, birds, and beasts interacting with each other. This ecology was in a relatively low-rattling state, at least if it had not been influenced by human logging.

In Figure II.5.5, the advancing fire front divides the image into three distinct parts. At the left is a region of the earlier ecology, fairly stable and not much rattled. At the right is a burned area that will be largely dead and unrattled when it cools down. Between these two areas of low rattling is a blazing, chaotic, inferno which at the moment is experiencing extremely high rattling. Treated at a higher resolution, there are several small areas behind the fire front that are still burning, and are subject to high rattling.

For what is this a good metaphor? If we are to believe in the perceptual control hierarchy, which I hope you do, having read this far into the book, it is a very good metaphor for the process of creating and using new perceptual functions whose inputs are perceptions created by already-built perceptual functions. We might call the synaptic region containing the synapses currently involved in this learning process a "learning front", behind which lies the organization containing the synapses used in the newly created perceptual functions whose resulting perceptions initially are not well controlled but are improving with experience (still hot, in the metaphor), and that become ever better controlled as they are used within a sufficiently stable Real Reality (RR) environment (they cool down). Controlled and controllable perceptions are calmer, lower in rattling measure, than whatever patterns of neural impulses might have been produced in the original ecological tangle[17].

Ahead of the "learning front" lies the tangle of excitatory and inhibitory synaptic connections, which form an ecology of interactions that do not contribute to any controllable perceptions. Its organization presumably would change toward lower rattling states independently of any learning relevant to the perceptual control hierarchy. Synapse connection patterns presumably would continue to evolve toward relatively lower rattling configurations, somewhat analogous to the autocatalytic "soup" of Chapter II.2. These changes might have been observable by Okabe (2020), but they would presumably have been slow, in the sense of involving relatively few synapses compared with the many that would be involved in the high-rattling region of changes inside the learning front.

If the metaphor is a good metaphor, it should map the small regions still ablaze behind the front in Figure II.5.5 into some analogous localized high-rattling region within the "already-learned" area. And so it does. Different stretches of the fire-front were set ablaze by different regions that were afire earlier, There are occasional gaps where the fire passes both sides of small areas, leaving those areas untouched. In the synaptic structure, these regions correspond to parts of the original tangle that remain unused in the development of new perceptual functions.

Though unused in the newly created perceptual functions, their connections still send and receive signals from the areas that passed through the learning front. The signals between the "old ecology" and the "under control" domains at the main learning front apply equally to local synaptic ecologies, creating a high-rattling interface or "interfacial learning front" that tends to consume the randomly connected ecology. Eventually, all these localized regions will be consumed.

The same argument applies to the rattling caused by the side-effects of active perceptual control on other perceptual control loops. Any novel perception created in the passage of the learning front is likely to be controlled differently than are pre-existing controlled perceptions, and to rattle the control loops with which they interact, which is another way of saying that they reduce the quality of control of these "older" perceptions. This is a situation we considered when we discussed reorganization, and there is no need to replay that discussion, which we will be extending in Chapter II.7. Instead, we will introduce another useful metaphor, "crumpling", which offers a quite different slant on the learning process.

II.5.9 The Free-Ranging Child and the Cosseted Child

In Chapter 11 (especially Section 11.5), we pointed out that for an individual organism to grow up capable of survival in a particular variety of environments it must be exposed to experiences in that variety of environments. The experiences of a jungle-dwelling child may allow it to become highly skilled in tracking prey animals and avoiding predators that might eat it, but those experiences would be no help in telling the transplanted person that it is unsafe to cross a busy city intersection, especially when the traffic light shows red. The same applies, *mutatis mutandis*, for a child brought up in a city centre. And it applies to informal organizations that self-organize on top of a common frozen supporting environment (such as a set of laws or moral norms). A city child who experiences a more natural quasi-wild Nature in

17. The passage of relatively incoherent synaptic regions through the "fiery" learning front into much calmer regions of actual control is surprisingly reminiscent of the action of Mozart's opera "The Magic Flute", in which a harried Tamino and Pamina pass through trials that include fire into the calm domain of the temple of Zarastro.

summer holidays may develop ways of controlling perceptions experienced in both environments.

Chvykov et al. (2021) make the same point from a mathematical angle. Inadequate driving variety of direction and magnitude does not disturb high-dimensional "leaves" enough or in enough directions for them to find their ways into relatively calm but possibly distant regions of the high-dimensional space. The city child and the jungle-dwelling child are limited in their abilities to develop novel perceptual control structures and atenfels, compared to the child who experienced both kinds of environment when growing up.

In the more restricted environment of Chvykov et al., calmer organizational regions would eventually be reliably approached by the relationships of one of their smarticles to others in an existing organization if it were to find itself in a random part of the space, exposed to drivers in directions it had not previously experienced. In the city, the predators are probably after the prey's money, not flesh, a concept perhaps unknown to the person who as a child learned the spoor of a tiger on a muddy trail, but not the on-line spoor of a scam artist.

A newborn baby abandoned to fend for itself is always in an unfamiliar region of the space. It has no means of countering even mild disturbances from almost any direction, and will die. A well nurtured young child survived long enough to have some experiences, but they usually involve disturbances only in directions within the bounds of normal practice in its family, whether the family be a young mother alone or a caring multi-generational family with siblings, grandparents, aunts, uncles and cousins.

The family, of whatever kind, shields or tries to shield the child from disturbances expected to be too great for the child to oppose or of a kind against which the child has not learned opposing control actions. The family is an organizational module, a resilient tensegrity structure, which may be buffeted as from a variety of directions, with none of its members individually experiencing much disturbance. If the family shield is too strong — what had been called "helicopter parenting" — the child does not experience disturbances from those directions, and cannot learn to oppose them. A tensegrity structure is too rigid to be resilient, since it has developed too few different directions of wires.

An uncosseted proto-adult who has been given free time to explore and figure out how to deal with increasingly difficult situations only gradually enters into a mature life in which it becomes increasingly more responsible for itself. As a mature adult it will live in an environment with more dimensions than that in which it grew its control hierarchy, but it will have already developed perceptual functions for a variety of things that it is likely later to want to control.

Children who over time have gradually been less and less strongly shielded will have been "blown" in this way toward calm places by variable but relatively gentle "winds" from different social directions. They learn to perceive and control against different kinds of disturbance, since the family shielding is sufficient for the child not to be damaged by over-strong disturbances that come from a variety of directions. It will probably have made many different kinds of social encounter with a variety of persons and Roles, a possible reason for the long-ago development of the aphorism: "*It takes a village to raise a child*".

Bad things that happen elsewhere in society or that did happen historically are not hidden from a child who will become a successfully self-reliant adult. Imagining the bad effects did happen to someone else allows the child to imagine consciously "what might I do if that happened to me?" As we described in Chapter II.6 of Volume 1, this kind of conscious imagination may allow the initial creation of atenfels for controlling against that kind of disturbance, should it ever happen

A "cosseting" family has the opposite long-term effect on the eventual adult. Cosseting blocks the wind that might disturb the child-leaf from most directions and reduces its force from other directions in a high-dimensional space. The parents control to prevent the child from getting hurt by acting for itself. The leaf (the growing child) is almost always in a calm region, well shielded by the family. Other than by emulating the actions of the family adults, he or she has no opportunity to develop a variety of ways to counter any of the disturbances it may encounter in later life. But if a leaf eventually moves out from that shielding, as a late teen-aged child moves out from the family, it becomes subject to greater winds and winds from new directions it cannot counter.

A child abruptly removed from such a cosseting family is suddenly in the position of a loose leaf taken from the leaf pile and left out in the open. The child may be unable to control against disturbances of kinds they have not experienced or of magnitudes against which they have had no opportunity to build “mental muscle”. Such a neophyte adult might well control not to counter novel disturbances, but to find shelter from the storm.

Previously cosseted children would probably find freedom difficult to live with. For them, control against disturbances from many unrelated directions at the same time would be hard to learn. For them, control would be easier in environments in which they had fewer directions from which new kinds of disturbance would be likely to appear.

A formal and disciplined school such as a military academy is an organization that provides such an environment. In a school with a formal structure, the new student is well shielded from most directions of disturbance, but not from disturbances imposed by Authority, by the formal rules of the institution. They may reorganize to tolerate such a specific kind of disturbance or to oppose or avoid it by a variety of means they had already learned when the Authority was a parent. The student is likely in later life to evade formal rules and formal Authority wherever their disturbances are encountered, because the corresponding perceptions have become well controlled in the student’s hierarchy by controlling their perceptions of the disturbance source.

When a “student” graduates from an environment of formal discipline, most shields from other parts of the environment are suddenly removed. Just like a child abruptly leaving a cosseting family, the new graduate becomes a “loose leaf”. One natural place for such "loose leaves" to resettle is in another sheltered place where like-thinking people have built their hierarchies so that they shelter each other in ways similar to the ways the parents sheltered the child, though very different in detail. If the child has learned that disturbances caused by Authority are easily avoided or by interacting with Authority so that the disturbance is withdrawn (perhaps by bullying or wheedling the persons playing the roles of Authority), the “leaf pile” group is likely to use similar atenfels.

We see the growth of "leaf piles" sheltered by each other because the control hierarchies have similar parametric locations in many dimensions of their universe. For example, a graduate from a military academy might find the corresponding branch of the regular military a suitable shielding “leaf-pile” (at least until ordered into battle or similar physically dangerous places).

Other previously cosseted children may have reorganized to form a "like-minded" community whose members mutually shelter each other from external disturbances by collective control that opposes those disturbances. We see the growth of scientific orthodoxies, people who accept news supplied by certain media and not by others, street gangs, and militia that invade the Capitols of more than a few nations.

Members of street gangs and of violent militia organizations all have in common that they have no atenfels for survival against disturbances of a wide variety of kinds, other than by acting to eliminate the perceived source of the disturbance. The once-cosseted individual becomes a bully, the group actively opposes competing views of the world. Less actively aggressive individuals may form a group of sycophants to a bully, accepting the truth of the bully’s perceptions of some parts of the world as basically an evil to be countered. In the next Chapter, we will discuss the use of the “Big Lie” that helps the liar to take over control of such groups.

Eliminating the disturbance source works equally well no matter what kind of perception is disturbed, but achieving its elimination usually requires strong action and is likely to be opposed. The "source" may be in an analogue of a well-becalmed leaf pile built by people who were allowed to build their reorganizations on a wider variety of experiences — winds from different directions. The "Big Liar" is likely to identify some such calm leaf pile as the source of the disturbance, no matter what its true source.

The collectively controlled violence of the "source-eliminators" is, to the leaves on the "target" pile, a storm wind from a new direction, possibly strong enough to disperse the target pile unless its leaves are in some way shielded by something outside the pile. Jews were not so shielded from the Nazis once they achieved power by legitimate means. Christian communities (a form of group) in Europe had collectively tolerated their Jewish neighbours except when they were unable to control their perceptions of their economic circumstance. At such times, the source of the disturbance could conveniently be assigned to the local Jews, who could be eliminated in a pogrom, rarely opposed and sometimes assisted by formal or informally acknowledged Authority.

A set of mutually supportive and shielded leaf piles that has survived one such attack may have been moved, like a barchan dune pattern, toward a region calmer in the direction away from the attack. It will be the stronger for the experience, and would more easily survive another similar attack. On the other hand, the "moving leaf pile" might have been sufficiently disrupted as to be destroyed by subsequent weaker "winds" from the same or other directions. For example some Jews might have converted to Christianity (at least in their other-self-image if not in their self-self-image), to reduce the level of disturbances to perceptions they were controlling, thereby weakening the group as a coherent structure.

This possibility brings up a way in which the winter-leaf analogy fails, nested group structures. After an attack many modules may survive while the pile as a whole may lose coherence in the form of the mutual shielding offered by one module to another. Local Catholic groups, for example, survived in protestant Elizabethan England when the national Catholic structure did not, and was not reconstituted for at least another century.

Control of perceptions by other people may have side-effects or direct effects on our control hierarchy, assisting our control ability sometimes, especially if the other people are aware of our existence and control for perceiving us to be happy and healthy — or constrained and miserable. The former lowers the rattling measure for our individual control hierarchy, the latter raises it. According to Chvykov et al. (2021) our hierarchy is more likely to reorganize in such a way as to enhance the former set of interactions and reduce the latter.

A newborn baby in an ordinary family has a relatively small number of people with whom it seriously interacts. It is "rattled" mainly by the environment, and "calmed" mainly by the shelter (in a winter-leaf analogy) provided by its mother, who feeds it and tries to figure out what else might lead it to cry or giggle.

We discussed this stage in the development of a private language between mother Cora and baby Ivan in Chapter II.9, but here we have changed viewpoint, and see Ivan as sheltered from rattling not only by Cora, but also by the rest of the family (except possibly a jealous sibling for whom the novel presence of Ivan is a disturbance to perceptions he had previously been better able to control, such as always having his mother available to shield him or her from rattling of the intrinsic variables because of disturbances to as yet unlearned perceptions).

Ivan is not the only family member who is shielded from rattling by others of the family (unless he is the only child of a single mother without support). If Cora has a loving partner, each shelters the other, forming a "calm duplet" into which the perceptual hierarchy of the maturing Ivan will tend to grow. The family tends toward becoming a localized region of calm, a "module" of interactions with other families with which the interactions tend to be of similarly calming or rattling character.

If Ivan's parents largely agree with the parents of similarly aged Bernard, their module to module mutual interactions are likely to be relatively calm, and so will be the interactions among the family members across the inter-module boundary. Each family shields the other, at least against some directions of disturbance.

We like, and tend to approach in our behaviour, those who seem likely to treat us well, and reorganize to become less like those who do not. These, of course, are statistical tendencies that cannot be applied to any one case. All that rattling theory does is "tilt the scales" toward reorganization that approaches a form that improves our overall ability to control. Rattling reduction does not of itself drive reorganization except in the sense that each control loop is both a driver and driven. It simply makes some directions of reorganization more probable than others.

It is probable that these two families have had mutually shielding relationships with other family modules, and that they occasionally had enhanced rattling interactions with other family modules. The simplest form of interaction that enhances rattling is the direct conflict, in which the two modules collectively control the environmental state of one variable that exists in both perceptual realities, and therefor probably exists in Real Reality. Families, as modules, are likely to avoid interacting with other families on matters for which they control perceptions at different reference levels.

Whereas the cosseted child is shielded from experiencing the variety of disturbances that enables her to control against a range of kinds of environmental disturbances, the over-rattled child has a related problem. The cosseted child has insufficient experience of rattling from a variety of different directions (local environments), the over-rattled child has too, much. That child lives with a caregiver but also in an external environment beyond the family. The external environment would kill a newborn left alone, as it would rattle her in ways for which she has no perceptual functions to provide controllable perceptions, and no output actions that would allow her to control those perceptions if she had them.

A child acquires her mature perceptual control hierarchy slowly, by years of experience that allow her to develop control systems—clusters of ECUs—suited to control against the kinds of rattling disturbances that her family environment allows her to experience. A non-cosseting family allows her to experience mild forms of many kinds of disturbance, with slowly increasing levels of rattling as the years pass. All the time, her reorganizing processes are moving her internal organization, especially including her perceptual control hierarchy, toward areas of relative calm in which she can control against the variety of mild disturbances that do rattle her.

"Slowly" is a key word here. In much the same way as the resources of a grazing Commons can be rebuilt over time, but are depleted by actual grazing, two competing processes occur, reorganization and variety of rattling. We could add crumpling here, which is not applicable to the winter-leaf metaphor. What crumpling does is to keep adding categories to the child's repertoire of object types that exist in the world.

Remember that in our "winter leaf" analogy, it is not simply wind strength that varies, but also wind direction. Every crumpled "crease" in a facet of the high-dimensional "paper sheet" represents a new type of perception for the child to perceive as different, and possibly to control for in a Search, or to locate in an Exploration. The crease adds a new direction over which rattling might influence the child's reorganization. Potentially, it adds to the total rattling experienced by the child, because it adds to a perception not immediately controllable by the child. What we saw, however, in Chapter II.7, is that this is a situation likely to present a puzzle solved consciously by putting together a suite of lower-level control loops as atenfels to initiate control of this new kind of perceptual discrimination.

The cosseted child has little opportunity to develop a usefully variegated view of the world outside the family circle, or of means to control what variety of perceptions he might nevertheless develop. But neither does the over-rattled child, the topic of this Section.

The child must develop the ability to control low-level variables to some degree of stability, so that each individually passes up to the next level only a small part of the variation in the disturbance that affects the perception it controls. As is often pointed out to a PCT novice, when control is good, variations in the disturbance correlate very poorly with variations in the perception, but correlate very well with variations in the output.

What is passed up to the next level when control is good does not correlate highly with the detail of the disturbance variation, but rapid variation in its amplitude envelope does correlate well with the amplitude envelope of the disturbance. When the disturbance changes unexpectedly, so does the perception passed up the levels by any single control loop, at any level. When control is poor, much of the detailed variation of the disturbance continues to appear in the perceptual signal passed up the hierarchy chain, along with the slower-varying amplitude envelope of the disturbance.

Several level N-1 controllers contribute their perceptual signals to any one level N controller. If the level N-1 controllers have not reorganized to control well, their fine variations are passed up to level N along with their relatively slowly varying amplitude envelopes. Since the level N controllers now have to deal with variations that are too fast for the processing loop delays their extra level imposes (and probably additional transport lags as well), the fast variations may prohibit the level N controller from controlling at all, and lead to phase-related positive feedback loops, the first intimation of a growing Bomb in the Hierarchy, or at least a period of what an observer sees as either random movements or a temper tantrum.

The basis of this Bomb is not the same as was hypothesized in Section 6.5, which was that something external removed negative feedback connections in the external portion of a control loop in the hierarchy, exposing a latent positive feedback loop that might propagate up the hierarchy. This Bomb is internal, caused by building a level N control loop on the foundation of poorly controlling loops at level N-1.

Returning to the parable of the starving sheep in the overgrazed Commons, we see reorganization as being the rebuilding of the commons, and rattling corresponding to the variably sized herds of sheep grazed on it by selfish and community-centred herders. So long as the variation of herd size is not too great, somewhat selfish herders can be accommodated within the steady by relatively slow process of regrowth. So long as rattling variability is not too great, slow reorganization is fast enough for the low-level controllers to become stable, removing a high proportion of the disturbance variability of variance before it gets to the next higher level.

If we look now at the family, not the external environment, each family member has survived the rattling to which he or she has been subjected as they matured, and has perhaps not been able to develop high-level controllers if the lower- and mid-level controllers have reorganized too slowly because of having been subject to high rattling in childhood. They do not control perceptions of their environment well, and quite possibly do not control well against the disturbances caused by a new baby or a "no-good child".

Their actions rattle the child and slow the child's reorganization if the rattling is too severe. They quite probably also are not good controllers against disturbances caused to each other by family adults in the household—if the environment is stable enough to include a house or apartment in which the family can live. Furthermore, poor control implies increased disturbances to the biochemical homeostatic loops—the dynamic intrinsic variables—and that almost inevitably results in poor physical health.

The general result is a high likelihood that poor environmental conditions in one generation will have rattled some members of that generation sufficiently that they have not learned how to deal with disturbances created by each other, and their children will be similarly over-rattled and unable to control much very well.

The issue could easily produce the same problems through the generations. A child who could not stabilize control of some highly rattled perception through experience over a stable foundation might instead reorganize by acting as the same-sex parent acted in similar situations.

Chapter II.6 Crumpling and Perceptual Control

Most people have at some time crumpled a piece of paper before throwing it away, usually without giving a moment's thought to what actually happens in the paper during the crumpling process. It may seem an odd digression to talk about the physics of crumpling paper in a book about the widespread ramifications of Bill Powers's discover of the perceptual control approach to psychology, but I hope you will find that it is not. Crumpled paper is not uniformly distorted from its original smooth flat form. Rather, it has regions that remain flat, separated by creases or folds, some of which are sharp and well defined. We will call the flat regions "facets".

We base this Chapter on a study by Andrejevic et al. (2021), who gave crumpling rather more than a moment's thought, using the actual data from the sheets crumpled by Gottesman et al. (2018). We use the work of Andrejevic et al. as a metaphor for the co-evolution — development by reorganization — of the categorical and analogue perceptual control systems in an individual. In Chapter IV.1, we will extend these concepts to the modular organization of social systems.

When you crumple up a sheet of paper before throwing it into the wastebasket, what actually happens to the sheet? Many things happen to it, but Andrejevic et al. concentrated on the fact that it acquires creases, lines along which the paper is sharply folded, and which form boundaries between patches they called "facets" that remain uncreased. They measured such variables as total crease length, facet area, and numbers of neighbouring facets that appear as a constrained piece of paper was crumpled, straightened out, and crumpled again in a standardized manner.

These facets will be important when we relate crumpling to perceptual control, because we will treat a facet as a metaphor for a range of values of a single perceived property of an object that has many properties, some accessible to already created perceptual functions that produce values of their perceptions, some not. But that is for later in the Chapter, which begins with crumpling of thin sheets of what we will call "paper" (actually thin plastic) as an observable process.

We begin by considering the physics of crumpling, at first when a flat sheet is crumpled, and then when different dimensions of a high-dimensional object are individually crumpled. The rest of the Chapter uses crumpling as a metaphor for the development of the categorical perceptual control hierarchy (Section 9.7) and its relationship to the analogue perceptual control hierarchy. The view using the crumpling metaphor highlights some aspects of the perceptual control hierarchy that are not evident from any of the various approaches we have taken thus far in the book.

II.6.1 Crumpling a Flat Sheet

Andrejevic et al. (2021) treated crumpling as an applied physics question, as Chvykov et al. (2021), did for rattling. However, whereas Chvykov et al. introduced a new statistical measure, "Rattling", the uncertainty of a distribution of change velocities, Andrejevic et al. did not introduce a new measure of "crumpledness". They used only conventional measures such as crease length in their analyses.

Andrejevic et al. analyzed sheets that had been crumpled in a study by Gottesman et al (2018). In their study, Gottesman et al. inserted a thin sheet that I will call "paper" (actually a thin Mylar sheet 10 cm square) into a cylindrical tube slightly over 10 cm long (Figure II.6.1). The paper curved around the tube perimeter, and the tube was wide enough that the edges of the sheet did not overlap). The paper was then crumpled by a piston pressing from above to a prescribed depth, after which the paper was removed and the measures of interest taken. The re-flattened paper was then reinserted and crumpled again to the same depth of piston depression. This cycle of crumpling and measurement was performed a total of 24 times. They did this with various values of the depth to which the piston was pressed.

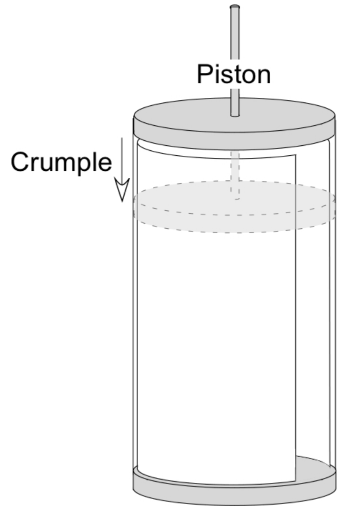

Figure II.6.1 The experimental arrangement,
showing the uncrumpled "paper" in the tube.

The interesting fact for our purposes was that after the first two or three crumpling events, the starting condition and the depth of the depression of the piston had little or no effect on any of their measures, or on the visual appearance of the pattern of as yet uncrumpled facets. For example, a sheet originally folded and creased into squares resulted in measures that were indistinguishable from those that started with a smooth sheet. After a few crumpling cycles, even sheets that on each crumple were subjected to greater or less compression (piston depth) were indistinguishable visually or in their measures.

In other words, if crumpled, the general properties of an oft-crumpled sheet will be the same, though the details will vary, no matter how the sheet started. In this, crumpling is similar to the processes that led Boltzmann (1877) to his statistical concept of entropy, the distribution of energies among idealized gas atoms that bump elastically into one another in an isolated chamber. Boltzmann found that the formula entropy $H = -\Sigma p \log p$ agreed with the measure of entropy based on bulk properties. Seven decades later, Shannon (1949) found that the same formula described the uncertainty of a variable if "p" were the probability that the variable had a particular value among a defined set of possibilities. Since we will be using crumpling facets as the analogue to discrete ranges of values of perceive properties of objects, it is interesting to observe similar processes at work in their construction.

Andrejevic et al. (2021) were interested in properties of crumpled as opposed to flat surfaces in various industrial applications. We, on the other hand, will be interested in their observation that each crease might become a line of discrimination between the parts of any facet it crossed. The distinct regions of a facet would then become separately identifiable entities. We will also be interested in something Andrejevic et al. could not measure, the bending of the sheet before an actual crease is fully formed.

Considered in the two dimensions of an *ideal* paper sheet, creases are lines that would cut across the paper at least over the width of a facet. But no physically realizable, bendable, and creasable sheet is ideal. In a real material with internal strength to resist bending, a crease may not reach entirely across a facet. Not all of them did in the original observations of Gottesman et al. 2018, or of any of the previous students of crumpled paper or mylar. Instead, a crease might terminate in a sharp bend that might continue along the line of the crease and become flatter with distance across the sheet (Figure II.6.2). When we begin to investigate crumpling as a metaphor for perceptual control reorganization we will find that this "bend-termination" of a crease has ramifications for controllable perceptions. But that is for later.

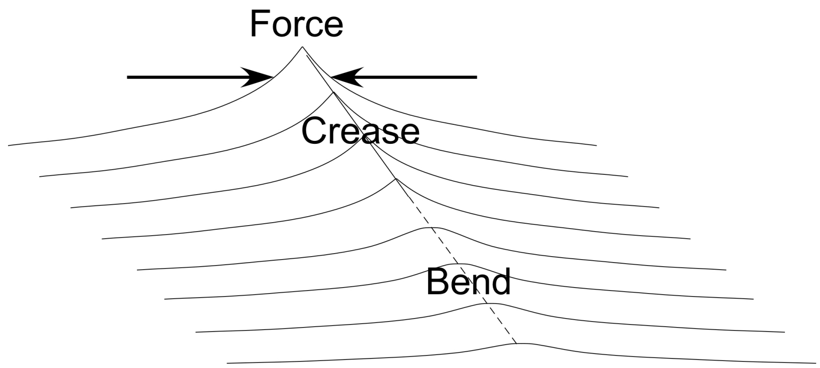

Figure II.6.2 A crease dissipating into a bend. The force is applied only to a localized part of the paper. Nearby, it creates a crease. Further away less force is transmitted through the paper and the reduced force is sufficient only to create a bend

All the experimental measurements of crumpling had perforce to use real physical material. No matter how thin the sheet might be, it had internal strength that initially resisted the crumpling force and (apart from the first one or two crumpling cycles when the piston and bottom of the tube would press all along an edge of the paper) transmitted it from the pressure point along the line of the crease for a distance that depended on the thickness and material of the "paper".

We will use this property of resisting bending and creasing when we "crumple" the objects whose properties we perceive and use in the perceptual control hierarchy. For now, we should simply note that the smaller the facet, the more probable it would be for any crease that starts in the facet to propagate as a crease completely across the facet without a "bend" extension to the next crossing crease facet boundary.

Across a crease, the boundary between facet fragments is sharp and clear. Over a bend, it is not. For both crease and bend, if they were topographic hills, energy would be needed to climb to the peak, some of which might be recovered when descending the opposite slope. In paper or mylar, this energy is not gravitational potential energy, but the energy stored in internal stresses that resist the bend. To flatten a region of the paper near the bend to the general level of the facet requires force, and the use of energy. More to the point, we will find that it takes energy to change a categorical perception from belonging to one category to the category on the opposite side of the crease-bend hill.

A sheet of paper, once creased, is no longer flat in our three-dimensional world of observation, though internally it remains two-dimensional. If this seems contradictory, think of the two-dimensional surface of a sphere, which we see from outside as a three-dimensional entity. When we live on the surface of the near-spherical Earth, we usually see it as flat. Only by observation over large flat areas such as the sea or a flat prairie can we see directly that it is not, as, for example when we see a ship sail over the horizon or a line of telephone poles beside a highway bending downward away from us until even the pole-tops are hidden below the horizon. Creasing increases the apparent dimensionality of a sheet of paper in our 3-D space, because we view a the paper from outside the sheet, but internally the sheet remains two-dimensional.

The forces mentioned in connection with the crease-bend transition are within the sheet. Suffice it to say here that the sheet has some resistance against bending, that increasing the sharpness of the bend increases the energy stored in that part of the sheet, and that converting the bend to a crease releases some but not necessarily all of that energy in the form of a sudden shock that must dissipate somewhere. The amount of energy released is proportional to the length of the new crease, plus some fractional amount due to any bend extension along the line of the crease. This will remain true in the perceptual control analogy despite the big difference in mechanisms. In the perceptual control version the energy released in a crease event will link crumpling to the development of tensegrity structures of hierarchic control systems, at least of categorical ones, and from them back to the analogue hierarchy.

Each new crease that crosses a complete facet creates two facets where only one previously existed. Andrejevic et al. (2021) found that the mean number of edges of a facet was 4.2, suggesting that most facets have between three and six immediate neighbours, which argues that creases, in physical paper at least, do not preferentially choose either the shortest or longest distance between facet edges. They also found that the total crease length increased logarithmically as a function of the number of crumpling events. This total length is a measure of the total energy used in crumpling, which on average becomes incrementally less with each crumpling event.

As we saw in Chapter II.5, energy is a critical resource in rattling. The distribution of variation in its rate of use over time and over the space of the rattled entities is typically similar to a declining exponential, with most measures being low, while a few high values belong to the long tail of the distribution. Now we see the same kind of effect in crumpling. Why?

Should we somehow relate the rattling space to the physical representation of crumpling paper, and more particularly to the changing values of perceptions during control? The crumpling process has no sense of velocity, whereas velocity is central to the rattling measure. It sounds very different, but perhaps we can bridge the gap. The biggest difference is the lack of a velocity measure in the crumpling. Velocity of what? In the rattling analogue, the velocity is a rate of change in the value of a perception.

We will, somewhat arbitrarily, imagine a crumpling event, the creation or extension of a single crease, as analogous to the tick of an old clock, and measure crumpling time as the number of "ticks" that have occurred since the start of a measurement, rather than as the number of clock seconds that have passed. Using this "tick" measure, we will be able to measure velocity, not as a count of events, but as the rate of change of total crease length per tick, which Gottesman et al. (2018) found to be logarithmically related to the existing total crease length. A rattling measure for crumpling would then be the uncertainty of this velocity either over tick-time or over clock time, or across parallel crumpling procedures.

II.6.2 Crumpling, Pressure and Energy

"Pressure" is a concept that we did not mention when we introduced and explained tensegrity in its physical form or its control form. But we will be using crumpling and the intrinsic pressure of tensegrity later in this Chapter, so let us begin to deal with "pressure" now.

Why are we at all interested in these physical parameters of the crumpled paper? It is because we will associate a facet with a perceived property of an object, and a facet boundary crease with a change in some perceptual property of that object, a transition from one category of object to another. The physical properties of crumpling translate almost directly into analogous properties of control. The two-dimensional "paper" object represents a perception of a single property of an object. Some ranges of values of that perception allow an object to be used as an atenfel for one set of perceptions that might be controlled, while other ranges allow it to serve as atenfels for other perceptions. Objects differ in the set of their potential uses as atenfels. Furthermore, any object may be crumpled across one dimension independently of any crumpling in other properties (dimensions).

The crumpling "pressure" we will examine from different viewpoints will be seen as "rods" in a high-dimensional description space opposed to "wire" tensions implied by the tendency for organizations to move toward low rattling in this description space. But all this is for later, after we do the necessary conceptual analyses.

The paper crumpled by Andrejevic et al. would probably not crumple into an array of creases without the boundary constraint provided by the tube walls. Without the tube, if you stand a sheet of paper on an edge and push down on it, usually all that will happen is that the point where you push will go down to the floor, and the paper will bulge out in a big bend, but not into a fold with a crease. The constraint of the tube side-walls is essential to forming a crease.

Within the constraining walls of the tube, the paper pushes back at the piston (and at the tube walls). Pressure is force per unit area of a surface, so if we know how hard the piston pushes and its surface area inside the tube, we can compute an equivalent *average* pressure within the tube. Unlike, however, the situation of a tube filled with gas, the force against the piston is very non-uniform across the piston's surface, being concentrated only at points and lines where the piston touches the paper. There are few such points initially but they become many and well-distributed across the surface of the piston after a series of crumple events have created many creases between mostly small facets.

Usually, when "pressure" is discussed, at least in classical thermodynamics, that pressure is a measure of a global property of a volume of homogeneous fluid (liquid or gaseous). The pressure can be measured by how hard the fluid presses on the walls of its container, in units of force per unit of area. The gross magnitude measured this way was all that was available to a human-scale observer. The concept of "pressure" was as elementary as that of "Volume" or "temperature", which figured in the thermodynamic equation "$P \times V = nR \times T$" or "$PV=nRT$", where R is a constant that depends on the gas in question and n is a number of molecules in the container.

Boltzmann (1877) explained pressure as the summed effects of the impacts of the random motions molecules of a gas, one at at time, on the container walls. The change of direction of motion of the molecule in the bounce implied the application of force by the wall on the molecule and by the molecule on the wall. Boltzmann noticed that the container walls moved imperceptibly slowly compared to the motion of the average gas molecule, so the force could be averaged over a huge number of these impacts.

In the case of pressure in crumpled paper or a tensegrity structure, two conditions are different. In both, the forces are static rather than momentary in the absence of external events, and in a tensegrity the compression forces are applied not by boundary walls but by wires that connect the ends of the rods. The wires compress the rods of the tensegrity structure, and there is no continuous boundary "wall" to form the equivalent of the forces that oppose the molecular impacts of a compressed gas. How, then, can we measure a pressure within a crumpled paper or a tensegrity structure and reasonably give it the same name as we do in thermodynamics?

There is actually a straightforward conceptual approach to this question. Let us imagine that instead of the crumpled paper, a gas is introduced into the cylindrical tube of Figure II.6.1. As Boltzmann proposed, the summed impacts of the gas molecules apply force to the surfaces of the tube. This force is locally orthogonal to every surface of the tube, and hence is an outward force that would push the piston and the bottom apart if the piston were not held steady by the application of downward force. The pressure P would be force per unit surface area of the top, the bottom, or the cut. Pressure — force per unit cross-sectional area — is a measure of compression, and compression is what a tensegrity rod, or the gas in the tube walls.

The pressure of the gas in the tube is increased by reducing the volume of the tube, such as by pushing the piston in a distance Δh. The pressure is inversely proportional to the volume, so in this case the pressure is increased by a factor $h/(h-\Delta h)$. If we think now of the Boltzmann view of the pressure as being the action of many individual atoms of the gas bouncing off the tube wall, including the piston and the base, each atom provides a force that is proportional to its change of momentum perpendicular to the surface off which it bounced.

When we are dealing with crumpled paper, that force is applied by the piston to a number of crumpling peaks that touch at point- or line-like regions of the surface. Whereas a gas allows for time-averaging of the impacts of a multitude of atoms, the paper stays in contact, avoiding the need for time-averaging, its resistance to bending pushing against the force depressing the piston. The calculation of pressure does not change, at least in principle. It will change on average by the same ratio $h/(h-\Delta h)$, though on different occasions, the relatively small number of points of interaction both within the paper sheet and between the sheet and the tube surfaces means that the change will differ from event to event, according to some distribution we need not consider in detail. It is equivalent to having a very small number of very heavy atoms in a gas-filled tube. The mass of those “very heavy atoms” corresponds to the R in PV=nRT.

In a gas, n is proportional to the number of atoms or molecules constrained by the enclosure (we will stick with atoms). The moving atoms sometimes collide with each other or with the container wall, exchanging energy and momentum at each collision. Each atom moves from one collision to the next with a certain kinetic energy, which an atom may gain or lose in the next collision. Boltzmann (1877) showed that this process eventually leads to the atoms having energies that are distributed exponentially, most having little energy while a very few are moving fast and have a lot of kinetic energy. T (Temperature) represents the average kinetic energy of an atom. In the crumpled paper, T corresponds to the energy stored in the bends of the paper, dissipated into the surrounding layers of paper when the bend suddenly transforms into a crease, like a shockwave dissipating into a gas.

With these concepts in mind, we can imagine how they apply to the crumpling of the paper sheet. Let us first consider the energy transfer from the movement of the piston to the paper sheet. The pressure P is the total force on the paper on the piston surface divided by the piston area. Energy is the force on the piston times the vertical distance the piston moves. How is this energy transmitted to the paper? Through the points where the piston actually touches the paper — mostly the peaks of the projections of various creases out of the plane.

These points must represent the atoms of the gas. Each cycle of creasing therefore adds new, smaller, “atoms” into the tube, splitting the larger “atoms”. As n increases in PV = nRT, the average R correspondingly decreases.

II.6.3 Networks and Their Duals.

The sheet that was crumpled was initially a plane, meaning that it dimensions could represent only two orthogonal perceptual categories, whereas the perceptual control hierarchy is based on senses that report many different dimensions of data. Unless some of these dimensions are highly correlated over evolutionary time or the time of reorganization, the perceptual hierarchy is likely to retain the same numerical dimensionality at whatever level of perceptual complexity. In what follows, we consider only two dimensions at a time, so that our discussion can be kept in an easily visualized plane that can be “crumpled” so that peaks and valleys protrude into a third dimension that might represent some non-perceptual dimension such as energy usage in neural firing in the creation of the crease or its predecessor bend.

Every network in a plane has a dual network, and we will use the dual of the crease network extensively in what follows. We start by describing planar dual networks in general. Figure II.6.4 shows a network (solid lines) and its dual network (dashed lines). Inside each bounded region of either is a node of the other. Each network is the dual of the other. The interpretation of what either one signifies depends on how the other is interpreted as a mixture of values in the two perceptual dimensions represented.

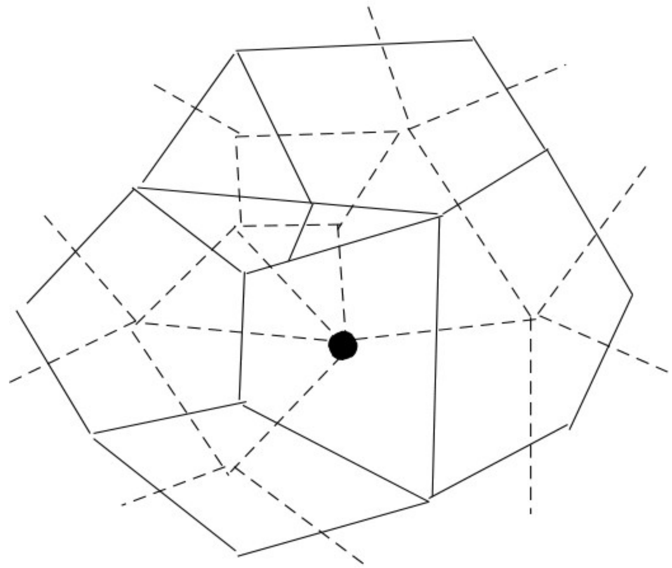

Figure II.6.4 A network on a plane (solid lines), and its dual network (dashed lines). Each network is a dual of the other.

Earlier, we interpreted a network of creases in a crumpled sheet as separating facet-objects that differ across some perceptual property. Those boundary creases might be represented by the solid line network in Figure II.6.4. The dashed links connect nodes of the dual network. Each such link represent transitions across perceptual value boundaries that distinguish between namable properties, such as "red" and "blue". But notice that a facet must have at least three boundaries, and therefore differs from its neighbours in at least three ways defined by different perceptions that have distinct characteristic values across the boundary. If the facet marked by the black dot represents "red" its four neighbours might be "blue", "yellow", "green", and "purple".

Andrejevic et al. (2021) found that their average facet had 4.2 boundaries, many having five. Each of these boundaries would represent a difference between two objects in a different characteristic value of some perception. This value of 4.2 is applicable to crumpling in a single plane, but one plane is inadequate to describe perceptual variations of an object with many perceptual functions. For simplicity we will, however, restrict our initial discussion mostly to the plane.

Each crease, no matter how high the dimension of the space, creases only one dimension, creating a split between two ranges of values of one perceived property. Consider, for example, the multitude of different properties into which the five categories of Figure I.9.16 (Greek as opposed to Roman script, or upper as opposed to lower case in either script, for example). Each is actually a two way split across values of exactly one property. We will soon see a facet as representing one surface of what we called a "White Box", an Object of Object Oriented Programming, or an object that we consciously perceive, but for now we concentrate on creasing one property at a time.

The dual network exists across this one-dimensional set of crease splits whether in "paper" or in high-dimensional "objects". The dual network and the crease network have a one-to-one relationship, the key difference being that creases separate whereas links of the dual associate. As the property collection grows, induced by ever more creases, it must grow in total length at approximately the same logarithmic rate as the crease network.

This close equality inevitably suggests the question of what the dual network represents, and the obvious answer is that it represents uncertainty. Using Mackay's (1953) distinction between metron (what is it?) uncertainty and logon (how much of it?) uncertainty. Knowing on which side of a crease a perceptual value might be (which sub-category it belongs to) is metron, while knowing its actual value to some degree of precision is logon uncertainty.

Each node represents information that might be gleaned by an observer. What would this uncertainty and information be about? The answer would be that it would be about restricting what category of object would have that perceptual value of this property in the range defined by the facet. The more crumpled the paper, the smaller and more numerous the facets, so the greater the initial uncertainty and the information gained by identifying a particular node of the dual network, which uniquely identifies the facet in which it lies.

A crease creates two facets in the crumpled paper where there was only one, reducing metron uncertainty by as much as one bit. A bend does not cleanly distinguish the two new facets, so to identify membership in the facet-category cannot provide as much as an entire bit of information, though it may signify a difference of membership value in fuzzy classes such "a tallish man" and "a tall man". What a bend does is bias the perception away from what will become the crease toward flatter regions of the facet on either side of the bend.

The crease effect still exhibits the category-edge bias in many perceptions, which we have seen before in a variety of contexts, such as that a phoneme near a category boundary will be perceived as a slightly modified version of the central category perception, the centroid of a facet in that direction, such as a /b/ or a /d/. In Chapter I.9 we called this effect "lateral inhibition" but did not associate it with hysteresis. With hysteresis, the strength of lateral inhibition is increased over that discussed in Chapter I.9.

These thoughts may lead you to think of a facet as representing two perceptions of properties of an object, distinguishing one category of object from another. The value of a perception would then be a point in one or the other of those category-discriminating regions. But this assignment of two arbitrary spatial dimensions ignores that the paper in the tube was crumpled only by a one-dimensional vertical force applied by the piston. No matter what direction a crease forms on the paper sheet, it is only one-dimensional, and separates the facet into two only in the direction of a link in the dual network, along the guide lines in Figure II.6.2.

One link between two nodes poses a one-bit question "Which possibility is it?". There are two possibilities, but is the answer exactly one bit? No it isn't. If link length summed over the whole network or any sub-net represents uncertainty among an ever-increasing number of possibilities, the same is true of the length of a single link, and of parts of a link. The uncertainty of a link is partitioned according to the ratio of lengths Included in the two facets or nodes that it joins. One corollary of this is that the information available about the entire perceptual space by a choice of facet is proportional, on average, to the area of the facet. The smaller and more precise the facet, the more information is communicated by its choice from among all the other facets in the crumpled two-dimensional sheet.

II.6.4 Facets and the analogue hierarchy

We now have two distinct approaches to category distinctions, the span of links of the dual network over crumpling creases, and the identification of perceptual functions as category recognizers. On the surface, these views may seem to be contradictory, but instead the apparent contradiction offers a guide to understanding the development and refinement of the perceptual control hierarchy over the lifetime of an individual of any species of living thing[18].

Using the Powers "neural current" representation of neural firing rates, the neural current output from a perceptual function represents the degree to which the firing pattern of its inputs matches the pattern to which the function is attuned. Each neuron in its "neural bundle" has hundreds or thousands of synaptic inputs from other neurons, some excitatory, some inhibitory. Just as our fingerprints and our DNA are supposed to be unique (other than those of identical twins), so are these precise patterns of where do the excitatory and inhibitory inputs come from in other neurons. The bundle consists of neurons whose input patterns are similar, but not identical, over time. The bundle current is virtual variable created by a neural collective (we discuss collectives and virtual variables more generally in Chapter III.1).

Now we take a different tack, and think about analogue variation within a category facet. We have identified a facet with a range of values of a perceived property of an object or a White Box, which implies we should use the two hierarchies we talked about in Section I.9.7. Figure II.6.5 (Figure I.9.17 reproduced) shows their basic interconnections on the perceptual side of both hierarchies. The polyflops represent the different perceptions controlled within an object at a given level of the analogue hierarchy.

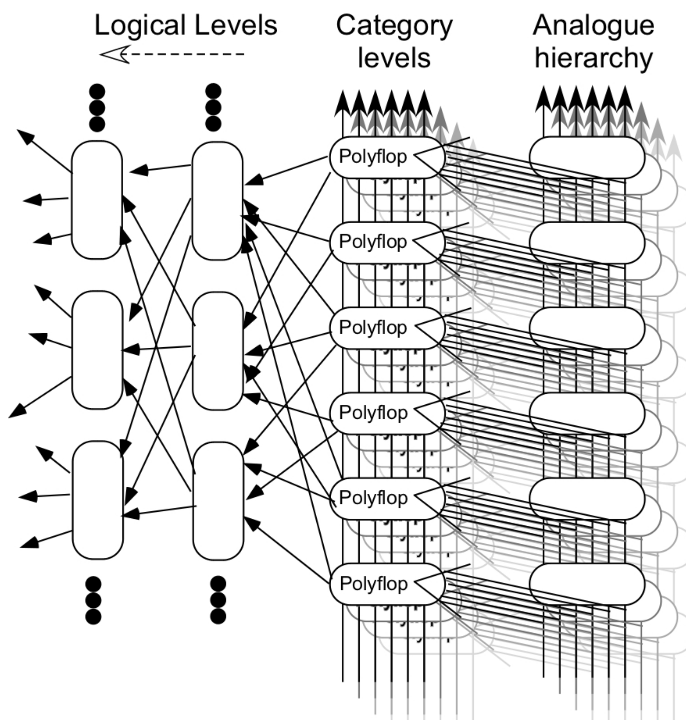

Figure II.6.5 (Figure 9.17 reproduced) The same interface structure as Figure 9.10 emphasizing the multiplicity of analogue perceptions at each level that contribute to a polyflop at that level. On both sides, you should imagine that each analogue and categorical perception is distributed upward to many perceptual functions at the next level, analogue to analogue, categorical to categorical.

In Section I.9.6 and Section I.9.7 we developed the concept of a "category interface" between the two control hierarchies, one based on discrete categories and their interactions, the other on the analogue

18. An interactive version of the detailed current understanding of the complete Tree of Life is at <http://www.onezoom.org/> (Retrieved 2022.03.07).

hierarchy of Powers. Now we can see that the analogue hierarchy deals in category definitions in the form of the perceptual input functions and the similarity of the data to the expectation for the category, but the category interface deals in the differences between the categories, and answers such questions as "Which category IS this?", using the polyflop construct based on lateral inhibition to enhance contrasts among categories that share property categories.

If we think now of a facet as a label for a kind of category, such as "Dog", its clear identification in Figure II.6.5 is clearly on the left side of the diagram, whereas the "dogness" of the entity perceived is on the right side[19]. As we saw when developing Figure I.9.17, the category selection is created by both history and the current value of analogue values of the two different perceptual categories. In the two-dimensional dual network, each link represents the variation of analogue values of two perceptual categories both within a facet and across the crease between the facets.

In this example, across the crease, the uncreased category might be labelled "Canine", the crease might separate the category "Wolf" from "Dog", but the separation between "Dog" and "Cat" would be at a less differentiated level, at which the ancestor categories had suffered fewer crease-making crumpling events, but whatever first distinguishes Canine from Feline within a possibly ancestral "Animal" category would be one that had happened. In a particular maturing child not yet able to walk —call her Mary — her experience might not yet have created a perception of the "Animal" category, but she distinguishes between the family cat and the family dog by some property such as size or fluffiness, which she can sense, and which might form a crease across a category such as "moving things not Mummy or Daddy". Only later might Mary encounter other entities that as an adult she will perceive to be examples of "Animal".

This example may illustrate that major categories can be built from lower-level categories just as readily as lower-level ones can be built by fragmentation from undifferentiated higher-level categories. Indeed, if we remember that the perceptual functions in the perceptual control hierarchy are the category recognizers, and that in Volume I we more or less assumed that higher level perceptions are built from the outputs (perceptual values) produced from stabilized lower level perceptual functions, building by fragmentation is a process we have sneakily added to the repertoire of possibilities for reorganization.

In the transition across a crease, the two possible outputs of a flip-flop exchange on-off values as the analogue value of a perception changes continuously along a dashed line between two facets of a network like that of Figure II.6.4. The centroid of a facet then represents the "platonic ideal" set of perceptual values for its facet-object, considering only those perceptions represented by transitions across the crumpling boundary creases that distinguish "this" kind of object from "that" kind of object. Following any of the dashed network lines from the centroid of the marked centroid in Figure II.6.4 to the centroid of its neighbour would trace the analogue value of the perception that distinguishes the two categories.

Figure I.9.14, reproduced here as z18.6, illustrates the change of category between A and H as the analogue angle of the side lines change from parallel in the H to meeting at the top of the A. If the "Contextual or Task Stress" is low, then the transition is smooth, from "A" to "more A-like than H-like" to "More H-like than A-like", and finally to "H". The analogy is a crumpling bend not enough to stress the paper to form a crease. A bend does not separate the perception of the configuration into discrete categories though it may bias the perception one way or the other in the neighbourhood of the boundary.

19. The choice of sides in the orientation of this and related diagrams is by no means coincidentally that of the analytical, difference-oriented LEFT track and the similarity reporting RIGHT track proposed as a model for reading in Taylor and Taylor (1983).

If the stress is high, the shift is abrupt like a crease, but with hysteresis, which is not a factor in crumpling physical sheets of paper.

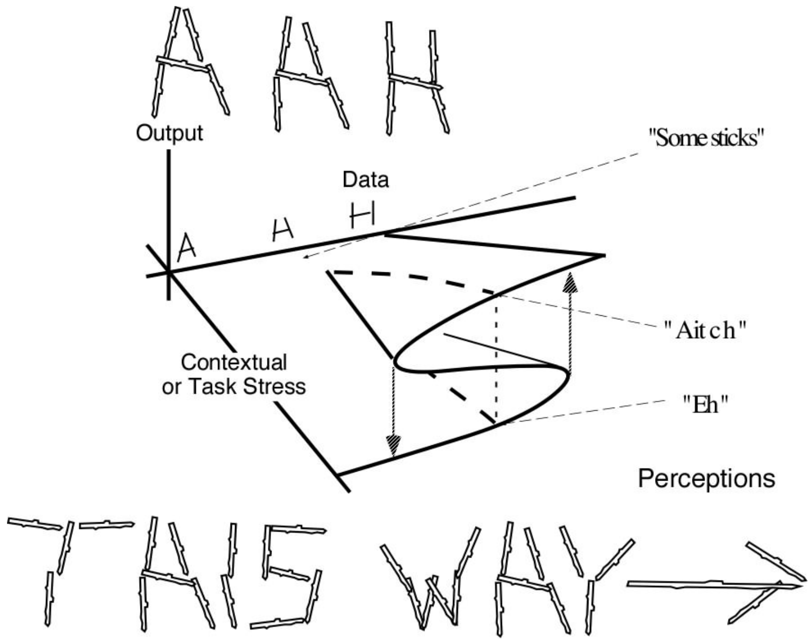

Figure II.6.3 (Figure I.9.14 reproduced) The transition between analogue and categorical representations of a configuration of sticks found on the forest floor. If the first example of the middle configuration in "THIS WAY" is perceived as H in THIS and the identical pattern is then seen as A in WAY, the difference is only in the perceptions of the surrounding letter context, together with a higher-level category distinction between the two perceived words.

Figure II.6.6 illustrates a continuous transition between analogue (low contextual stress) and digital (high stress) perceptions of the same configuration of sticks. The stress is created by the fact that the different patterns of sticks are easily interpreted as coherent larger patterns — word-objects.

None of the boundaries of the "H-object" facet, however, are between a canonical "H" and a word of any kind. "H" can be seen as a White Box that is a functional component of a "THIS" category word White Box, whereas "A" is not. Nor are other possible letters that might be formed by accidental slight displacements of any of the sticks. Those displacements might well create creases between this "H" and other possible interpretations as different letters, depending on different distortions, but none of them would function effectively in the "THIS" White Box.

For example, if the bottom-right "leg" of the H were to be shortened, at some point it might be seem as "P". If that leg were angled outward, it might be seen as a "R". If a new connector stick had originally been partway across the bottom, but had been kicked aside by a hiker, the letter might be seen as a "B", and so on. All of these transitions cross different boundaries of the "H" facet, but their analogue forms generate locations within the "H" facet so long as they are not so extreme in their deviation from an analogue "H" as to favour the perception of the letter on the opposite side of the crease. None of them, however, would function in a "TxIS" White Box, especially if the "TxIS" white box were itself a component of what is perceived as a direction indicator. All of them would probably be perceived as "H" when seen on the forest floor.

We must ask whether all these analogue variations of "H" act independently, and the answer is that often they do not. If the simple centroid of "H" to centroid of "A" path crosses, as it must, a crease, might there be a situation in which non-canonical deviations of "H" toward other options could lead to an uncertainty over whether the current form is an "H" or an "A"? Indeed there can be. Consider the central form of the layout of sticks in Figure II.6.6. Figure II.6.7 shows a slight modification to that arrangement of sticks, as the central ambiguous one from Figure II.6.6 (the left one in Figure II.6.7) is modified to look a little more like an "R".

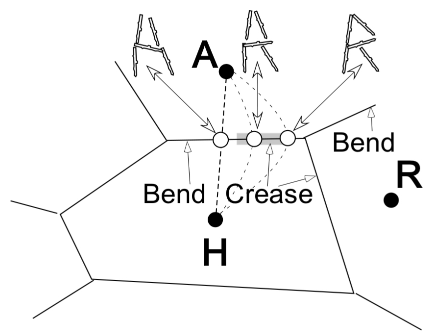

Figure II.6.7 A crease terminating as a bend because of the influence of a neighbouring possibility. The dashed lines indicate different possible routes between "H" and "A" stick patterns through the analogue map that includes the similarities with neighbouring possible categorical assignments of the analogue perception produced from perceptual values at lower levels of the analogue hierarchy. The more "R-ness" produced by the "R-configuration" analogue perceptual function, the less like either an "A" or an "H" is the perception at the peak of the crease or bend.

If you do not have to decide whether the stick arrangement is intended to be an "A" or an "H", you might say it was neither, but an "R" in which the bottom-right stick had been slightly moved. In the language of bends and creases, the relative nearness to "R" affects the clarity of the discrimination between "A" and "H" around the point of maximum ambiguity. If there is forcing from a higher-level context, such as the words "THIS WAY", the bend would progressively approach the "R" boundary as a crease, uninfluenced by the possibility of "R". We see here an interaction between levels of the analogue hierarchy, and perhaps of the digital hierarchy as well.

This example can be replicated with any category types. For example, consider a colour space within which objects of different utility can be described in large facets, such as "greyscale" versus "coloured", These categories are separated by a single crumpling crease. Each facet can be crumpled to create smaller sub-facets, such as "black", "grey" and "white" or "red", "green" and "blue". These can be further subdivided into colours identified and recognized especially by artists and decorators, and the cross-crease interactions in the analogue spaces can be followed, such as between shades of grey and different kinds of "red", such as "deep pink", "pale rose", "scarlet", "carmine", "maroon" and so forth. Similar fragmentation of large facets into ones with subtler analogue perceptual values across creases can be pursued whatever the category type.

Facets, as objects, can be distinguished not only by their analogue perceptual values, but also by their functional performances, which are defined by the linkages of the usages of control of their analogue perceptions as atenfels in the perceptual *control* hierarchy. We hinted at this in the discussion of White Boxes within White Boxes in the "THIS WAY –>" arrangements of sticks, but there is more to it than was brought out there. We shall pursue this possibility shortly, but not yet.

II.6.5 Facets, Facet-objects, and "Searching For"

An object is something we perceive consciously, as a constituent of Perceptual Reality, but objects as such do not appear as perceptions in the perceptual control hierarchy. Properties of objects do, each as a scalar perceptual variable such as the location of the object, its weight, its hardness, and so forth.

Any consciously perceived object has many such perceived properties, many of them used by other objects and other kinds of objects. Each of these perceptions of properties is the analogue output of some perceptual function, a scalar variable. Different kinds of object have different values of their properties. For example a small sparkly transparent object with many little flat faces might be a diamond, but if you perceive that it is not very hard, you will perceive that it is not a diamond.

Class objects are categories; object Instances are individual entities. At any one moment you may see or hear several Instances of a Class. There may be many Chairs around a table, some with arms and some without. The property of having arms is something you can perceive and use if you want to perceive yourself to be seated. It has possible values "Yes" and "No", but not both at once.

The property of "arm-ness" is split by a crumpling crease between these two values, creating two facets where there was only one. The original facet was a perceptual property of the Class "Chair", but that property now distinguishes two subclasses "armchair" and "arm-less chair", which do not have the "arm-ness" property as a variable, because in each subclass it has a fixed value, perceivable but not variable. On the analogue side, chair-like objects could be made with upsweeping sides to the seat, that in the extreme might be called "chair arms" but in chairs with smaller amounts of upsweep would just be aids to seating stability and comfort. The Analogue hierarchy would have a value of arm-ness that when applied to the crumpled facet would translate into one of the two possibilities for the chair "armchair" and "arm-less chair".

Extending this simple example, we can see that a facet generally represents a defining perceptual value range for a Class object. If you are looking at something, and your perception of that property is out of that range, the other side of a facet boundary crease, what you are perceiving is an Instance of a different Class. It is in a different category. A facet represents just one property of what we have variously called an object or a White Box in Perceptual Reality (PR), an Object in Object-Oriented Programming, and a Black Box in Real Reality (RR). Objects have a collection of functions that include perceptual controls for the values of different properties of the object. An "object" is by no means required to be tangible. A political party is just as much an object with many perceptible properties as is a door-handle.

If we identify a facet that we can diagram on a plane (e.g. Figure II.6.7) as representing an object or White Box, the only variations available are boundaries and transitions between internal creases and bends. Boundaries indicate that something is perceptually different between what lies on one side of the boundary and what lies on the other. The facets on the two sides differ in some way that makes them represent different Classes of object.

In the crumpling study of Andrejevic et al. (1921), the boundary between two facets is formed by a crease. We have also discussed soft boundaries formed by bends that they could not observe. We will consider soft boundaries shortly, remembering that we have already talked about them as fuzzy boundaries between perceptual identities such as "tallish" and "tall" when applied to a person. Boundaries separate category identities, whether they be classically sharp as formed by a crease, or fuzzy as formed by a bend.

A facet might have the identity of "chair", with fragments or sub-facets "armchair" and "armless chair” separate at one of the facet boundaries. Another fragmented part of the facet might become a “wooden chair” with “upholstered chair” as a possible counterpart fragment. At a different perceptual level, Martin Luther’s contribution to the pressure already imposed on the facet called “Christianity” was the final “straw that broke the camel’s back” that caused a bend in a facet to become a crease that fragmented “Christianity” into “Catholicism” and “Protestantism”.

Before the final crease in Christianity occurred, various dissident priests before Luther had denounced many of the same practices of the Papal structure, creating the bend that preceded the crease, the final break, across Roman Christianity (a previous crease had separated Roman from Orthodox Christianity). Even after the break, whether Catholic or Protestant, an individual would still be a Christian, and what the person sat on would still be a chair, whether wooden or upholstered. In that sense, across a crease the progenitor facet retains its identity as an element that participates in the definition of a class of object.

One of the “mantras” of PCT is “many means to the same end”, meaning that many if not most controlled perceptions have a variety of actions available for their control. In the conventional perceptual control hierarchy, each action output contributes to the reference values of possibly many different control loops for lower-level perceptions, just the perceptual function receives input from possibly many lower-level perceptual signals. In the crumpling metaphor, this implies that there is a many-to-many between level connection of both perceptual facets and action output facets, just as is shown in the conventional diagrams of hierarchic control systems. We will be concerned about this mantra later, when we discuss individual development, maturation, or reorganization.

The crumpling metaphor automatically agrees in many ways with the conventional hierarchy, but in other ways they differ. The Powers hierarchy contains “perceptions” but does not contain “objects”. The crumpling sequence contains both, if we interpret each facet boundary as a separation between ranges of value for some particular perception.

Whatever the facet represents, its object has several different controlled perceptions. We have identified an object or a White Box by their inclusion of a variety of properties, each of which is potentially perceivable and has property values within a range appropriate to the object type. A multidimensional object is multidimensional in that its different perceptual property boundaries distinguish the sub-facets on opposite sides of the boundary “crease” by their differing potential uses as atenfels. Each perception is a value in a different dimension of the description space we discussed in Section II.5.4.

Another difference between the crumpling metaphor and the development of the perceptual control hierarchy is that the hierarchy has been assumed to develop from the bottom up, building new kinds of perceptions (descriptive dimensions) that depend on consistencies among sensory inputs and action outputs. Following Wiener (1961), we visualized these dependencies and the corresponding input and output functions in the form of White Boxes that emulated the functional effects of the unknowable workings of the Black Boxes of Real Reality (Chapter 11).

On the other hand, the crumpling metaphor creates new perceptual types by refining categories as subclasses of broader categories. A perception of “Dog” may be partitioned into “Husky”, “Greyhound”, Pekinese, and so forth, without any of them being any less a Dog. In our crumpling metaphor, these sub-classes of Dog are connected by being parts of the same “Dog” facet broken by creases across one or more perceptual dimensions. They are smaller White Boxes that individually and together perform the functions of the larger “Dog” White Box in ways appropriate to the breed.

The differences between a Corgi and a Greyhound are in the acceptable ranges of several perceptual properties of objects that are parts of the animal, such as legs, torso, muzzle, and so forth. These affect the different abilities of the various breeds. However, while the kinds of Dog differ in many properties of such objects, all Dogs share more properties, which in the crumpling metaphor are represented by facets broken only by creases between one kind of dog and another, not by creases that distinguish Dog from Cat, both from Fish, and all living control systems from inanimate Rock.

II.6.6 Conflict, Action, Choice

As we mentioned above, PCT has a mantra "Many means to the same end". There is often more than one way to control any given perception, but one cannot use two different ways at the same time without creating conflict. How is the choice made? That has been an open question, at least in my mind, since I first learned of PCT, in spite of discussing it with Powers in person and having puzzled over it in the back of my mind for nearly 30 years. Crumpling, or rather the fragmentation of category perceptions, seems to offer an answer. Consider the example in Section II.6.5 of the sparkly thing that was not a diamond, but for some purposes could be used as though it was.

"Diamond" is a label for a class of object. That diamond-like thing that was not a diamond was an instance of a different class of object. The distinction between Class and Instance is a feature of Objects in Object Oriented Programming (OOP). A Class Object has property values that pertain to all the instances of the class that are not over-ridden by explicit parameter values set for that instance.

Let us suppose that one had the diamond-like thing that was not very hard, and declared that this instance of a non-diamond was nevertheless a diamond for your purposes, which might be to create a sparkly piece of cheap jewellery. Because controlling the perception you want to control uses an atenfel provided by the object does not involve its hardness, that property was over-ridden, not by anything about the object itself, but by the use to which it is put as an atenfel.

It is an instance of some superclass that is shared by the class to which this object belongs and to the class of true diamonds. That superclass was fragmented by the perceived difference between controlled perceptions that can be controlled by using true diamonds and those that cannot be controlled by using objects like this sparkly thing. The crease is created not by an intrinsic property of the object in isolation, but by the set of atenfels it could provide, if they were needed.

The separation could have been made without reference to the uses of the object, simply by noticing some difference in the statistics of the complex of perceptual properties associated with some objects and not with others. Just as the builder of Wiener's White Boxes can glean information from correlations among the output terminals of Real Reality (sensor input data) even when the inputs to RR are held constant, but can learn much more by deliberate manipulation of the input patterns, so it is much more likely that categories such as the diamond and sparkly non-diamond can be fragmented by trying to use them and finding that there are some things a diamond can be used for that a non-diamond instance cannot.

In the dual network of crumpled paper or fragmented categories, there is a crease across the hardness perception or property. That crease is most probably created by the different uses of instances of objects with two different ranges of acceptable properties on the hardness perceptual dimension. When one wants to control a perception that can be controlled only by using a very hard object, the result is a lot of contextual pressure in the sense of Figure II.6.8. This object in my hand is either hard enough to do the job or it is not. There's no fuzzy boundary corresponding to a bend. My perception of hardness is on one side of the crease or the other. The object is hard enough to be useful, or too soft to be useful. Far enough either side of a bend, the same decision is possible with nearly as much clarity.

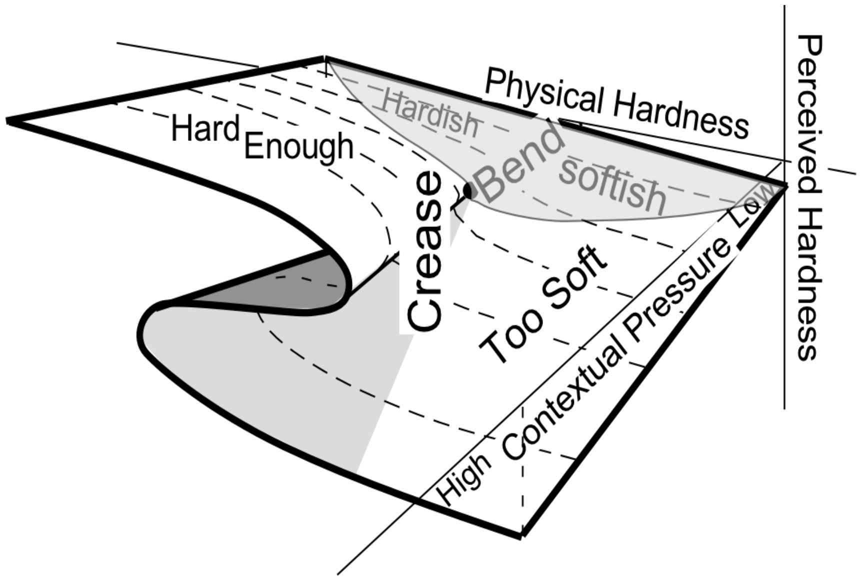

Figure II.6.8 The difference between a bend and a crease separating a harder analogue physical property value (upper and leftward section) from a softer value (lower and rightward section). When the contextual pressure to make a decision about whether the object can be used is low (shaded area), the perceiver is likely to see the pattern as more "hardish" (higher in the diagram) or "softish" (lower) than as either "hard enough" or "too soft". When the contextual pressure is high (white area), the perceiver is more likely to see one or other of two categories, "hard enough" or "too soft", depending on both the data and the history of what has been seen.

One of the axes in Figure II.6.8 is labelled "contextual pressure". How much does it matter whether the perceiver decides whether the object instance is hard enough for a purpose or too soft to be useful for that purpose? That depends on the effective gain of the control system that might use the hardness of the object as an atenfel for controlling its perception. There is a difference between really, really, wanting to perceive one's name etched into a piece of glass and idly thinking it would be nice if one's name were etched into it. If you don't really care, you are unlikely to fragment the category of hardness into hard enough and too soft, but might nevertheless note that the object is pretty hard or rather soft.

Scratching with a diamond is not the only way to etch glass. You can do it with hydrofluoric acid, but if you don't know how to use the acid or don't have any, that approach is not available. Suppose you don't have any hydrofluoric acid, which is pretty nasty stuff to have around the house, but you have some other acid. Could you use it? You have a choice between two things that have a property "can etch glass". Does this property correspond to a category perception that can be creased? They belong to very different objects, "gemstone" and "acid" (but don't try touching hydrofluoric acid!). Acid doesn't even have a "hardness" property to be creased.

So what distinguishes the actions of controlling a perception of the glass being etched by a diamond or by this dangerous acid? I suggest it might be a perception of personal safety, between a class of actions on one side of a bend-crease, that are more or less safe as opposed to dangerous on either side of a bend or safe enough versus too dangerous across a crease.

If this sounds like control of action as opposed to control of perception, it is not; it is control of perception of a property of an action perceived in situational context. This has no directly obvious parallel in crumpling, because crumpling has no active components in the way rattling does and as PCT does. But the perception of a property of an action in situational context has much in common with the perception of a tangible or an abstract object defined by its perceptions of its many properties.

That the object may be defined by the many properties of the situational context of an action is simply a case of an object that has both tangible properties (perceptions of the environmental context) and abstract properties (relationships between those perceptions of the environment, such as that a liquid might be spilled but a solid cannot spill). For the purpose of etching glass, the property range “safe enough” would be required, which implies the existence of a crumpled facet (a fragmented category of object) and the pressure to perceive the difference between what makes *this* situation safer than *that*.

We have identified a crumpling facet with a specific range of a particular perception. The perception is of a property of an object, and its range is a static entity, a noun if we use a linguistic analogy. It does nothing by itself. In the White Box view, a perception of a property is of an active entity, a verb. It takes a certain pattern of input variables and outputs a single value that is the value of that perception. It is a function, and in the perceptual control hierarchy described by Powers based on his own intuition, one of only eleven different kinds or classes of function, defined by their level in the hierarchy. How many different kinds of perceptual function there may actually be in the hierarchy of any particular individual is, at this stage of the science, anyone’s guess.

We do not need to know what kinds of perceptual function inhabit a particular hierarchy to know that different kinds or classes do exist, defined by what kinds of signal their input terminals accept and to what other classes of perceptual function they may send their outputs. Objects (objects or White Boxes) are defined by the list of their properties, or rather, by the perceptions of their properties[20] in terms of what the objects can do or have done to them — what *use* they might be to a living control system at a particular stage of development.

20. Much of what follows was originally suggested to me in an e-mail exchange with Eetu Pikkarainen during February 2017, during which I was being introduced to the ideas of action semiotics.

Chapter II.7. Crumpling and Rattling: Development

Using crumpling alone as a metaphor for the ever finer discrimination of classes or categories of perceived objects is all very well. But every new crease adds to the number of perceptions that might be rattled. On the other hand in itself, the varying disturbances that rattle analogue perceptions generally do not rattle category perceptions unless the analogue effects move a perceptual value across a facet boundary. At first sight the fragmentation of a facet might either increase or decrease the rattling experienced by the entire perceptual control structure.

The intuitive increase or decrease of rattling by further crumpling fragmentation of the facet structure turns out to be irrelevant in practice, since each crumpled crease leaves the ancestral facet as it was. If control does not require the specificity of the new, smaller, sub-facets, the original is still available for use. If it does not matter to someone whether a new dog will be a greyhound, a corgi, or a terrier, so long as it is friendly, then the person will not control for perceiving any of these particular breeds, but will control for perceiving "friendly dog". Because the ancestral category remains available, the fragmentation caused by crumpling need not contribute at all to the rattling of the organization of the perceptual control structure. The perceptual controller is undisturbed by changed in the perceived sub-category.

II.7.1 Crumpling in Evolution and Individual Development

In the experiment analyzed by Andrejevic et al. (2021), to crease the paper requires the application of force to the piston in the tube (Figure II.6.1). The piston moves down a distance Δh and applies a force F in order to do so. Energy can be measured as force times the distance over which the force is applied, so the energy used in shortening the paper by Δh in the direction of the tube axis is $F\Delta h$[21].

Energy is never lost from the Universe. It simply changes form. In the experiment the energy of the piston's movement is used to bend the sheet before the bend becomes too acute and actually breaks into a crease. The energy used in creating the bend is stored in the springiness of the paper and could be released as kinetic energy if the bending force applied by the piston were to be suddenly removed. The bending energy is instead released when the bend turns into a crease, initially in the form of kinetic energy of the paper near the new crease, but when the paper is stable again, none of that energy remains in the form of kinetic energy.

Where does it go when the piston is pressed down and crumples more creases into being? Before the break, the energy was all in the paper, and the greatest energy density was where the bend was sharpest. That is exactly where the crease appears, but after it does, where and in what form is the energy that was in the springiness of the bent paper? Is there an analogy in the formation of a crease that distinguishes perceptual facet-objects?

In paper form, at least some of the bend energy in the paper is likely to be released as heat, but in analogous situations related to PCT, thermal effects are not where energy is dissipated, at least not initially. It appears as the energy of disturbances to other perceptions, causing increased error in those perceptions. The process of energy distribution was discussed in Chapter 8 in connection with the development of tensegrity structures as normal motifs in a hierarchic perceptual control system. The "crease" between the newly separated parts of a facet appears abruptly, creating a shock disturbance to perceptions that were using the unbroken facet.

Energy is used by the controllers of those perceptions to oppose the effects of the sudden shock. The now distinct parts of the facet are represented in the hierarchy by new perceptual functions, and the triad

21. Actually $\int F dh$, since the force is unlikely to be uniform over the whole piston depth change.

of the original facet and its two parts is akin to the inverse process of the creation of a “chair” perception from perceptions of its legs, seat and back. Now, rather than building the chair from its parts, the perceiver had previously learned to perceive “chair”, but has just become *conscious* of the independent existence of its parts, perhaps because one broke off when someone sat on the chair.

The process is the same as for word recognition when someone familiar with the whole sound sequence becomes *aware* that it contains components (syllables) that are used in other familiar sound sequences. I emphasize the words “conscious” and “aware” because these new “syllables” are not controlled yet in the non-conscious reorganized perceptual control hierarchy. They are perceived in a kind of “I’ve seen one of those before” non-specific manner that becomes labelled by the actual syllable identity (e.g. “man”, found in “gentleman”, “handyman”, “seaman” and a host of other words). Later, some maturing individuals might perceive enough similarities between certain syllables such as “man”, “cat”, “dog”, “bug”, “tag” or “bat” to perceive them as members of a superclass we could label as “stop-go-stop”.

Each new creasing event allows the perceiver to identify smaller often observed fragments of facets, such as phonemes within syllables or curved forms carved into wooden chair legs that experts learn are associated with the style of a particular furniture maker. Where a novice might see a category “ornate furniture” or perhaps “Danish modern”, an expert might see “Chippendale forgery”.

Analogue perceptions of variations within facets leads to the creation of these discriminations, especially if the differences imply differences in action. For example, subtle differences in the ornamentation of a chair leg may be the difference between the chair selling for tens of dollars rather than thousands of dollars at auction. As we shall see, the creation of new facets that are fragments of an ancestor leads to and forms part of a feedback loop that includes the construction of new analogue perceptual functions.

When, much later, we come to consider crumpling in social situations, that released energy of facet fragmentation might be manifested as a change as small as varying some parameters of a different crumpling crease, such as a shift in its location within the facet, or as large as the internecine violence of a civil war in which one faction controls for perceptions to have analogue values within ranges that belong to one sub-facet while the other controls for the perceived situation to belong in the other sub-facet. But that discussion is many Chapters away. In this Chapter we are interested in much simpler things. One of those simpler things is the concept of aging or maturing.

We earlier suggested that for crumpling, “time” might be measured in ticks, each tick representing the development of one new crease. Now we ask what corresponds to an increasing number of crumpling episodes? We answer: The maturation of an individual exposed to an environment that induces rattling, as they observe different kinds of repeating differences between this and that in the environment. A city-born child who has only recently learned to walk does not perceive a meaningful difference in the environment when standing beside her mother who is waiting for the sign to say “WALK” and the light to turn green, and then those things happen and her mother walks. But then, neither might those subtle changes in a tiny part of the environment be noticed by an adult who grew up in a jungle and is in the city for the first time.

Why? What changed when the light turned green and the sign changed to “WALK”? Surely a lot changed, but which of those changes mattered to mother enough to persuade her to cross the street? To the child, what changed was mother’s action, not the colour and shape of the WALK sign among the multitude of changes available to the child’s perception. Something in the environment led to mother’s change of action, but the child initially does not perceive it. Later, the child will differentiate this relevant portion of the perceptual environment in similar situations into to two facet-objects that are associated with two different kinds of control action, one in which there is a white “WALK” sign and another in which there is a red “DONT WALK” sign.

Some of the more frequent "crumpling" episodes of this kind will have been experienced by most of the ancestors of this newborn living control system, whether it be a human, a snail, a herring or a tree. Rather than have the individual relearn the facets commonly distinguished as useful in every generation, evolution would be likely to favour the propagation of "creases" that define categories of perceptual values in the newborn that enhance the fitness of the species. We discussed this from an engineering point of view in the early Chapters of this book. Now we are seeing the same thing from a different viewpoint.

"Fitness" is not a matter of what discriminations separate facets by "creases" in the perceptual possibilities. It matters what the individual does to control them in its wider environment — what lower-level controlled perceptions are available for use as atenfels given the circumstances. A deer that is born with as little ability as a human newborn to control its leg muscles to run would not survive long in a world full of predators with a taste for baby deer. The human child has no such problem, provided it lives in an environment in which it is protected by its mother and by the society in which she lives. One would therefore expect the individuals of a species to be born ready-equipped with whatever types of perceptual control have added to the Darwinian fitness of members of the species.

One characteristic of low-level controlled perceptions as atenfels is that they are likely to be more useful — contribute more to "fitness" to propagate the individual's genes — if they are usable as atenfels in a wide range of more complex tasks. From the crumpling point of view, each such low-level mechanism for control represents a large facet, differentiated into sub-facets in more complex ways of controlling higher-level perceptions. The "Crumpling" metaphor thus works between levels of action output side of the control hierarchy, as it does within levels of the perceptual part of the hierarchy. Does it work similarly between levels of the perceptual hierarchy?

The categories are multidimensional entities, types of object, and the plane of a facet is planar only conceptually for the purposes of being able to diagram the relations between facets on a page, as in Figure II.6.7. Nevertheless, before being separated by a crease, neighbouring facets were continuous with respect to the perceptual variable that crosses the crease boundary between the now distinct facets.

A perception whose categorical identities (labels, if noted in language) change across the crease is one-dimensional, though it fragments entire multidimensional objects that share that perceptual property.. But consider a physical one-dimensional line. Although the physical paper is only two dimensions, when the paper is crumpled, lines on it project into the third dimension. What this indicates is that across a crease, nothing keeps two adjoining sub-facets coplanar across a crease in the high-dimensional space of description used for an organization of controlled perceptions.

In the analogue perceptual control hierarchy, the changing values of a perception within a facet are simply changes in the value of a controlled perception that is of a property of as object — a White Box. When it trace crosses a crease into a new facet, it remains the same perception, the output of the same perceptual function, but that perceptual value is now of a property of a different object that is distinguished from the first object by the different uses to which it can be put, different internal connections within the White Box or between White Boxes.

For example, in the A-H transition illustrated in Figure II.6.6, an "A", however shaped, has an entirely different linguistic function from an "H" of the same shape, and in "THIS WAY" the two identically arranged set of sticks laid out on the forest floor could not reasonably have been intended to spell "TAIS WHY". Similarly, among sub-classes of "DOG", there is no point in entering a corgi in a race with a greyhound, nor of matching a greyhound against almost any kind of terrier when a farmer wants to catch the rats who are eating the corn he keeps in a barn. The different breeds were originally developed by guided evolution because of their ancestors' ability to serve different functions.

II.7.2 Crumpling, More than a Metaphor? 1: first "words"

Up to this point, we have treated crumpling as simply a descriptive metaphor for processes related primarily to the category interface of a Perceptual Control System that consists of an analogue hierarchy and a logical hierarchy that are separated by a categorical interface. On one side are analogue perceptions of how much of individual properties of tangible or abstract objects exist in the environment of an Elementary Control Unit, and the strength of belief about that value (Figure 1.4). These single properties can take on a continuum of values depending on how the inputs to their perceptual functions change. On the other are discrete perceptions of entire objects or of their individual properties. These can be, but need not be, labelled individually for communication with other individuals who perceive similar categories.

These perceptual values, whether analogue or discrete, are of perceptions of properties of objects that each have many perceptible properties, many of which offer controllable perceptual values. Many different kinds, or classes of object (such as hammer, kite, shoe, hailstone, or plank) share some properties (such as hardness, weight, or flexibility), though they don't share common value ranges of those properties. Other classes of object (mirage, political party policy, piano sonata, concert audience) have none of those three properties, but share others. All of these latter classes of objects are what we call "abstract entities" but perceptions of their properties are no less real because the entities3 are not tangible.

In the crumpling metaphor, the crumpled paper has a number of "facets" that are separated from each other by creases. We have identified a facet with a perception that has values that for some reason are distinguished by which side of the crease the current analogue value of the perception lies. Facets represent perceived properties of objects, not entire objects that each have many properties. Facets, in the crumpling metaphor, are separated by creases or bends that each distinguish ranges of values of a particular analogue perception that are appropriate to two different classes of object.

Thus far, however, we have made no attempt to suggest just why "these" analogue perceptual values are on one side of the crease while "those" values are on the other side. We have said only what amounts to "that's just the way the cookie crumples". Now we will take two steps further, to argue that functionally, crumpling is close to enough aspects of perceptual control to call it a model rather than a metaphor.

We start with something that we will address in more detail in Chapter II.9, the first language of a baby interacting with its mother. We will call the mother "Cora" and the baby "Ivan". What Cora needs to learn is how Ivan acts differently when he cries because he is hungry and when cries because he has a pin sticking into him. The cry may have a slightly different tone, and the movements of his arms and legs may also be different. Initially Cora does not perceive these differences as meaningful, although in her analogue perceptual control hierarchy she controls for Ivan to be happy, which she perceives is not the case when he cries.

Both audible and visual differences may be subtle, but if Cora is able to identify the differences with a better than chance probability of being correct, she will less often need to search for what is troubling Ivan than if she cannot perceive the difference between the patterns of sound and motion. What she is learning is a distinction between categories that are perceived by Ivan. Ivan acts differently to control values of analogue perceptions that distinguish the categories. Cora need not consciously perceive the difference, but there is learning feedback between her and Ivan, because Ivan is (probably non-consciously) learning that when he does thus and so, he gets fed or is relieved of the pin-prick.

With luck, the action chosen by Cora is the one that reduces the error in the perception controlled by Ivan by having Cora perceive his actions. He is also unconsciously learning to control his perceptions of the tensions in his muscles, and his control improves as he slightly changes his movements and cry styles to enhance the differences that Cora, perhaps non-consciously, picked up on initially.

Cora and Ivan co-reorganize through a positive feedback process that winds up with a distinct two-word "language" emitted by Ivan. Cora, on her part, will use different words spoken in "Motherese" while she feeds Ivan or while she is looking for a pin sticking into him. Her actions and sounds combine into "words" interpreted as patterns by Ivan's synapses, as do Ivan's in Cora's non-conscious hierarchy. The more consistent are Ivan's sounds and movements and Cora's words and actions, the easier it becomes for both to control their perceptions, of Ivan's contentment for Cora, and of his own perception of things being as they should be for Ivan. For both, and for the organization that contains them both, rattling is substantially reduced.

We can use crumpling, still as only as a metaphor, to trace this interaction. Cora and Ivan both perceive as very different the actual situations to which their two evolving "words" refer. Early in the development Ivan's crying rattles Cora's control of her perception of his general happiness, which approaches its reference when Ivan stops crying. When Cora learns to act appropriately for hunger and for "pin", she is rattled less, so she has contextual pressure to learn any signs that distinguish the two states.

The critical distinction is in the distinctly different *actions* Cora must use to control Ivan's crying in the "feed me" pattern of sounds and movements and the "get this pin out" pattern. In the flip-flop representation, the positive feedback strength between the two flip-flop perceptual outputs "hungry" and "pin" is initially weak, but increases as feedback biases both Cora and Ivan in the ways they modulate their actions. Cora more quickly stops the crying because she becomes more accurate in identifying the action appropriate to Ivan's problem. At some point, the perceptual flip-flop feedback becomes sufficient for it to flip cleanly into one of its two states, which for Cora mean "He wants feeding" or "He wants to get a pin out". In the crumpling metaphor, at that moment, a bend has broken into a crease. Two facets "please feed me" and "please remove pin" have fragmented from one ancestral facet "Something is wrong. Fix it.".

I deliberately put labels on these two facets, as well as the larger progenitor facet (Ivan's discontent) that got creased from the "unhappy" fragment of the ancestral "Ivan's contentment" facet. These labels are the *meanings* to Cora of the fragment facets she now can see as distinct because she has created a new analogue perceptual function, though nobody else can perceive the differences at all because they have not developed such a perceptual function.

In crumpling, the physical way the fragmentation happened is that more and more energy was stored in the paper, the tighter the bend, and eventually that energy was dissipated when Ivan made the distinction between his *actions* reliably clear to Cora, with the result that the error in his controlled perception is rapidly reduced by her *actions*. The formation of the crumpling crease reduced the rattling and the energy expended by both Cora and Ivan in controlling their main perceptions, both of which have reference values that amount to perceiving Ivan to be content.

The meanings of Ivan's or Cora's "Words" are in the actions that affect the perceptions controlled by their use, not in the labels themselves. As we saw in the "H↔A" example (Figure II.6.6) an increasing contextual stress increases the likelihood of a choice being forced. If Cora is not controlling for a reduction in Ivan's crying, his variation in his sound and movement patterns will perhaps be interesting, but would not affect her actions. Only as her stress of uncorrected error increases beyond her tolerance limit will there come a point at which she will act one way or the other. Her perception moves from bend evident in the low-stress segment of Figure II.6.6 or Figure II.6.8 to the crease regime in the higher stress area.

The above discussion treats one facet as similar enough in two people that they can each use their version of it to act cooperatively. This facet has no direct connection with the complex context in which Cora's interaction with Ivan happens, but it does have an indirect connection. That context is both historical and environmental, and affects what each actually perceives. So far, we have not addressed the context, because we have not needed to know more than what each other's actions do with respect to the fragmentation of the facet in question.

At this point, Cora is able to put words on her discriminative perceptions of Ivan's different crying and moving patterns, but if she is entertaining a friend (who we call Fiona), Fiona cannot perceive the distinction. When Ivan cries, Fiona may say that he needs feeding, but Cora may say "No, that doesn't mean he is hungry; he is saying that he must have a pin or something stuck into him". Fiona cannot perceive the distinction between Ivan's "words", but she perceives that a distinction exists. That distinction has the function of separating for Fiona two states that Cora fixes by different actions, even if Fiona cannot perceive what Cora can.

With sufficient opportunity to observe these two states simultaneously in both Cora and Ivan, Fiona might be able to learn Ivan's language of sound and movement. She will do so only if she perceives that there is a discrimination to be made, and moreover that she will be less rattled overall if she does learn to make it. If she does not perceive that there exists a distinction to be made, she is unlikely to perceive from simple statistics the distinct patterns that to Cora are clearly different "Words" performed by Ivan.

Fiona is experiencing what Ivan will experience many times as he grows up. He will encounter words that distinguish states that he does not discriminate, and will have the opportunity to observe actions by others that allow him to build the requisite perceptual functions to control for those same discriminations. The example of the coherence between a green light and a "WALK" sign as opposed to a red light and a "DONT WALK" sign would be an instance.His mother's different actions that change what he is able to do (perceive himself on the other side of the road) form "pressure" for Ivan to discover the perceptual patterns, usually multi-sensory, that correspond to the different actions he observes others to produce. That "pressure" is the rattling associated with Ivan's inability to control whatever as yet undeveloped perceptual function it might be that allows others to act differently in some, for him undefined, context.

Although this little story involved a social setting, the focus is on Ivan and his development of words through the different sensory patterns of Cora's different actions. Baby Ivan may or may not perceive Cora as an independently acting agent. That is irrelevant. What matters is that different things he does with his muscles have different effects on perceptions he controls, and that allows him to refine his action pattern to improve the precision of his control so that he can usually get fed or get a pin removed quite quickly when his perception moves away from its reference value of "content" into either of those other states.

II.7.4 Crumpling, More than a Metaphor? 2. Shockwaves

In the crumpling situation, one fold, whether bend or crease, exists in the context of the rest of the sheet, which has other bends and creases. Together, these bends and creases transfer the force applied in the experimental situation analyzed by Andrejevic et al. (2021) between the piston and the bottom of the cylinder in the form of physical pressure (Section II.6.2). A bend can transfer vertical force between the upper and lower sheets it connects, because the energy that was supplied by the piston is stored within the bend.

A crease does not have this internally stored energy and cannot transfer vertical force in the same way. It can, however, transfer force between its crest and a sheet on which its crest rests, or that rests on the crease. When a bend suddenly becomes a crease, the change in the application of vertical force to the neighbouring sheets happens suddenly, causing a kind of shock-wave that spreads in all directions through the paper. The changed stresses might well create new creases from bends elsewhere in the sheet that were already tight and near their critical points, in an avalanche process like that of the "Bomb in the Hierarchy" (Section 6.5).

Let us consider the effects of the shockwave of crease creation in paper a bit more closely. We start with the finding by Andrejevic et al. (2021) that the total length of creases in the crumpled paper grows logarithmically as a function of the number of episodes of crumpling. If the total crease length increases while the average facet size decreases, the average crease length, a variable not recorded by Andrejevic et al. (2021), must also decrease. Since the energy released when a bend converts to a crease is proportional to the length of the resulting crease, the energy of the average shockwave that must be dispersed through the rest of the paper sheet decreases with the age of the sheet, if age is measured in the number of crumpling episodes it has experienced.

So it is as babies become children, children become teenagers, and teenagers become adults, to the extent that mature individuals may never perceive any shockwave effects large enough to involve the creation of new category discriminations. The more different environments they encounter, the more categories of object and of object properties they can develop, and the less likely they are to find that their previous categories need revising as a consequence of some new category being created in their mind.

Accordingly, as is commonly observed, the older a person is, the more probable it will be that they become "set in their ways". They tend to hold onto old beliefs, old ways of controlling well developed perceptions, and old lifestyles insofar as environmental changes permit. When circumstances do not permit, the older a person is, the more interior rattling they are likely to experience and the less able they are to reorganize so as to move their structure to a calmer location in their changing environment.

If one or more new creases are created by a shockwave resulting from crease-formation — facets fragmented — these in turn might propagate their own shockwaves that induce further new creases, in an avalanche similar in principle to the "Bomb in the Hierarchy" (Section 6.5), which is likewise fundamentally induced by propagating the effects of positive feedback loops.

Avalanches may happen at any age (of a sheet of paper or of living control systems), but tend to become smaller and less frequent with age. In humans, an avalanche might be given the name of "revelation", "epiphany", "insight", and the like, which usually are accompanied by feelings of great calm and stress relief. One suddenly perceives consciously that one "knows",which is something many teenagers and young adults seem to experience[22].

The shockwave avalanches of creasing or of epiphanies have something in common with the creative inventiveness of autocatalytic processes (Chapter II.2) when the increasingly complex soup of inventions suddenly forms a new homeostatic loop, perhaps in the form of a technological revolution. At that moment "everything just fits" and a new stable structure has been created. The new and revolutionary autocatalytic loop was not created by an avalanche, but its effect on rattling is very similar. Rattling surges and then is reduced below its earlier stable value.

Both the epiphany and the completion of an autocatalytic invention loop are reorganizations — of interactions in the soup and of the human' perceptual control system — that very suddenly lead to a much less rattled, calmer, configuration. In palaeontology, this effect is called "punctate evolution". It happens when a configuration such as an ecology is on the verge of instability or chaos and a small amount of energy applied at one place can cause big changes elsewhere. There is a short phase of much instability and chaotic interactions and events, followed by a much longer period of relative stability. Emotionally, this often happens with what are experienced as great insights — short instability followed by enhanced calm.

22. And, I might add, research scientists (from personal experience).

II.7.5 Crumpling, More than a Metaphor? 3. Teenagers

Avalanches cannot be large in a sheet of paper that has had but few prior crumpling episodes, or in a baby or toddler that has perceived very few distinctions that discriminate between categories of objects. Avalanches will also be unlikely when an old sheet of paper has had many crumpling episodes and has very small facets, or when an old person has become "set in their ways" because of their finely intricate network of interlinked categories.

There is, however, an intermediate stage between pristine newness and aged stiffness. In humans we sometimes call this stage early maturity or the teenage years. In paper crumpling, this is a stage at which the paper still has some fairly large facets but also many smaller ones. In humans, this is the age at which most great mathematical discoveries are made, and at which persons tend to experience religious epiphanies and conversions[23]. These are also the ages at which individuals tend to take risks with their own person, and take up dangerous sports that were not family traditions.

Returning to the question of the propagation of a shockwave when a new crease suddenly is formed from a bend, the initial shock becomes less and less energetic the shorter the length of the new crease. The same would be true of the shock of fragmentation of one perception into two that are distinguished by different kinds of control actions used in perceptibly different contexts.

A newborn baby may have some genetically built-in biases for perceptual types, but whether it does or not, its first categorical splits will be of much wider-ranging categories than will its later ones that are splits of categories already split, perhaps into so many generations of smaller sub-categories that the ancestral relationship is essentially invisible. Provided the shock of fragmenting a category does not cause an avalanche, the energies involved in the formation of later creases will have more localized influences, measured in terms of how large a proportion of the growing perceptual hierarchy they affect strongly.

A newborn baby might also have a large catch-all category akin to the original sheet of paper to be crumpled, We could perhaps label this category "Mystery". This "Mystery" category corresponds to no analogue perception, but if a new perceptual function is developed in the analogue hierarchy, the actions to control the new perception might well fragment a new category-facet from the "Mystery" ancestor. In later life the Mystery category could be ancestral to the sciences and religions, but we will not worry about that here. We will use "unhappy" and "content" as our examples.

The "unhappy" category might fragment into "hungry" versus "in pain", or some such. That discrimination is a big deal, because different sensations accompany the actions Mother (Cora) uses to move a perception into the "content" category. If Cora doesn't provide a "nipple in mouth" sensation when Ivan's perceptual category is "hungry", but instead fumbles around his person seeking the source of Ivan being "in pain", or vice-versa when Ivan is "in pain" but not "hungry", Ivan's perception of his own state will not easily move to the "content" category.

Ivan's ancestral categories do not vanish when he makes finer and finer discriminations among the the perceptual complexes that are suited to different control actions. Later in life, Ivan will still have his initial large categories as ancestors to the many smaller fragments created by Ivan's experiences. The "in pain" category might have differentiated into such fragments as "scraped myself" and then into "scraped my elbow" and "scraped my knee", while another fragment with "in pain" ancestry might lead to "stomach ache" which would have a different ancestry from the category "skin damage" that was an earlier ancestor to the "scrapes" which are now distinguished from "cuts". All these are Ivan's discriminations that depend on different actions used in control of the subtly different ranges of

23. I have no statistical data on this, but it is my impression of what those whose personal stories I have heard or read have said about when these religious events happened to them.

perception on the two sides of a crease. Different successful control actions, early in Ivan's life actions by Cora, and later by himself, would lead to category fragmentation.

Now let's look further into what we might call "the problem with teenagers", a time when parents begin to feel bereft of the lovely child they had been able to help through crises who now has turned into a surly creature seemingly bent on conflicting with its parents, and insisting that it doesn't need or want their help. Perhaps it also plans to leave home very soon for an independent life (except for purposes of laundry and some occasional financial contributions). Not all, or even most, teenagers fit this pattern, but enough do that it has become a category, a popular caricature of that stage of growing up, like the "Terrible Twos".

Why does this apparent antagonism toward the parents so often happen at that age? The shockwaves implicit in the crumpling metaphor, or its PCT equivalent in perceptual category subdivision, may be part of the answer. It's a question of whether there are enough fragments to support an avalanche, but not so many that they are so widely interlocked that they can quickly and widely redistribute the energy of the shock through the structure so far crumpled.

In the teenage years, the number of categories distinguished by both language and by alternative means of controlling the corresponding analogue perception has grown large enough to sustain avalanches but is not yet large enough to be tightly linked by membership of categorical perceptions in objects, or more particularly in the abstract syndromes that are modular objects, components of the teenager's social environment. Earlier, only two such modular objects, "family" and "school friends" have been available to children in wealthier societies, but the family now has little new to offer, and school friends often scatter. The teenager will encounter new groups of people, and from them will learn to make new creases across what had been stable fragments.

Each new crease will generate a shock, and each shock has the potential to launch an avalanche of subsidiary shocks, resulting in new revelations or insights. Having achieved a new insight, the teenager is likely to perceive that her parents are totally ignorant in her new area of knowledge. After all, they did not offer the information that led to the new epiphany, so they probably did not know it. Parents therefore are not the reliable supports they had seemed to be. They are frauds. This is likely to happen more than once before the cascade of large and small avalanches subsides because the remaining category facets have mostly become too small, and too interlinked into novel abstract objects, to generate large shockwaves when a bend in one collapses into a crease.

When Ivan reaches a mature middle-age, he has developed many precise categories from his infant large categories. With small categories, the shock of discovery is usually small, and an avalanche will usually create small insights, not life-changing revelations. The latter can happen at any age, but the more linkages Ivan creates as he gets older, the more stable are his categories, and the less likely it is that a small shock will propagate far beyond the category initially fragmented.

Among things Ivan has learned from the parents rather than from general experience meeting people of all kinds is a set of control actions we call social mores, which he is therefore likely to distrust. We are here verging on the kind of sociological discussion better left to Volume 3, but to some extent we have derived some "crumpling" PCT reasons why the caricature of the Terrible Teenager might be valid for some teens, but not all.

Why do not all teens fit the caricature? Some parents have allowed their child to have more experiences free of parental supervision, and to learn to interact with them when they already have categories fragmented by their own discoveries through experience in dealing with new situations and people, rather than having been fragmented by interactions only with parents and school friends with similar backgrounds. The Teenage Years are not as new (and therefore not so strongly rattling) to widely exposed children as they are to children who are always protected from adverse (and rattling) experiences by their parents.

II.7.6 Crumpled Humour: Riddles, Jokes and Buns

"A Bun is the lowest form of what"? How do we crumple that? Is what wheat, wit, or what? Is a Bun a pun? No it isn't, but the whole sentence is, because Bun sounds very like Pun, and what, through the intermediary of wheat, sounds like wit, of which a pun is said to be the lowest form. Cockney or Australian rhyming slang is similar. An Australian and I had a mutual friend called Gerald, but my Australian friend called him Sydney, or sometimes Sydney Morning. Why?, or should I instead ask by what train of intermediaries did Gerald become Sydney? I leave that puzzle to you, unless you want to cheat by reading the footnote[24].

Two objects may be separated by one crease or by many across perceptions of different properties of the objects. The objects have many properties, all of them in common except for the distinction created by the different uses available to the objects distinguished by values of the creased perceptions. A bun can be eaten, but a pun cannot, except perhaps in the metaphorical context of "eating one's words".

In the case of the *words* "Bun" and "Pun", they are phonetically distinguished by the sounds of their initial consonants. These small acoustic differences allow the possibility of a crumpling crease that creates a phonemic difference in a facet-object that was a wider phoneme class—bilabial aspirated stop. The acoustic difference is in the timing of the aspiration and the moment when the voicing starts of the vowel that is written following the B or the P. Creases based on analogue phonetic perception differences within phoneme classes are quite deep into the linguistic categorization of classes, the result of many crumplings and much shortening of the shortest side of the resulting facets.

However similar the words, a Bun you could eat is a very different kind of object than a Pun you may laugh or groan at. That very significant difference in the semantic or pragmatic referents of the two words in English is the result of many perceptual differences that have nothing whatever to do with language. The crease that forms between edible objects and language objects is a very early one, perhaps even generated by evolution because it allowed a newborn baby to distinguish between human voices and other sounds from the environment around mother and baby. You can yourself assuage hunger by eating an edible object, but you can assuage hunger using language only through the actions of another person.

Within each class, language and not-language or voice and not-voice, many creases across different perceptions occur before the baby learns to distinguish the words "Bun" and "Pun". It is a kind of distinction that exists in some languages but not in others. The same creases may be learned at the same time within, say, "Bowl" and "Pole" or "Bin" and "Pin", or, more probably the B-P contrast is learned separately as a crease between each pair of words before it is recognized to be available to distinguish a variety of word pairs, and long before long-trained linguists — scientists, not random native speakers of English — identify such characteristics as aspiration or vowel onset. The distinction between "Bun" and "Pun" *as words* is a long way down the creasing tree.

On the non-language side of the pun, the difference is much deeper. Probably very early in the creasing maturation process, things you can feel by touch are differentiated from things you can talk about but not touch. This might well happen soon after the distinction between language noises and non-language noises. "Pun" is an example of a language noise but it does not label an object one can feel by touch. "Bun" is a language noise that labels something you can touch, and an example of a language noise very close in the phonetic language creasing tree to the noise of "Pun". The Bun object class, which

24. The sequence is "Gerald" rhymed with "Sidney Morning Herald", abbreviated to "Sidney".

includes both its language-noise and its non-language form, is very close to the Pun object class in one perceptual dimension, but is separated by deep creasing trees in most other perceptual aspects.

Many, if not most Puns are of this type, using a wide semantic difference but a small phonetic difference between the phonemes it contains. Crossing such a boundary crease is energetically cheap, as compared with crossing a boundary between semantic classes, such between the meanings of "bun" and "pun", or "beer" and "peer". Semantic facets are in distantly different multidimensional sub-spaces from phonetic facets, but both start from the same first few creases in the clean sheet of a language being newly learned by a baby or a foreigner. Neither baby nor foreigner will understand the pun as humour until much later.

Labels are different. A label is something the sounds of Bun and Pun have in common with their semantics. When one hears "Bun", one is likely to also perceive both the sound and the semantics of both sound and semantics of "Bun", and perhaps imagine eating one. The sound of Pun may not be evoked by the sound of Bun, but just as was the case of the A-H example (Figure II.6.6), the higher level context of the semantics of "wheat" may induce a perception of the sound of Pun because of the very short crease that separates them. The sound of Pun induces a perception of the semantics of Pun, and from there to the perception of "wheat" as an phonetic neighbour of "wit", both of which are in the same semantic domain as their mates in the parallel sentences Bun-wheat and Pun-wit.

The incongruity between the close phonetic distance between the two phrases using Bun and wheat versus Pun and wit on the one hand, and the huge semantic incongruity between them on the other is the essence of much humour. The phrase "Things are not always as they seem" applies, but that phrase also applies in situations far from humorous. It is, instead, one way of describing a complex set of crumpling creases that recur in very different circumstances, including camouflage and deceit (Chapter xII.15). It is a motif of crumpling added to the catalogue of perceptual control motifs we have been building and will continue to build throughout this book.

The set of perceptual property boundary creases limit the kinds of perception about which that particular facet could potentially provide information. Andejevic et al. (2021) suggest that the average number for crumpled paper is about 5. Without any other evidence, it is tempting (and no more than tempting) to suggest that a facet represents the list of properties that *consciousness* may take into account when defining a new perceptual function represented by a crease. In turn, they suggest that categories are at least sometimes an aspect of conscious rather than non-conscious perception. This list is discussed in the section on consciousness as a tool for the "Mechanic" of reorganization (Chapter II.7).

Puns may be the lowest and simplest form of wit, but their inverse, the long and complex Shaggy Dog Story cannot be far behind. Whereas the pun reverses the semantic domain at the level of label, the Shaggy Dog Story does the domain reversal at the start of the story, usually but not always siting the whole story it in some fantastic Universe that exists only in the mind of the reader or hearer after it has been carefully built up over a long rigmarole of embellishment and side-tracks. The punch line, in contrast, seldom exceeds a short sentence.

The essence of the Shaggy dog is in two significant features, the length of the story to set the stage, and the speed of the switch in the punch line. The switch may take a variety of different forms, from a moderately simple pun to a shift from a complex fantastic premise to the destruction of the premise in the punchline. This latter form is by no means as simple a form to devise, perform, or analyze as is the simple "Bun". Here is an example, much shortened[25].

A mid-level major league baseball team in spring training was being watched by a horse from the sidelines. After a short while the horse approached the team manager and said that he could do better

25. I have no idea of the provenance of this particular Shaggy Dog, so I cannot assign credit.

than the players he had been watching. The manager pooh-poohed the idea, but the horse persuaded him to at least give him a tryout.

So the horse was provided with a glove for one hoof and was sent to the outfield to see if he could catch a fly, which he could, while covering much of the complete outfield with the speed of his running. And he could accurately and quickly kick the ball back to the infield to catch a runner off-base. Next, the horse was fitted with a way to hold a bat, and put against the team's best pitcher. The horse hit a home run off a majority of the pitches, had a perfect eye for whether the pitch would be a ball or a strike, and usually hit long line drives from those within reach, whether ball or strike.

Agreeing with the horse's self assessment, the manager got him signed up as a team member. The human team members were so buoyed up by the thought of having this incredible player on their roster that they greatly improved their own performance on the field. The fact that they had a horse signed up as a team member, however, had to be hidden in case the powers that be made a new rule limiting play to humans. The team, nevertheless, found themselves in the World Series. They came to the bottom of the ninth one run behind with two on base, and the horse was told to go and pinch-hit.

After a long argument with the umpires, nobody could find a rule prohibiting a horse from playing, so they let him hit, and he laced the first pitch for what might have been a triple for a human player. But the horse just stood at home plate and refused to move, while the team members urged him to run at least to first base. Eventually, the horse was put out at first, and the team lost the game and the series.

After the game, the manager asked the horse why he didn't run, and he replied: "Don't be absurd. Who ever heard of a horse running bases."

A Shaggy Dog Story has a much more subtle twist than does a simple pun, but the essence is the same, switch the pragmatic (in most Shaggy Dogs) field between the lead-in and the twist that the reader-hearer is required to work out. The humour is in the failure of an expectation that has been built up to a state of near certainty.

Jokes come in a wide variety of forms, including the "Three XXX walk into a bar" motif, which often consists of the rabbi, the minister and the Imam, or the English teacher, the philosopher, and the mathematician, or the sheep, the snake and the crow, all exaggerating what is generally assumes about someone like them. The joke is again in the failure of an expectation, such as the mathematician proving that there are no people in the world, or the crow asserting that since the sheep can't fly, it is not a living creature. No matter whether the premise or the punch line is absurd, it is a failure of expectation. The perception built by the body of the joke is not the one presented at the end.

The building of the perceptual complex is an analogue of the Big Lie we will discuss when we deal with politics, in that both are built on and in a fantasy world that at the end turns out to be different. The main difference is that the listener-readers of a joke are not supposed to believe the fantasy world to be true, whereas the targets of the Big Lie are.

Whether Riddles should be treated as humour depends on something other than the puzzle aspect at the core of each instance. To solve a riddle involves a search through the dual network of the crumpled space to find a perception with all the properties demanded by the presentation of the premises. If the answer, once discovered, is in a different semantic or pragmatic field than expected it is likely to be humorous, especially if, as with a pun, the semantic of pragmatic field requires starting many crumpling levels above one or both of the anticipated and the correct answer's fields. "What walks on four legs then two legs and then three legs" is an ancient classic, with the answer "a man" treating an old man's cane as his third leg.

All in all, the essence of humour from the viewpoint of PCT and Crumpling Theory seems to be the creation of the perceptual properties of a variety of objects that abut in some perceived properties, and the switch from one selected by the premise to an adjoining one selected in the switch. Rattling does not seem to be involved except insofar as the twist may rattle a part of the perceptual control hierarchy of the listener-reader. A good PCT explanation of humour remains to be developed — at least I have not heard of one.

We return to more serious considerations, linking rattling with crumpling.

II.7.7 Rattling and Crumpling: Tensegrity and Pressure

Rattling, as described by Chvykov et al. (2021), is a measure over a partial or complete organization of any size. Crumpling fragments objects into classes, as does category perception. Both can be described in a high-dimensional space of description, but the axes of the two spaces differ. Rattling tends to move the organization as a whole into calmer parts of the description space, and the main axes of its description space are each related to how control of one perception (one-dimensional) rattles the control of another. Crumpling does not occur unless some facet sustains a crumpling pressure that pushes or pulls the components together, while simultaneously creating new sub-classes of pre-existing classes.

Rattling changes the organization of existing components of an organization, while crumpling adds more components to an existing organization. Can we reconcile these relationships between crumpling and rattling effects? Particularly, can we identify a pressure source due to rattling that applies to an entire organization created by crumpling and White Box building?

Perhaps we can, at least in organizations built from perceptual controllers. Rattling tends to lead organizations toward calmer regions of their description spaces. These relatively calm regions are places where two things co-occur. One is that adverse side-effects of control of one perception on control of a different perception by the same or a different individual tend to reduce through reorganization of one or both of the controllers within their control hierarchies. This applies to all the pairwise interactions among the controllers in the organization. The other is that the controllers that belong to an organization are individually calmer if they are shielded by each other from external sources of rattling. Both effects tend to pull controllers toward each other if they interact.

“Pulling toward” is equivalent to “being pushed toward”, so far as the participating entities are concerned. It opposes a tendency to separate entities, which, as we saw above, is a feature of creasing. As Scarr (2014) pointed out, the opposed pull of two wires creates a point at their junction that can act like the end of a rod or a wire connected to the junction. We have seen this as the “stiffness” motif in perceptual control (Section 8.1). It is easy to see as a force when the wires are not collinear, and indeed “pulling” and “pushing” both have the connotation of using force.

Rattling within an organism can form tensegrity “wires”, that dynamically seem to pull the organization and its parts toward low-rattling regions of the description space, whereas, as we saw, crumpling provides “rods” in the network dual that oppose being pulled together. Perhaps we should add also the hysteresis effect, which tends to keep things on whichever side the happen to be of a bend or crease. The hysteresis effect is not the equivalent of the tendency of crumpling to push apart perceptions that naturally fall of opposite sides of a crease. There is no “pushing apart” in hysteresis. There is just a tendency to stay put longer than the raw data might warrant.

Hysteresis inhibits items that would naturally be situated near the crease from tipping to one side and then back again repeatedly, in a mode called “chattering”. Rattling, on the other hand, moves things around. More rattling implies more uncertainty of how fast the forces affecting perceptions change, which implies that the higher the rattling the more likely is a perceptual value to cross a crease and “fall over the crest of the hysteresis loop". A perception is likely to be easier to control if it is in a natural place, typically a place nearer its reference value—as a choice of category.

Even though rattling causes the location of an entire organization to tend to move toward regions of lower rattling in its description space, rattling does not exert any actual force anywhere in the organization. The units of the organization change if the organization is to change, and when the units are control systems, their rattling is in the form of low Quality of Control.

A bundle of rods and a bundle of wires do not a tensegrity structure make. They have to be appropriately connected. Several wires, on average at least one per dimension of the space, must pull on the end of any one rod. Do rattling and crumpling within an organization necessarily link controllers in this way? Are there conditions that enhance or reduce the likelihood of such connections occurring?

Well, yes, rattling and crumpling do link controllers this way quite naturally. If you think about the facet-objects of a crumpled high-dimensional "sheet", every crumpling boundary represents a perceived property value difference across the boundary. As we discussed earlier, the crease creates two distinct "daughter" categories from one, distinguished from each other by an analogue perceptual value range difference. These are distinct categories, not only of perceptions of properties, but of entire objects or object classes.

If one is controlling for perceiving one of these categories rather than the other, one is controlling the analogue value within a two-level hierarchy. Although after the crumpling event there are more components in the perceptual control organization (three, where there had been one) control within each sub-component is likely to be better after the shock of separation has subsided than it was within the larger category. Rattling "pulls" the organization toward this development of a perceptual control hierarchy, while crumpling pushes the newly created subcomponents apart. Overall, the hierarchic structure is calmer than the structure of large undivided facet-objects.

These are the conditions in which we might expect tensegrity effects to occur, even in a high-dimensional space of description. The total force compressing a physical rod must be the total force of the pulling wires in the rod direction. We should keep in mind that each rod and each wire is one-dimensional, no matter how high the dimensionality of the space of description, and that if a wire meets a rod, they meet at some angle α, a scalar value. If that wire pulls with force F at any instant, the effective force in the rod direction is Fsin(α), which could be positive or negative in the direction of the rod. It is the sum of these wire forces that compresses the rod.

At any moment, most sums will be near zero, or so it would seem at first sight. All is not as it first may seem, however. The crumpling pressure induced by rattling reduction is the sum of the average rod compressions in every direction, and unless parts of the organization escape rattling by the main body of the organization (as did the American Pilgrim Fathers by leaving England), rod compressions in any localized part of the organization as well as over the entire organization will be distributed in the same long-tailed way as all the variables we met in our discussion of rattling.

Within an individual organism, parts of the organism cannot easily escape rattling by physical separation, but if there is a local source of high rattling, such as gangrene in a limb, amputation may serve to reduce the rattling the limb causes the main body of the individual. This is rarely done, and almost never is it done by the perceptual control actions of the gangrene-rattled individual.

Rattling, Crumpling, pressure, and facets are just metaphors for components and functions of a living organism. Like all metaphors that seem useful, they will be flawed. The question to ask is in what way these particular metaphors might be flawed. Are their problems central to the foregoing arguments or not? The fundamental reason rattling tends to move organizations to calmer regions is simply like the winter-leaf that is less likely to blow far from a sheltered place than from a windswept area. That phenomenon, moving more easily from a highly rattled place than from a less rattled place, is independent of what is being rattled, so it applies just as much to the relationships among control loops and categories in an organization as it does to the "smarticles" of Chvykov et al. (2021).

Chapter II.8. Reorganization and Consciousness

Beginning in Section 6.4 we have from time to time mentioned consciousness, and noted that what we perceive consciously differs from anything the Powers non-conscious perceptual control hierarchy controls. We consciously perceive entire discrete objects and complex abstract entities, but the perceptions in the perceptual control hierarchy are simply the magnitudes of scalar variables, variables that we interpret as properties of the objects we consciously perceive.

We perceive bodily states and moods such as contentment, nausea, pain, joy, hunger, fear, fright, love, unease, satisfaction, and so forth. Such states are consciously perceived. How do they relate to the operations of the perceptual control hierarchy? Are they associated with particular states of the external environment reported by sensors? Associated, assuredly. Cause or are directly caused, assuredly not.

The same sensory input can lead to very different conscious perceptions, as attested by the probably apocryphal story of an aide-de-camp of a general, both on horseback in a fierce 19th century battle. The aide says to the general: "*By God sir, your leg has been shot off*". The general looks down, and answers "*So it has*". He had not consciously perceived the sensory inputs that must have accompanied the sudden loss of the leg, his consciousness being entirely occupied by the progress of the battle. Under calmer circumstances, could such a situation ever be sufficiently plausible to become a long-lasting story?

Figure II.8.1 suggests how biochemical homeostatic loops among the "intrinsic" (biochemical) variables might interact with the perceptual control hierarchy to allow such an extraordinary response by the general. We assume that interactions of these kinds, involving different hormones, enzymes, and bio-active chemicals are responsible for these varying moods and perceptions, perhaps affecting loop gain values or other intra-loop parameters, perhaps affecting the inter-loop parameters of the hierarchy and causing reorganization conscious or non-conscious.

Even when we consciously control a perception, it is never a perception of an entire object. Rather, it is a perception of something about the object, quite possibly a property for which control has not yet been incorporated into the non-conscious hierarchy, or for which the current environment provides no atenfels yet in the hierarchy. We addressed this problem in Section 6.4 by invoking Friston's Predictive Coding, or Free Energy" (Friston, 2010). Conscious perception is much richer than the perception of single properties in isolation, or than the constellation of properties of single objects.

In this Chapter we will be more interested in the problem-solving, perceptual control aspect of consciousness than in the fact that our experience tells us our local environment is a space full of separable objects or abstract structures. Clearly, however, these two aspects of consciousness cannot be well separated.Conscious perception may be concentrated by what we call "attention" onto a focal set of perceptible properties, but if this control aspect of consciousness works to allow control of a perception not controlled within the non-conscious hierarchy, those are properties of objects that are likely to have "White Box" functions that include atenfels likely to be involved in solving the problem at hand.

Our problem solving uses the interacting properties of the objects we might be able to use to solve the problem. Looked at this way, conscious problem-solving is building White Boxes in our imagined Perceptual Reality to emulate some feedback process we hope to find available through Real Reality (RR). These imagined White Boxes may well incorporate atenfels (in the form of smaller White Boxes) that do exist in the non-conscious hierarchy — they tell the problem solver "I already know how to do that, so don't bother with tedious calculation how to do it."

We are "inventing" or "remembering" what happens when we build new OOP objects (White Boxes) from components that are functions provided by (properties of) existing OOP objects (White Boxes). In biological organisms such as ourselves, each recurrence of the same solution to a problem results by a

Hebbian-anti-Hebbian (HaH) process (Chapter 9) in building a new ECU (Elementary Control Unit) into the Powers hierarchy — reorganizing the hierarchy. We start, however, in a quite different place, somewhere in a physiological homeostatic loop.

***q20.1 Sleep and reorganization

II.8.1 Reorganization in the Individual Revisited

Ignoring the rattling and crumpling of the last two Chapters, we look a bit more at reorganization of the non-conscious "reorganized" perceptual control hierarchy within an individual. This time, we take into account the biochemical and microbial homeostatic loops and networks that we discussed in the previous three Chapters and ignored in our discussion of rattling. Chapter II.2 and Chapter II.3 fill in some gaps in the earlier discussion of reorganization, though we still leave other gaps to be filled.

As, from a very different viewpoint, did the "rattling good" discussion of energy, entropy, and rattling in the last Chapter, they offer some notion of the biochemical aspect of the body of an organism, including our own, as being filled with many levels of external and internal control loops and homeostatic networks. The variables in these loops are concentrations of essential components of our internal environment, the "intrinsic variables" that Powers took to be the driving factors in reorganization. They are produced by and used by our cells, and by bacteria in the many complex ecosystems that constitute our microbiome.

We will not disagree with Powers about the importance of intrinsic variables, nor about whether they are all concentrations of biochemicals (for example Powers, quite reasonably as we saw above, treated Quality of Control (QoC) as an intrinsic variable). We can, however, take yet another different point of view, which was presaged in Figure 11.4 to Figure 11.6. This view takes the perceptual control hierarchy to be the external environment of a homeostasis hierarchy that exists within the biochemical, including the microbial and intracellular, structure of the body. Now, however, we link the homeostatic biochemical environment more closely with the processes of the perceptual control hierarchy.

The actions performed by the outputs of the various control units in this hierarchy vary the parameters of the the perceptual control hierarchy on both short and longer time scales (Figure II.8.1). Likewise, inputs to a homeostatic loop from outside it, such as the signal from the Output Function of an ECU marked "Reference" in Figure II.8.1 can also be seen as disturbances, because they have aspects of both — more strongly as References for the immediately preceding loop signal and more strongly as disturbances to the immediately following one (see Appendix 1).

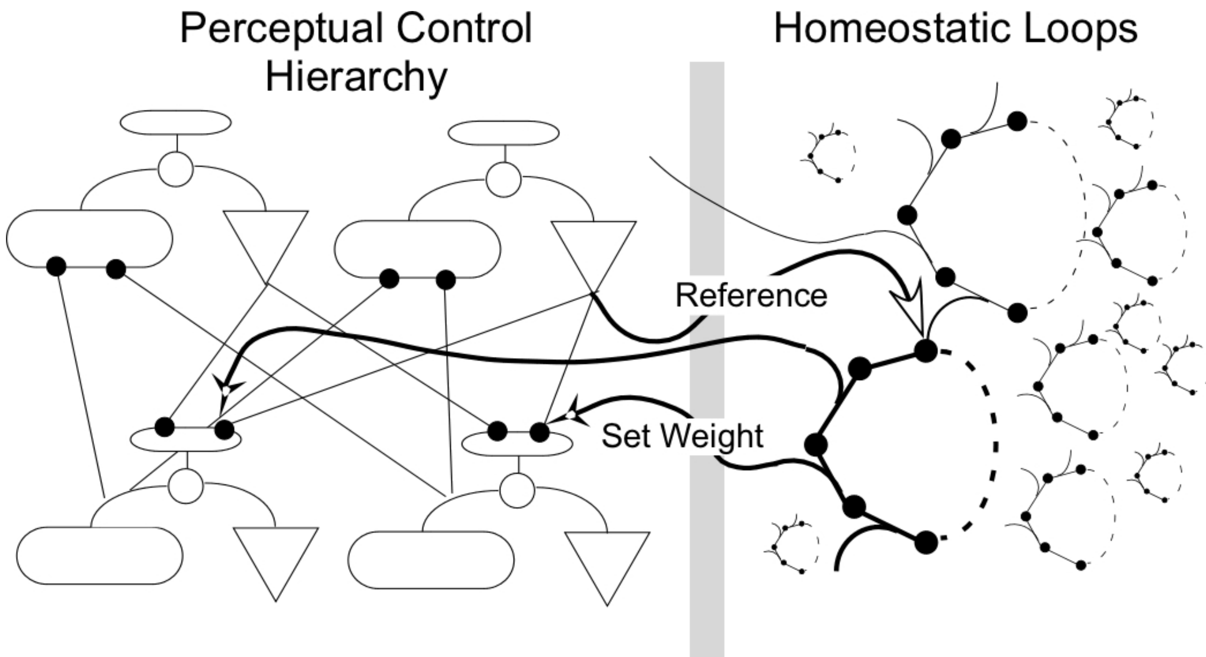

Figure II.8.1 Suggesting how a homeostatic loop hierarchy might interact with the neural perceptual control hierarchy. In this illustration, one of the outputs from the perceptual control hierarchy, along with a side effect of another homeostatic loop, contributes to the reference value for an action in a homeostatic loop, while another of the homeostatic loop variables adjusts the weight of a reference input function in the perceptual control hierarchy. "Reference values" in the homeostatic loop affect the concentrations of biochemicals participating in a reaction. Many of those biochemicals are among Powers's "intrinsic variables".

Later in this Chapter, we will take the hierarchy as the "inside" rather than as "the environment". These views are not mutually contradictory. Rather, they illustrate Powers's frequent exhortation to look from the viewpoint of the area in which you are interested. Figure II.8.1 suggests one variable in a homeostatic loop being disturbed by the changing output of a perceptual control loop (that variable is also disturbed by an output from another homeostatic loop). In stabilizing the homeostatic loop the outputs from two other of its variables affect the sensitivity of the reference values of two lower-level perceptual control loops to changes in the output of the higher-level one[26]. The outputs from the homeostatic loop array to the perceptual control hierarchy and vice-versa form hybrid loops, some of which might be control loops that affect the parameters of control on a moment-to-moment time scale rather than only by secular reorganization.

In Section 9.7 we suggested that there might be two parallel hierarchies of perceptual control, one creating analogue variables at several levels of complexity, the other selecting which of mutually incompatible possibilities was being perceived and controlling "logical" variables in ways we might call analytical thinking. The logical hierarchy was drawn as linking to all levels of the analogue control hierarchy and extending its levels "sideways" to perceptions of increasing complexity.

---

26. Such a set of connections would create a non-linear relation between the output of the higher-level control loop and the loop gain of that control loop with a magnitude of effect that would depend on the output from the other homeostatic loop. Whether such a relationship would ever be useful is anyone's guess. It's just a possibility that might be useful as an example.

Now we add a third leg to this array, a biochemical leg that consists of homeostatic networks in which many of the constituent products are the Powers “intrinsic variables”. The perceptual control legs communicate with the external environment by way of muscles and sensors, whereas the homeostatic leg communicates through material transport in both directions. The Logical hierarchy communicates to the environment only by way of its interactions with the analogue and homeostatic hierarchies.

Figure II.8.1 suggests that the homeostatic biochemical structure can affect the weights of the perceptual control hierarchy, which implies that it could do the job attributed to “reorganization”. But more than this, the biochemicals that interact with the neural synapses, where this presumably happens, include hormones associated with mood, emotion, energy, enthusiasm, as well as depression and malfunctioning of the hierarchy such as in schizophrenia or Parkinson’s disease.

Micheva, Weinberg, and Smith (2020) mention 42 different biochemicals associated directly or indirectly with the synapse, all of which presumably have their own effects singly and in combination. In the other direction Xu et al. (2020) show that, at least in mice, neural effects can even influence immunity. Both papers suggest by direct experiment that there is a complex two-way linkage between the operations of the neural and homeostatic hormonal structures, possibly akin to the sketch in Figure II.8.1.

As Figure II.8.1 shows, some of these homeostatic biochemical loops are affected in the other direction by neural processes through the operation of glands. The obvious implication is that these “trans-membrane” neuro-biochemical loops allow for the perception and to some extent control of feelings and emotions, and even of general health. .

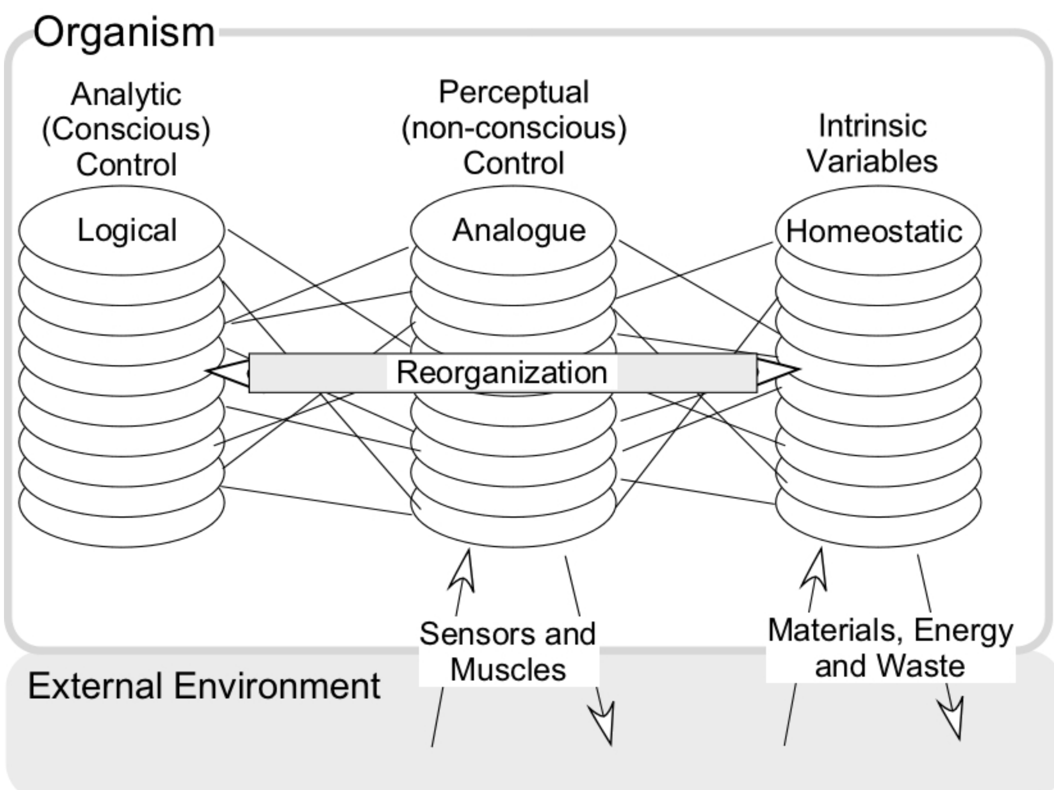

Figure II.8.2 Schematic of the tight relationship between homeostatic loops and the analogue and digital pillars of the control hierarchy that were discussed in Chapter 9. The Logical hierarchy accesses the environment only by way of the Analogue perceptual control hierarchy. The question mark on "Conscious?" hints that perceptions in the logical hierarchy may be often consciously controlled.

The biochemical structure consists of myriads of homeostatic loops that constitute a vast network. Some of these are among the various complex molecules, but at higher levels they can be considered to be among cells, including both fixed and mobile cells. Among the non-fixed cells, we must especially include the complex ecologies of the microbiome, which have a large contribution to the processing done at most exterior and interior surfaces of a mobile organism, such as the skin and the gut. The microbiome ecology might be imagined as threading intimately through the "homeostatic" column in Figure II.8.2, or as forming a fourth column of its own.

At the microbial level, the network begins within any one bacterium, which takes in some source of energy as food and emits waste that might serve as food or catalyst elsewhere in the microbial ecology, but not only there. The products of the various microbiomes on and in a complex organism are complex and may be food for the structural cells of the body, whose waste may reciprocally serve as food for microbes and other cells. These interactions are likely to form autocatalytic loops, and except in the case of infections and diseases, those autocatalytic loops will be homeostatic, in an interlinked network.

As we noted earlier, when one homeostatic loop buds off another, the two loops begin intermeshed, sharing some links, but there is a good chance that the continued production of new complex entities will short-circuit one of the loops, perhaps bypassing the original shared links. If that happens, the two become independent, each being the source of its own sub-family of homeostatic loops, some of which survive, and cooperate or compete with existing loops, perhaps by providing novel catalysts or anticatalysts for reactions that may or may not already be included in homeostatic loops. The global structure tends toward modularity, low-level modules being interconnected and forming larger modules of modules between which the boundaries are fuzzy

At all levels of this hierarchy of homeostatic networks of biochemicals there are waste outputs, which we also sometimes label as "side-effects" (as we do in Figure II.3.3). These side-effects or outputs of the many homeostatic loops change in concentration as the loops are disturbed, or as those that cycle rhythmically over time go through their cycles. The concentrations of these products are among the "intrinsic variables" that affect reorganization. Some of them may influence the firing probabilities of nerves with which they come into contact. Some of them are influenced by nerve firings. As in Figure II.8.1, the effects go both ways, as well as through the environment from the side-effects of perceptual control back into the homeostatic loops.

II.8.2 The Homeostatic Interface to the Control Hierarchy

When we added a logical hierarchy beside the analogue hierarchy, the domains were separated by a conceptual "membrane" that from the viewpoint of the reorganized perceptual control hierarchy consisted of the polyflop category perceivers. What might be the "membrane" or interface that separates inside from outside of this new hierarchy of homeostasis? I suggest that glia cells such as astrocytes, together with organs such as liver and pancreas, glands, synaptic transmitters such as dopamine, and nerves that terminate other than on muscles and sensors form the conduits that allow the conceptual membrane to be permeable. The conceptual membrane is the difference of kind that exists between nerve firings and biochemical concentrations.

A sketch of the parasympathetic nervous system in Wikipedia[27] shows 24 different glands innervated by the neural system that emit different biochemicals, each of which may in turn influence some aspect of the working of part of the perceptual control hierarchy and the sympathetic nervous system. In all, the biochemical homeostatic networks form a third structure to add to the interacting analogue and (conscious) logical control hierarchies.

What is the external environment of this third structure, which might be a hierarchy of networks? It is not only the perceptual and reference values of the other two hierarchies (though sometimes that could happen), but it is also their parameters, which are manifest in the strengths of synapses and the firing tendencies (polarization) of nerves, neither of which appear in our functional descriptions of the analogue and digital perceptual hierarchies. They are manifest in the parameters of the individual elementary control units (ECUs) and of the connections among ECUs both across levels and within levels.

Some of these membrane-crossing "sensors" and "actors" are widely distributed through the bloodstream and extracellular fluid, some are delivered very precisely to specific places or precisely depend on specific places. Widespread variations cannot influence specific neural sites, so cannot be candidates for modulation of localized parameters of the perceptual control hierarchy. If they interact with the perceptual control hierarchy at all, it would be through a generalized effect on substantial components, such as a generalized increase or decrease of gain in one or more types of interconnection.

For example, perhaps adrenalin might enhance the cross-link gain of the categorical interface between the analogue and logical hierarchies, enforcing choices such as "fight or flight", or more mundanely by enforcing a choice of meaning for an ambiguous word (Figure 9.10 through Figure 9.13) by increasing "task stress" or "contextual stress". Within the kind of "neural bundle" into which Powers grouped neurons so that their firings could be summed into a "neural current", a widely distributed biochemical could shift a global bias that would change the excitability of the individual nerves differently, changing the current in the bundle, the tolerance of a comparator (Section 4.6 and Section 5.1) or the gain of a processing element. Variations in the concentrations of other widely distributed chemicals might be perceived as changes in emotional states.

On the other hand, targeted cross-membrane effects, such as through glial cells known as astrocytes, might well modulate localized parameters. In the other direction, both the sympathetic and parasympathetic components of the autonomic nervous system influence the production of a wide variety of biochemicals, and are known to be implicated in many homeostatic processes.

Homeostatic loops of the kind discussed in this and the previous Chapter do not necessarily tend toward static values of their variables. As we said above, some may cycle rhythmically in the absence of input from outside. Even the circuit of Figure II.3.1 might do this if the time-constants of the integrators happened to be appropriate. Loops of this kind could serve as clocks or timers for intervals ranging from milliseconds to months or even years. When the nervous system is part of a cycling homeostatic loop, the properties of the neural parts of the circuit cycle with the rest of the loop, and the products output by the loop might change the properties of the perceptual control hierarchy equally rhythmically. Cycling phenomena are easily synchronized by small changes in external variables, so such cycling homeostatic loops may well be synchronized by the environmental rhythms of the days, the work week, or the seasons.

In this context, we next take a very simplified look at a review paper by Röder et al. (2016) on the

27. https://en.wikipedia.org/wiki/Parasympathetic_nervous_system#/media/File:
1503_Connections_of_the_Parasympathetic_Nervous_System.jpg Retrieved 20.12.08

metabolism of glucose. The paper goes into the detail of a wide variety of interacting homeostatic loops that influence insulin and glycogen levels that respectively reduce and increase the level of glucose in the blood. The loops they discuss involve different major organs of the body and a host of different molecular concentrations. Of these, they single out the pancreas, and specifically the Islets of Langerhans, as central to the stabilization of sugar levels in the blood, since it is they that secrete the insulin.

Of the other organs they discuss, however, the one of most interest for the purposes of this book is the brain. In this we see homeostatic loops that include in the loop both neural and biochemical activity. Introducing what they call "The brain-islet axis", they say (their references omitted):

> *Just as insulin exerts its effects on other organs and tissues, other organs interact with the pancreas to modulate insulin secretion […]. One of these interacting organs is the brain, which comprises the mutual brain–islet axis that interacts with the pancreas and vice versa. The pancreas is highly innervated with both parasympathetic and sympathetic nerve fibers from the autonomic nervous system. At the same time, insulin receptors are widely distributed within the brain, including the hypothalamus, cerebral cortex, cerebellum and hippocampal formation in humans, as well as the olfactory and limbic areas, hypothalamus—particularly the periventricular nucleus and the arcuate nucleus—hippocampus and the choroid plexus in rat brains. Lesions in various brain regions were shown to affect pancreatic hormone secretion. The destruction of the ventromedial hypothalamus results not only in insulin hypersecretion due to loss of the ventromedial hypothalamus-mediated inhibitory impact on pancreatic β-cells but also in higher glucagon levels.*
>
> *Glucagon secretion may also be modulated by the hypothalamic brain-derived neurotrophic factor via efferent nerves, whereas the melanocortin system directly reduces basal insulin levels by pancreatic regulation of glucose homeostasis stimulating sympathetic nerve fibers via α-adrenoceptors. Acting via α-adrenoceptors, norepinephrine also inhibits insulin secretion, which is an important aspect of the fight-or-flight response. The neurotransmitter Neuropeptide Y (NPY), which is mainly expressed in the sympathetic nerve fibers of the autonomic nervous system, also blunts insulin release, and the loss of NPY's inhibitory action results in elevated basal and glucose-stimulated insulin secretion as well as in increased islet mass. …*

In this first paragraph about the brain relationship, they mention "insulin receptors" in four regions of the brain in humans and five in rats, of which two are are the same in both species (hypothalamus and hippocampus). We might call these "insulin receptors" the sensors of perceptual control loops or homeostatic loops that cross the conceptual membrane between biochemical concentrations and neural firing rates.

Röder et al. are less specific in this paragraph about the efferent nerve connections that we might say form the outputs to the environment across the conceptual membrane from nerve firings to biochemical concentrations, since much of the evidence comes from brain lesions that might affect any part of the neural section of membrane-crossing loops. Nevertheless, they do mention "*pancreatic regulation of glucose homeostasis stimulating sympathetic nerve fibers via α-adrenoceptors*", which I am guessing to be in the adrenal gland. Those fibres could be seen as sensors on the biochemical side of the membrane or as outputs on the nerve-firing side.

Röder et al. do not discuss behaviour except to talk about what happens after a meal is eaten. We, however, are interested in the perceptual control loop in which the action output is the eating of a meal. What perception might be being controlled, and where does the reference value for that perception come from? If we look entirely from the perspective of the Powers hierarchy, in which all perceptions ultimately derive from sensors of the environment outside the skin, we get one answer: food is visible. It would fly in the face of most people's subjective experience to accept this as *the* controlled perception, with a permanent reference value of wanting food to be visible, and to eat when it is.

So we ask whether sensations from inside the skin might lead to a set of action outputs that result in the complex sequence that we call "eating a meal". Here we come up with an answer more compatible with experience — we feel hungry when we happen to have a reference value for degree of hunger that is less than our current perceived hunger. But this turns out to be one of what Powers called "dormitive principles", named for the idea that we go to sleep because we have an excess of "dormitive inducers". It explains nothing. Why do we ever feel hungry?

It may often be true that we feel "hungry for chocolate" when we see chocolate, even in a picture, but this is more likely to be an atenfel for the method of reducing the feeling of hunger than a component of feeling hungry. On the other hand, the physical contraction of the stomach is something that can be sensed, if not consciously perceived, and this sense might possibly be transformed into a perception of a level of hunger that exceeds a desired reference value. If we consider the passage from mouth to anus as a continuation of the skin, these sensors are just another member of the wide-ranging class of skin sensors, and control of stomach skin pressure can be achieved by filling the stomach —by eating. That approach suggests one kind of answer.

The quoted section from Röder et al. (2016) suggests another answer entirely. Translated into the language we have been using, variable values relating to insulin concentration are outputs from, or variables within, different homeostatic loops that involve insulin. These homeostatic loops affect and are affected by sugar concentration. They might be completed through the brain by way of outputs from the nervous system to organs that participate in those same loops, to organs that participate in other homeostatic loops that connect in some way back through the original loop, or, and this is the critical point, through behaviour that acts on the external environment, the action output being to eat — "classical" perceptual control.

Röder et al. show a loop (their Figure 2) that uses insulin level even though the loop passes through several organs. The effect of the loop is control of blood sugar concentration, since it produces an increase or decrease of insulin level that either enhances or reduces the rate at which the insulin depletes the blood sugar so that a more or less constant level of sugar concentration is maintained in the bloodstream. As we saw in Section I.4.6 on the Comparator, when a variable value cannot go negative, a control loop must be split into a section that reduces too high values and a different section that increases too low values. The same is true for the loop shown for insulin control, though both sides of the loop use the pancreas and the liver.

In Appendix 1 we show that from the viewpoint of any one of its nodes, a homeostatic loop acts as a control loop for the variable within the loop that enters that node. The external variable that enters the same node might be called a reference value for the loop variable, which would then be a perceptual variable. The corresponding external variable that enters the homeostatic loop at the immediately preceding node acts as a disturbance to the same "perceptual" variable (Figure II.8.3).

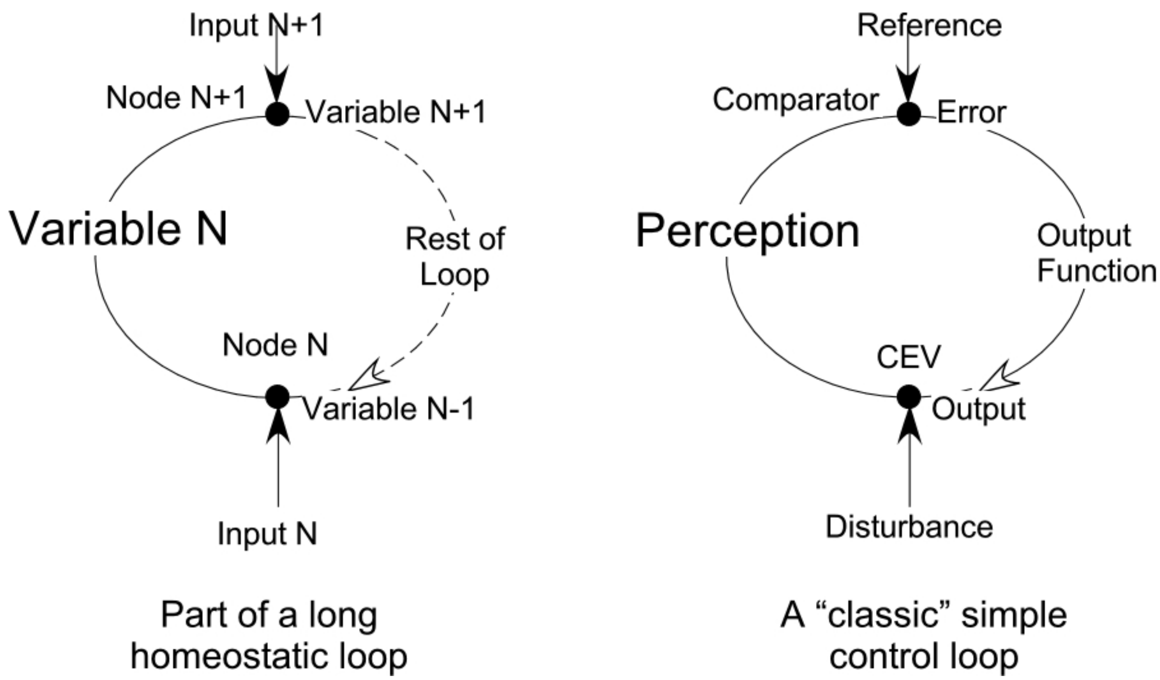

Figure II.8.3 A section of a homeostatic loop (left) matched to a corresponding simple control loop (right), illustrating how an input from external sources may act as both a reference value for one loop variable and a disturbance to another. In the control loop, the loop gain is concentrated on the output side. In the homeostatic loop it is likely to be greater in the rest of the loop than in any single leg, such as between Node N and Node N +1. In that case, Input N is a disturbance to Variable N, for which Input N+1 is a reference value input (while acting as a reference input for Variable N-1.

Figure 2 of Röder et al. could be seen as showing a control loop in which sugar concentration may be disturbed when one eats, but the salient controlled variable consequent on that disturbance is insulin concentration. How does this make sense within the conceptual framework of PCT? Should not the disturbed environmental variable be one that directly corresponds to the controlled perceptual variable? Well, yes, if there is only one. But in a homeostatic loop that is longer than two nodes, the loop variable entering any node has an effective reference value set by some following external variables and is disturbed by preceding external variables (Figure II.8.3 shows only one of each).

If one variable in the loop is called a perception that is controlled to match its reference input, then all round the loop the other variables are also controlled to match their local inputs, rather than to correspond with the arbitrarily chosen one labelled “perception”. Of course, for stability there must be an odd number of sign changes around the loop. In the control loop there is exactly one, at the comparator, after which the next variable, the error, is e = (r - p) in the classical loop, a sign change from the entering loop variable p.

In practical control loops, the sign change might occur elsewhere in the loop, in which case the error would be e = (r + p), r being the sign-reversed value of what it would be in a classical loop. In the homeostatic loop the necessary sign reversal or reversals could be anywhere in the loop, with the sign of the appropriate reference/disturbance input being set appropriately for the loop to function properly in its larger context.

As we saw all through the discussion of autocatalysis and homeostasis, the outputs of the nodes of a loop may be as important as the stasis itself. The more a loop variable is disturbed, whether you call that disturbance a reference input or a disturbance input or something else, the more the output of the preceding node around the loop must change in compensation. Any side effects of that output presumably will vary similarly. The output is at least a variable value that might appear and be used elsewhere, either focused in one place or diffused globally through a wide region. I think here of chemical concentrations transmitted by constrained tubes in the way that blood is distributed around the body, as compared to a diffuse "soup" that has no targeted locations. According to Röder et al (2016), the pancreas delivers different biochemicals by both means, some by targeted ducts and some by diffusion through the bloodstream that goes everywhere in the body.

Consideration of these "cross-membrane" interactions offers a new look at reorganization, which now can be seen as a process that could function on all time-scales from a spur-of-the-moment choice of output modes to a lifelong process of maturation and increasing development of ever more complex perceptions. Later in this Chapter, we will look at a possible relationship between consciousness and ongoing reorganization, as one gains skill in controlling a particular perception.

Above, I suggested that "eating a meal" might be a consequence of the perceptual control hierarchy acting as an integral component of a big homeostatic loop that extended though the perceptual control hierarchy, the external environment, and the biochemical homeostatic system. But in Figure 11.4, we showed a quite different kind of loop that traced a path around the same three domains. That path included another component between the "Intrinsic Variables" and the perceptual control hierarchy, a component labelled "Reorganizing System". Perhaps we might now be in a better situation to think about what that is.

II.8.3 Quality of Control and Reorganization Rate

Figure 11.6 showed a hypothetical control loop that set parameters for a single elementary control unit, using the e-coli process as its output and the quality of control of the subject loop as the perception being controlled. By giving this "control of control" loop a reference input from elsewhere, the subject loop might be adjusted to control optimally, after a long period of e-coli search, or to control "just well enough" after a short learning session.

Although to vary the required QoC is very similar to introducing a tolerance zone in which finite error is treated as zero error, a change in the tolerance zone width would seem simpler and easier to reverse if it were directly influenced. Now we have proposed a means to do just that, using an output of some stage of a homeostatic loop as the reference measure "from elsewhere" of the QoC that is affected by the tolerance zone width, which can simply be a bias on the firing threshold of a neuron in the comparator. To do this requires no specialized control system, but it is not unreasonable to suggest that some of the membrane-crossing homeostatic loops might be control loops like that in Figure 11.6.

Functionally, the CEVs of such control systems are parameters such as the gains of analogue control units, the inhibitory strengths that affect the internal loop gains of polyflops, as well as the untold myriads of structural parameters that determine the connection pattern within a control hierarchy. These parameters are analogous to the myriads of properties and relationships in the real world that come together to form Objects, simulated by the CEV Objects of the perceptions we control. The logical hierarchy has as its inputs and its CEVs patterns of the perceptual values in the analogue hierarchy. The proposed biochemical hierarchy has as its CEVs patterns among the parameters of both the previous hierarchies.

Figure 11.6 used the e-coli process to adjust the gain rate, the leak rate, and the tolerance parameters of the subject control loop so as to optimize QoC. Now, however, we have a quite different possibility, or rather a quite different set of possibilities, because we have a relatively small set of hormones, neurotransmitters, and other biochemicals whose concentrations in various parts of the body must be kept stable or cycle appropriately, and the mismatch minimized between these few actual values and their optimum values, while the mismatch is used to adjust a very large number of parameter values.

Where have we seen this kind of mismatch before? We saw it in the perceptual control hierarchy itself, when a very small number of slow-moving muscles were available to permit control of a very large number of possible patterns of variation in a complex environment by controlling their expression as internal perceptual variables. How did the perceptual control hierarchy develop to handle this problem? By developing multiple levels of perception to produce perceptions of coherences in the environment, and by deploying the muscles selectively to control those perceptions that most required control at different moments.

According to what measure did one perception have a greater need for control than another at a given moment? Ah, that is what we are talking about now, the varying values of the intrinsic variables that participate in various homeostatic loops.

The process described in the previous three Chapters would be expected to create a very large number of *different* homeostatic loops, though not in the trillions that would be needed to work with the trillions of synapses in the brain. To work with individual synapses, as astrocyte glia cells apparently do, is not strictly necessary. We can revert to thinking about Powers's "neural bundles", but now in a different context. The nerves that cohere so as to provide the effect of a "neural bundle" act as stochastic collective controllers when they participate in a bundles control loop. They behave similarly, and it is no stretch of the imagination to suggest that they might do so in part because they participate (to different degrees, perhaps) in common homeostatic loops.

At this point you might be asking "But what makes the collective decision for the nerves in a bundle to join the same homeostatic loop?". The answer is that it is simply a natural result of the kind of autocatalytic invention process that we have been considering. If the nerve contributes to the loop, those connections survive. If it doesn't, they don't.

Furthermore, perhaps we can now see that exactly the same argument applies to homeostatic loops as it does to the development and survival of novel complex perceptions. The homeostatic loops that work will survive. They work because the actions in controlling those novel perceptions create or maintain the internal and perhaps environmental conditions that allow important homeostatic loops to continue operating. The entire development of the control hierarchy is, from this point of view, simply a combination of the ongoing inventiveness of autocatalytic loops and natural selection, on both individual and evolutionary time scales.

Think about the situation now from another viewpoint. The life of the organism depends on the maintenance of all the homeostatic networks that involve a great diversity of biochemicals that are processed by whole ecologies of microbes, by physically distinct organelles within cells, by discrete entire organs such as the human pancreas, stomach, lungs, bones, muscles, etc., etc. This whole complicated structure is open to influences from outside that might destroy some of the homeostatic processes, and, as with the "Bomb in the Hierarchy" (Section I.6.5), such breakdowns sometimes cascade. How does the structure protect itself against this possibility? By acting on the outer world to reduce the entropic effect of the most damaging of those outer influences.

When we discussed the latent heat of homeostasis, we noted that real physical energy is used to create a homeostatic loop, and some of that energy is stored in the products of that loop, which act as the catalysts in the loop. The fact that the loop is homeostatic means that the rate of decay of the products exactly balances the rate of production, on average, so there is a balance between the heat stored in new product and that released in decaying product. Indeed, this is the way energy is transferred from food to the cells where the product "ATP" is used. If for some reason such as starvation such a homeostatic loop reduces its rate of productivity or ceases to work, that heat balance is broken, and the energy released by the decay of product exceeds that re-stored in the production of new product.

A homeostatic loop might change its rate of production if, for example, the concentration of an inhibitory component were to be enhanced. If the loop included part of a microbiome, for example, an increase in the number of bacteria that contribute to the anti-catalytic function would have that result, which we might see as an infection. Reduced availability of an essential mineral or vitamin might reduce the output of a catalytic component, again reducing the output from all the loop components. Physical obstruction of communicative pathways could have the same effect.

Depending on the degree of disruption of various homeostatic loops and the density of energy storage in the decaying products, other loops may take up the slack in the manner of Figure II.2.9 by increasing their production rate (the organism "acting feverishly" might be an observable result), or the temperature of the body might increase, increasing the loss of energy to the environment. These are possibilities that remain to be explored, suggested rather than demonstrated results of the application of perceptual control theory to the internal workings of the body.

II.8.4 Misperception and Reorganization

So far, we have considered how the side-effects of control influence the intrinsic variables of an organism, and we will continue to do so. Since the later part of this book is concerned with general cultural issues, we must consider how reorganization might be affected by "things that are not what they seem." To bring this within the domain of a single control hierarchy, imagine that a person or an animal is fitted with special spectacles that, unknown to the individual, contain prisms that alter the relation between space as seen and as kinaesthetically perceived. Much work with such prism distortion of vision was done by the Gestaltists (see Wagemans et al., 2012a, 2012b for an extensive discussion) in the 1930s, by Hein, Held and their colleagues in the 1960s, and by J.G.Taylor in the 1960s and 70s.

The main finding of these studies was that humans learned to relink their kinaesthetic and visual perceptions, but to do so required active movement by the subject while vision was unobscured (e.g. Hein and Held, 1963 and references therein). For example, in one study the subject either was pushed around in a wheelchair or was allowed to manipulate the wheelchair without assistance. Only in the latter case did adaptation occur.

Both J.G. Taylor(1963) and Hein and Held showed in different ways that the relinking of visual with kinaesthetic experience occurred not as a unitary transformation that corrected for the distortion imposed by the prism, but affected only one kind of perception or one part of the visual environment at a time. For example, Taylor, who walked with a cane, said that when he learned to walk wearing prisms that both displaced his vision laterally and expanded one side of his visual field relative to the other, thus making the floor seem tilted sideways, he came to see a level path perhaps a metre wide in front of him with the floor tilted on either side of that path[28].

J.G.Taylor also demonstrated that the reorganization of visual perception with distorting prisms was

28. For a few years in the 1960s I worked alongside J.G.Taylor (no relation) after he retired from Capetown. Some of what is described here comes from his face-to-face statements.

multimodal. One dramatic demonstration of this is in a movie[29] of a heroic Seymour Papert (with whom Taylor collaborated) learning to ride a bicycle while wearing fairly heavy prism goggles that inverted the visual world left-to-right. At first, Papert crashed as soon as he started to ride. After a while, he could ride uncertainly, like a child just beginning, and later he rode confidently and naturally.

Then Taylor replaced the inverting prisms with ones that left the visual world unchanged. When Papert first wore them (knowing what Taylor had done), he crashed, and had to relearn to ride again, which he did quite quickly. After some days of wearing on one day inverting prisms an on another day non-inverting, he learned to ride either way, but if he removed the goggles when they had inverting prisms, he crashed. Finally, Papert became able to learn to ride a bike whatever the condition of the goggles, and could switch between states while riding, with no apparent ill effect on his ability. Papert had finally been able to reorganize so that his perceptual world emulated his real world environment well enough to be safe on a bicycle.

J.G.Taylor. explained these results by noting that the goggles had weight and produced feelings of contact with the skin, and these tactile-kinaesthetic sensory perceptions combined with the visual to make a total experience when Papert first wore the goggles with inverting prisms. Only by much active training riding a bicycle with different correlations between the tactile-kinaesthetic perceptions and the visual was he able to create a perception of the visual scene that depended correctly on the visual senses and not on the weight of the goggles, inverting or not.

A similar kind of perception based on the combining of different sensory modalities is revealed by the so-called “McGurk effect”[30] (McGurk and MacDonald, 1976), which shows that the visual movements of the vocal apparatus such as the lips and the auditory sensation of the air-pressure waves contribute together to the perceived identity of the phonemes in spoken words. We might say as an easy overgeneralization that what we consciously perceive is seldom derived from input from one sensor system.

What does all this have to do with misperception and the intentional or unintentional creation of misperceptions? And what does misperception have to do with reorganization? Firstly, if what one perceives is different from what the real world provides, then efforts to control that perception, a perception usually valid for the real world, will have unpredictable influences on the intrinsic variables. In the extreme, controlling a perception that incorporates a misperception can kill. If one drives over a bridge that is not there, bad things happen. If you perceive that John is about to kill his wife, you might perhaps act to restrain (coerce) him so that he cannot. But if the perception was wrong, and John’s wife was about to kill him, that action abets the murderer.

These immediate effects of misperception may be self-evident. What may be less so is the effect of misperception on reorganization. Since the side-effects of acting to correct a perceptual error that does not correspond to a real-world error can disturb other perceptions, not only the misperception but other parts of the control hierarchy will be likely to undergo reorganization. In extreme cases. it is possible that initially correct perceptions may be altered to become incorrect, but in such a way that the complex can be controlled moderately well.

This kind of “false reorganization” is unlikely to happen when all the perceptions in question can be influenced strongly by the person’s actions, but is quite likely at higher levels of the hierarchy, where the

29. About 50 years ago, J.G.Taylor personally showed me this movie and explained what was happening at each stage both in respect of what had been done and the theory behind it. I have not been able to find whether the movie now exists in a publicly available form, but if it can be rediscovered, I hope that it will be.

30. A good description is in Wikipedia <https://en.wikipedia.org/wiki/McGurk_effect>

perceptions have been built from data that has undergone much transformation. For example, one may perceive one property of "Fred" to be that he is trustworthy, because you have been told so. If Fred tells you something, you may perceive it to be truly a part of the current state of the world.

If Fred really is trustworthy, that is all well and good. The world probably is the way Fred says, to the best of his belief. If you control a perception that includes the state of that part of the world, your control will probably be effective in moving your perception nearer to its reference value, and the side-effects of your control will tend to influence your intrinsic variables "properly".

If, however, your perception of the "trustworthy" property of Fred was a misperception and Fred was lying, to control the perception that includes what Fred told you will very probably not have side-effects that would tend to bring your intrinsic variables nearer their genetically determined dynamic reference values. Furthermore, what does happen in your perceptual reality is likely to include surprises. Surprises increase the rattling measure of your perceptual hierarchy, and are ordinarily accompanied by reorganization that tends to reduce that rattling measure. Your reorganization will continue, probably faster than before, but it is unlikely to consistently improve your ability to survive and prosper.

Reorganization, whether using the e-coli principle or not, is inherently unpredictable, since the essence of the process is that the organism lacks information as to how better to live in the actual environment, as opposed to the perceived environment. Nevertheless, as Powers argued, the end result is a tendency toward better control (which implies a lower rattling measure over the entire hierarchy) and "healthier" intrinsic variables. That argument, however, holds only if the controlled perceptions developed through reorganization correspond to consistencies in Real Reality (RR). If they are misperceptions, reorganization, always in RR, will be happening with respect to a world that differs from the perceived world (PR). If that is the case, there may be no way to control effectively, and using the e-coli analogy reorganization may zig-zag the action side of the hierarchy into situations that could be quite damaging.

We will pursue this issue further when we deal with political issues (mainly in Volume IV), since much political technique consists of creating the situation described in Handel's "Messiah" in the line "*And why do the people imagine a vain thing*". In other words, political lies are intended to ensure that the supporting population remain unable to find ways of making their own lives better, but instead reorganize to use support for the lying politician as a means to have it done for them. For now, we just note that misperception can create problems for reorganization far beyond those created by trying to find actions that allow a misperception to match its reference value when the value in the real external environment is quite different.

II.8.5 Reorganization on Many Scales

Let's look at reorganization from a different viewpoint, a viewpoint starting with the "What, Why, How" approach to the simple control loop in Figure 11.7, rather than starting with the control hierarchy that we have treated as the object that is reorganized. When we use this approach to look at reorganization, we find that the control hierarchy of the individual is far from the only structure that can be reorganized. Indeed, the use of the word in PCT is appropriated from its use in organization such as corporations or political parties. The "How" of reorganization may depend on the nuts and bolts of the mechanism being reorganized — to reorganize a company has little to do with concentrations of biochemicals in homeostatic loops — but the "What" (is to be achieved) and the "Why" (it is wanted) seem to be much the same in both cases.

The argument of this Section allows no distinction between reorganization and evolution. Functionally, they are identical. It is, however, convenient to distinguish them on the basis of scale, both of time and of the relative independence and interdependence of the components of Structure. If they are to be distinguished by the use of different words, “evolution” would refer to structure templates that extend beyond the lifetime of an individual or that extend over more than one ancestor-descendant pair. In the next few paragraphs, when we say “reorganization” it will usually be intended to include “evolution”. The context should make clear those occasions when this is not true,

The “Why” answer is the simplest. For reorganization or evolution, there is no “Why” in the sense of a reference value for a controlled perception. There is a “Why” answer, however. and it is the same as the answer given in Section II.3.5 to Nevin’s question “*What’s in it for the cell*”: What’s in it for the individual Structure is “nothing”, except that those structures that help their host organism or society to survive are more likely to be found in the future than are those that don’t. The “structure” in the answer to Nevin’s question is that of the cell, but now it is widened to encompass any structure that endures in a low entropy state while possibly changing in detail over time. It is the answer of evolution “What works, survives”.

The arbiter of what survives is Real Reality (RR). RR can be said to be like a judge of style points in a sport such as ice-dancing or a judge in a TV elimination show. The decision of the RR judge is final. What survives has been “authenticated” by RR thus far in the survival game, but it may not survive the next round of testing. If we imagine the “Why” of reorganization (or evolution) as a reference value in a control loop, the “Why” answer would be “*To satisfy the judge (Real Reality)*”.

The “What” of reorganization is some structure of interactions among dynamic processes. Richardson’s little whirls that feed on the velocity of greater whirls (Chapter II.4 intro) are such a structure. The investigations of Prigogine and his colleagues into self-organization addressed the thermodynamics of dissipating energy and gaining entropy through a cascade of ever smaller structures ending in the unstructured molecular motion we call “heat”. The “What” of reorganization is the description of that structure for a particular case, be it the structure of a society, the structure of a sporting league, the structure of a perceptual hierarchy, or the structure of a cell. In a perceptual control hierarchy it will be the organization of low-level perceptual functions into larger perceptual functions at higher and higher levels.

If we are talking about the structure of a perceptual control hierarchy, RR as we perceive it dictates (in its judge role) that control can be effective only to the degree that the feedback path remains stable within the time-scale relevant to that particular control loop. If the properties of the feedback path change much during the time-scale of the feedback loop, control loses quality the faster the feedback path changes on the time-scale of the effective loop transport lag. For reorganization to “satisfy the judge”, the required stability is in the environment. The many strands of the reorganization feedback loops pass through many different complexities of the environment, some more stable than others, with the real-reality Laws of Nature[31] being apparently the most stable of all.

Figure II.8.4 shows a generic division of the Universe into three components, with labels. At the level of the perceptual control hierarchy, the generic label “Structure” represents the perceptual control hierarchy, “Mechanic” represents the physiological entities and processes in the body, including the microbiome ecologies, and “Environment” represents everything outside the body, including other living control systems. All three are dynamic entities that can influence some aspects of the others.

31. As opposed to the Laws of Nature as they are understood at any given time, from capricious Gods and Goddesses to Newtonian mechanics and optics, to Einsteinian Relativity and “spooky” quantum entanglement, to who knows what in the next decades or millennia.

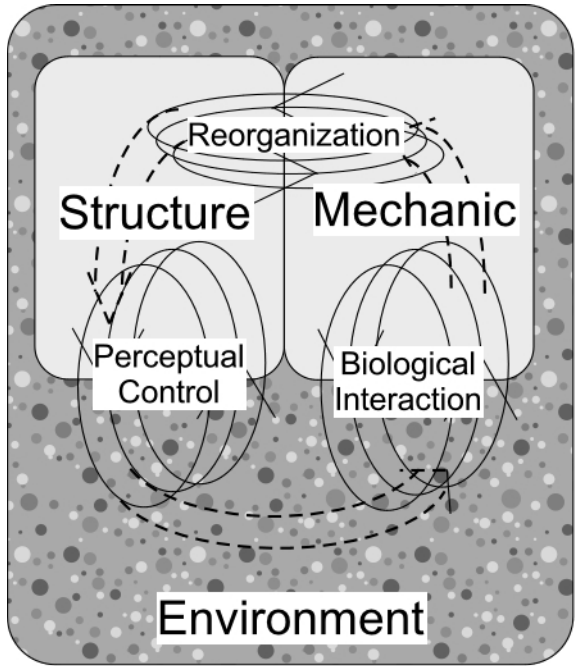

Figure II.8.4 A generic partitioning of the Universe. The three partitions are some structure (here the perceptual control structure in a single individual), a Mechanic that interacts with the Structure (here the physiological entities and processes of the body), and the rest of the Universe. The same partitions apply at a wide range of scales, with different realizations of the Structure and the Mechanic.

In Figure II.8.4, each component of the structure is shown as interacting in a feedback loop with each other component. These feedback loops act in concert, creating two larger loops, one of which is shown as dashed arrows in the Figure, while the other is the same in the reverse direction. The first, again at the level of the perceptual control hierarchy, shows that the effect of perceptual control actions stabilizes the local Environment. Side effects of those actions may affect the physiological entities and processes of the body. Changes in the external influences on physiological homeostatic loops would influence reorganization of the perceptual control hierarchy, completing the overall feedback loop.

As an example, dressing one way for an employment interview might lead to employment, while dressing another way might not. These have different effects on the way the mechanic will reorganize the perceptual control hierarchy, since being or not being employed allow different possibilities for control of perceptions such as having food to eat. The effect is of a braided feedback loop that passes through all three domains. Powers (2005) shows such a loop in more detail (B:CP Figure 2.1, p191).

The larger loop exists also in the reverse direction, because of the bi-directionality of each of the two-way loops of which it is composed. There is no need to follow it in detail here.

It is easy, and correct, to see the larger three-component loop as being built on the three two-way loops in the same way that loops to control complex perceptions are built from loops that control simpler component perceptions in the perceptual control hierarchy. When we follow this kind of structure to different scales, we will see again a self-similar hierarchy of feedback loops, forming modular structures from level to level (Figure II.8.5), just as happens in a perceptual control hierarchy with all its levels of complexity.

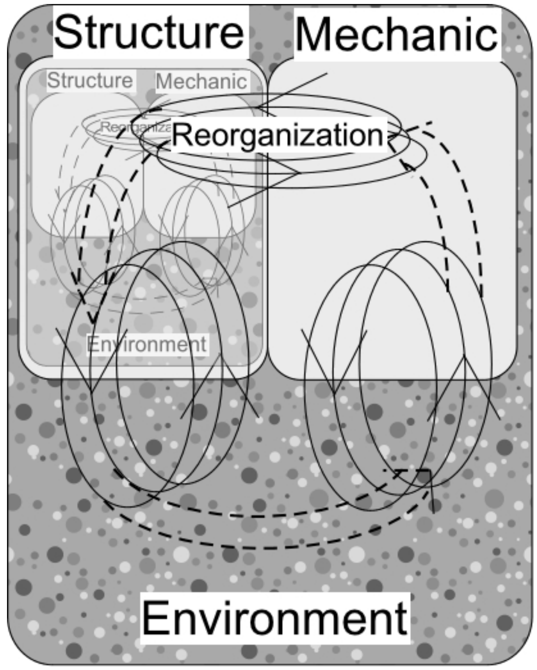

Figure II.8.5 A two-level hierarchy of reorganization scales, functionally identical at each level. Further levels may exist, both larger and more refined. At each level, the structure being reorganized is the way the smaller structure interactions work together, just as in a perceptual control hierarchy, the higher-level perceptions and outputs determine how the functionally identical lower-level control loops work together.

In these Figures, only the Structure changes by reorganization. In the conventional use of the word in PCT, that structure is a perceptual control hierarchy. The Mechanic is the means by which the structure — the connections and parameters of the hierarchy — are altered at all time scales. "*The Environment is Real Reality and evermore shall remain so*" at all scale levels of the diagram, but at different time scales, different aspects of the Environment are most important. In the "Black Box White Box" analogy, what is reorganized is the connectivity and parameter values among the various White Box Objects that have so far been found and constructed by Wiener's engineer to emulate Black Box Objects.

II.8.6 The REV hierarchy

In the perceptual control hierarchy, control of the more complex perceptions depends on control of the simpler perceptions of which the complex perception is built. So it is with the hierarchy of reorganization or evolution suggested in Figure II.8.5. The Mechanic that affects a social structure is constructed from the individual perceptual control operations of the people, animals, trees, or perhaps even bacteria that participate in the structure.

One must be careful here, because social structures can be imposed from above, such as by an imperial power enacting a law that affects only a colony. This does not contradict the concept that an evolving structure undergoing reorganization does so because of the actions of the members. Rather, such top-down construction and maintenance is analogous to the construction by humans of a so-called "autonomous" robot. The structure of the robot has nothing to do with the ability of the robot and its future copies to survive, unless the designer had that trait as a design objective.

We are not talking here about such arbitrary constructions. Here we are dealing with structures for which the judge of survival is the ever-changing Real Reality based on the never-changing structure of Natural Law. In the next two Chapters, however, we suggest that perhaps one role of conscious thought is to act as such a "designer" of the control hierarchy within an individual organism. In the hierarchy of "Structure-Environment-Mechanic" complexes, this possibility implies a possible parallel and interacting hierarchy, much as we argue that the perceptual control hierarchy may have interacting analogue and digital components and that biochemical homeostatic loops may form part of the Mechanic's toolbox.

In what follows, we will refer from time to time to a "REV" ("Reorganization-EVolution", or "REVision") hierarchy, often in connection with social structures rather than within the individual, although the same construct applies equally at the smaller scales. In this hierarchy, an Elementary Social Unit (ESU) consisting of one level of Structure-Environment-Mechanic takes on a role analogous to that of the Elementary Control Unit (ECU) complex of Perceptual Function-Comparator-Output Function in a perceptual control hierarchy.

There is, however, a significant difference between an ECU and an ESU, in that the ECU output sends signals to the Environment through lower-levels of ECUs, whereas the ESU is a complete loop through the Environment, living and non-living. It is not just the family, the sports club, the revolutionary clique, or a friendship group, but includes both the interactions among the individuals, which occur through the environment, and their collective control effects through which the group, as a group, influences the rest of the world. At higher levels, the lower-level groups are the "individuals", just as in a perceptual control hierarchy within a single body, lower-level perceptual signals form the inputs to the next-level perceptual function inputs.

The "signals" between the levels in a REV hierarchy are survival and continuation of the existence of the Structure, and these signals transcend the individual. Higher level structure affects the survival probabilities of lower level structures, while lower level structures optimized by their local "Mechanics" permit the continued existence of the higher levels that depend on them. Particular structural forms, such as bilateral form, or controlling head connected to sensing and acting body, that recur and survive long enough not to be considered transient forms become Motifs that exhibit emergent properties.

In a REV hierarchy, there is nothing to convey survival "signals" analogous to the "neural bundles" that carry the signals of the Perceptual Control hierarchy. These "signals" are abstract probabilities. Despite the lack of an obvious carrier for the "Neural Current analogue signals of survival probability, nevertheless the Structures at the different levels of a REV hierarchy are inter-related in much the same way that the perceptual signals are related in a perceptual control hierarchy.

Whether in a living being or not, a higher-level structure is composed of lower level structures that relate to each other in some way. In a society, the survival of a higher-level ESU depends largely on the survival of enough of the lower-level ESUs upon which its structure component depends. A sports club, for example, has a structure that consists of some administration, some code of conduct, some teams that play the sport, and so forth. The structure of administration may include committees, a system of fee collection, and so forth. The code of conduct may include components specific to the club, such as requiring members to wear the club colours when in the building, but that will also use structures of conduct that are used in other contexts, just as are the lower-level perceptions in a control hierarchy. The teams run by a sports club contain players who are organized differently depending on the sport.

If there are not enough players to form at least one team, the whole hierarchy that is the club may dissipate. If there is no administration, it may go bankrupt or the building may collapse. If there is no written or unwritten code of conduct, the members may argue in ways that disrupt the club and lead to its demise. The higher level structure of "team" depends on individual players. The higher level structure of "club" depends on "teams", and on "administration" and "code of conduct", each of which is structure in the ESU that defines a level.

The survival of an ESU depends on the “Mechanic”. In the sports club, at the “club” level, the Mechanic is an agreed set of rules, a “constitution” written or unwritten, by which changes can be made, such as to change the sports being played, to allow the club to interact with other clubs in one or more leagues, and so forth. Such changes are likely to influence the survival of the club as a distinct structure, for better or for worse. The functioning of that “club-level” Mechanic is implemented by the members or by outside influences accepted by the members, analogous to the homeostatic loops that may implement the Mechanic in a perceptual control hierarchy.

Although there are many analogies between a REV hierarchy of ESUs and a perceptual control hierarchy of ECUs, the REV hierarchy is not a control hierarchy. It just looks the same as one in many ways, while differing in the crucial way mentioned above: The ECU is only one part of a control loop, the rest being its environmental feedback function, whereas the ESU is a complete loop. Nevertheless, the modular structure of a REV hierarchy is likely to be similar to that of a control hierarchy.

Within an individual or a cultural group, modules are likely to group into ever larger-scale modules of modules, cohering finally into what is often called a “personality”of the individual or group. To some extent this nesting of modules into a personality corresponds to levels of perceptual control, but it also may include module nesting within a level. The “personality” of a group seldom is true of all the members of the group, any more than all Scandinavians are blond. Are all Italians great lovers? Does every “risk-taker” enjoy scaling cliffs? And what about people who display multiple personalities?

Always we should keep in mind that surprise and rapid change ordinarily increase the rattling measure over all the error variables in the controlled perceptions of all the control hierarchies in the social system, and that all organization is apt to change in the direction of lower rather than higher rattling. This applies just as much to the ecologies of living things studied by scientists called “ecologists” as it does to the ecologies of people who interact with each other formally or informally, and to the ecologies of perceptual control we have presumed to settle into modular hierarchies of interacting control loops. In Chapter xII.10 we will examine how conscious thought affects reorganization in humans and probably a few other species.

II.8.7 The Mechanic’s Toolbox

The reorganization process can be likened to the work of a human mechanic or developer maintaining and upgrading a complex piece of machinery. A human professional mechanic’s toolbox usually contains more tools, more diverse kinds of tools, and more precise tools than the toolboxes of most handymen-about-the house. We might expect evolution to have developed an analogous variety of specialized tools to reorganize structures in different situations.

In Chapter I.9, we talked about Hebbian-antiHebbian (HaH) learning through modification of synapses according to the relative timings of firings in upstream neurons and the downstream neuron. We have discussed Powers’s “e-coli” method of reorganization in simulations, and why the HaH process might implement the “continuing direction” aspect of the e-coli approximation to optimality (Section I.9.5). The “e-coli” approximation technique is, however, an analyst’s view of the result of the Mechanic’s work, not necessarily one of the tools in the toolbox, and the HaH process offers no convincing explanation of the “tumbles”.

A different kind of possible tool for reorganizing the perceptual control hierarchy was suggested in Figure I.11.6. This tool is itself a control loop, which has the job of optimizing the QoC of a single perceptual control loop by acting on the local parameters of the ECU. Above, for example, it was suggested that the perceptual control hierarchy might be reorganized by homeostatic loops that penetrate a porous membrane between the perceptual control structure of signal impulses along neurons and both local and widespread concentrations of different biochemicals such as hormones and enzymes. That would be another kind of tool, although a perceptual control loop is, as we saw in Section II.3.1, just a specific case of a homeostatic loop.

One problem is that optimizing control processes by reorganization must eventually deal with trillions of synapses, and the only way to do this within a time span shorter than the life of an organism is to affect relatively small groups of them together in a hierarchy of structures (in his toy Universe, Kauffman(1995) found groups of about 5 or 6 units were optimum), the bottom level of which interfaces with the perceptual control hierarchy through the external environment. This is precisely Wiener's approach to building White Boxes that emulate the functioning of a complex Black Box. By analogy with Kauffman, we may guess that few of these white boxes start with more than five or six input terminals, though they might grow more as time goes by and the operations of the initial ones become more practiced, and therefore instances of them become more stably entrenched as units within many White Box instances of appropriate types.

If this structure, or any similar structure of homeostatic loops that include the perceptual control hierarchy in their loop structure, is really represented in living things that reorganize, it must be subject to many of the effects that we have noted in the perceptual control hierarchy. In particular, it will be liable to experience avalanches caused by the "Bomb in the Hierarchy" phenomenon (Section I.6.5). While the avalanche progresses, rattling increases, and the Mechanic will tend to use available tools to reorganize the structure relatively rapidly.

In the perceptual control hierarchy, the Bomb avalanche is the result of switching some environmental feedback loop from negative to positive feedback. The initial Bomb may set off an avalanche of such explosions up and down through the levels of the hierarchy. Such an avalanche might be of any size from all-encompassing, blowing up the entire perceptual control hierarchy, to trivial, affecting only the control loop initially infected. Such an explosive avalanche of positive feedback propagation could plausibly implement the tumble feature of e-coli reorganization.

Why might such a Bomb blow up in a reorganization structure that encompasses neural and biochemical (and probably micro-biotic) processes? The Bomb is distinguished by feedback loops that are normally negative feedback loops, but that at some point switch to positive feedback. Something that normally helps, now hurts. That is precisely what would be the case if the "straight-line" path of the e-coli passed by the point where continuing ceases to improve the control performance of the perceptual control hierarchy, and instead begins to make it worse. The equivalent of the Bomb would be a scrambling of the directional parameters of the e-coli motion — a tumble.

Among the tools in the Mechanic's Toolbox is one mentioned by Powers and used explicitly in psychotherapeutic "Method of Levels" (e.g. Mansell, Carey & Tai, 2012). That tool is consciousness, the main topic of the next two Chapters.

II.8.8 The Magic Number Seven?

A famous paper by George Miller (1956) claims that working memory holds no more than "the magic number 7 plus or minus two" items. Over the intervening two-thirds of a century, the principle has remained viable though the detail may have been questioned. We should expect, then, that when a "Solution is here" in the form of available developed perceptual controllers (as in Figure II.10.11) is proposed, it would seldom if ever contain more than six or seven parallel steps requiring different skills, including the perceptual categories used in perceiving the form of the problem as well as the action categories required for its solution.

Since the categories developed from passive observation and from non-conscious experimental control actions do not invoke any memory other than that embodied in the strengths of individual excitatory and inhibitory synapses, there is no intrinsic limit to the number of inputs that may contribute to a perceptual category. All the same, there are hints in Kauffman's work (Kauffman, 1995) that a similar limit might well apply to interacting patterns of sensory inputs that form categories, or to side-effect interactions within and among modules of the perceptual control hierarchy. We will consider this possibility of a limit on the non-conscious development of perceptual categories later. For now, we are concerned with the development of delimited perceptual categories in conscious rational thought processes.

A limit to a single-digit number of useful perceptual categories in a single conscious thought may seem quite contrary to the enormous variety of things we can imagine doing or observing in a single scene, but observing is passive, and we do not usually imagine doing even three things at once, let alone six or seven. An octopus might, but with our numerically limited arms and legs, we may manage two, and in appropriate situations a third, which often consists of holding some state stable while varying the other two. Might the same kind of limit apply to the number of independent entities and relationships involved in a narrative fragment, which consists of one event that links changes in the values of individual instances of one or more perceptual categories? We will consider this possibility below.

We return to Stuart Kauffman's "At Home in the Universe" (Kauffman, 1995), which we used when describing autocatalytic networks. Now we look at a simulation experiment on a toy universe that Kauffman used as a view onto biological possibilities for solving what mathematicians call NP-hard problems. These are problems for which a unique best solution can be found only by searching the entire space of possible answers. Powers regarded reorganization of the perceptual control hierarchy as such a problem, though he did not look at it as a mathematician might have done. Rather, he treated it as an engineer would have, and as we, too, are doing in our various ways, as a problem in finding very good solutions and not doing a prolonged search for the very best. As the saying goes: "*The perfect is the enemy of the excellent*".

A "fitness landscape" based on all the interactions an individual of a species might have with individuals of its own or other species presents this kind of NP-hard problem. There is, presumably, a stable best way for each to improve its own welfare by interacting in this or that way with all the other individuals of all species that it encounters, and in Chapter II.5 above we suggested a "simulated annealing" possibility that Chvykov. et al. (2021) called "rattling" (Section II.5.1) for optimizing organizations of such interactions. Reducing the rattling measure improves the stability of an organization, much in the way that cooling a heterogeneous liquid is likely to increase its viscosity. But too much viscosity produces structural elements like glass, which, though technically liquid, move so slowly that over time-spans of generations they can be treated as though they are solids for perceptual control purposes.

Different arrangements of side-effect interactions among the elements of a social structure will act like the binary units of the toy universe studied by Kauffman (1995). Kauffman was looking at how the range of interaction among units affected the overall fitness of the whole system in fitness landscapes of different ruggedness. We define fitness as the evolutionary ability of a species to survive over time as circumstances, including the types of other species with which they might interact, change. Kauffman did not treat those interactions, but instead used a summary called a "fitness" landscape, in which a species survives better the higher its fitness — Darwin's "survival of the fittest" seen in a different light.

Fitness landscapes may be smooth or rugged, and the locations of their peaks in the appropriate descriptive space are expected to move as species and their interactions evolve. The fitness landscape for any one species is not the same as the fitness landscape for another. A species that can evolve to follow the movements of its fitness peaks is one that is likely to survive, though never in its precise original form. It will evolve new abilities and eliminate dangerous atenfels for controlling its perceptual variables, just as human drivers would tend to avoid roads that become bumpy and ill-maintained if there is a smooth highway built nearby.

Fitness landscapes are not simple. What Kauffman did was to strip all the complexities out of a system of interacting entities in a changing fitness landscape, by restricting the entities to "units" having only two states that he called "1" and "0", which would change depending on the combined state changes of all the units in a square matrix patch of N×N entities. In his simulation, the entire universe consisted of a patch of dimension 120×120 entities in each direction, a universe of 14400-bit uncertainty. The experiment was to determine the size of non-overlapping patches that when individually optimized would yield the best optimization of the fitness measure, low entropy overall, over the entire universe, if such an optimum existed. Under a variety of conditions that size was near 6×6. Larger, and the system as a whole became chaotic, and smaller it usually froze into a sort of optimum worse than the optimum a non-overlapping set of 6×6 patches typically achieved. Patches of 6×6, seemed to be on the critical edge of chaos.

One cannot take any numerical result from such a toy universe as applicable in detail to a real-world situation, but the character of many of the results is likely to apply more generally. If there are too few interactions among entities (side-effect interactions if the entities are living control systems), the whole set will be too stable to follow changing peaks in a varying fitness landscape. If there are too many, the whole set will change chaotically, not settling near any particular fitness peaks.

Why did Kauffman get the results he did, which can be characterized by investing his interacting binary units with properties like control units whose side-effects mutually disturb each other. A result that might be generalized was that if there are only a few mutual side-effects, things get locked up into too stable homeostatic loops, but if the side-effect interactions are too many, the system's organization might easily change chaotically. Such a system would not survive long. As usual, evolution seems to find an "edge of chaos" solution that allows rapid change but remains in control.

Let us consider why evolution tends to approach "Edge of Chaos" conditions, evolving from both sides — over-stability and chaos. Firstly, a state of chaos would involve a lot of rattling, rapid unexpected changes in the perceptual errors for controlled perceptions, whatever those perceptions might be. Rattling always increases the probability of organizational change, which always tends toward (but does not guarantee) a state of lower rattling. The more the rattling the greater the likelihood of organizational change. Side-effects disturb all affected controlled perceptions, reducing the quality of control, and the lower the quality of control, the more varied and unexpected will be the disturbances to other controlled perceptions. Hence, the tendency toward organizations with lower inherent rattling will trend toward a reduction in the number of strong side effects that influence any controlled perception.

This analysis suggests that evolution should trend toward isolation, since that eliminates all side effects. but let us not forget that the side-effects are the results of controlling by acting on the environment, and those actions will disturb perceptions that also must control by acting on the environment. It is impossible to eliminate side effects by isolation, if the environment is shared, which would make isolation of a controller that influences and perceives the real reality environment a self-contradiction.

Six may not be a safe number to extrapolate from Kauffman's toy experiment, but somewhere in that neighbourhood there is surely a limit to the effective number of side effects that can be tolerated if the system is to be both stable and flexible in changing conditions. "Effective number" means that it takes several weak side-effects to carry the influence of one strong side-effect, a difference not possible in Kauffman's toy universe. "Stable and flexible" is a description that might describe a control loop, as well, as well of any tensegrity structure. Too stable or too flexible will not do the job well in either case.

In Chapter I.8 we argued that a control hierarchy should be a tensegrity structure, and that a two-level hierarchy with three interacting upper-level control units would form a minimal control tensegrity structure. A minimal tensegrity structure is fragile, in the sense that if any of its components fail, the rest "collapses", each component dealing with stresses by itself. Redundancy is needed for long-term survival of the structure, because there will surely be occasions when a necessary atenfel is unavailable to at least one of the control loops involved in the structure.

We must ask how much redundancy is optimum, because there is a trade-off between involving more control units in the structure to improve redundancy and the interplay of side-effects that reduces the control quality of the affected control loops. The more mutual side-effects from other controllers there are, the worse each controls, which in the limits renders a hugely redundant control tensegrity structure ineffective in controlling any of the perceptions involved. All of its components become highly rattled, just from side effect interactions among the units of the structure. Evolutionary reorganization will ted to render such structures unstable.

If three control units at each of two levels can produce a minimal tensegrity structure, how many are needed for adequate redundancy, redundancy that allows both for bypassing a failing controller in the spreading of stresses throughout the system, such as a controller that fails because a key atenfel becomes unavailable, and for allowing non-destructive reorganization if the structure considered as a whole organization becomes strongly rattled.

The main result that I extrapolate from Kauffman's toy universe experiment is that local "patchy" optimizations of a fitness measure in an interactive dynamical situation can have results better than global optimization, but only if the local patches are neither too small nor too big. "Better" includes both stability and the flexibility to follow important peaks in different dimensions of a changing fitness landscape. In Kauffman's case, all the local patches were the same size, but one might anticipate that if the experiment were re-run using patch sizes distributed as facets are in a "crumpling" study such as those in Chapter II.6, a similar result would be found. In what follows, I will use "6" as though it were an exact number for all cases, but the reader should be aware that George Miller's "plus or minus two" always applies, and that under different circumstances, "two" might be just as indefinite.

I think the same arguments would apply to a perceptual control hierarchy within an individual, in which a "patch" is the convergence range of input lower-level perceptions into a perception at the next higher level, or the reference level fan-out to lower levels. This fan-out does not necessarily represent choices of action, but is more likely to represent concurrent or sequential actions that would produce the desired effect on the controlled perception.

These actions are likely not only to act together in support of their common higher-level controlled perception, but also to have side-effects on each other. Accordingly, we would expect Kauffman's "patch" experiment to bear on the mutual side-effects among the actions in control of lower-level supporting controlled perceptions, and to suggest that there may be an optimum fan-out that allows both for a reorganized and rather stable perceptual control hierarchy and adaptability in the face of changing conditions.

Humans are reputed to be the most adaptable of all living species, having evolved in different ways to live in all parts of the globe other than Antarctica.

We cannot rely on Kauffman's patches as precise analogues of biological processes, any more than we can take the precise perceptual control hierarchy described by Powers as literally true of biological processes. As early as Chapter I.9 we demonstrated reasons why lateral inhibition had to be incorporated in the perceptual control hierarchy, and some consequences of that requirement. We should expect similar effects from relaxing the mathematical precision of Kauffman's toy universe of interacting switch flipping. Nevertheless, just as with the Powers perceptual control hierarchy, Kauffman's findings can be taken as being very probably a useful guide for further exploration without taking the details as gospel truth.

It is, for example, by no means assured that the same numerical result would be optimum for the kind of perceptual category development suggested by Figure II.10.11, but the same principle probably would be. The value six or thereabouts is likely to be a per-dimension number, and since each category represents a single dimension no matter how many inputs the corresponding perceptual function might have, the conscious thinker of Figure II.10.11 is thinking of only one thing at a time, about what procedure using available control units would be likely to achieve the reference value for the newly designed perception — the consciously recognized Problem. The implication is that the newly defined category and its atenfels probably number around six.

Figure II.10.11 shows the development of a newly controllable category at the top of the non-conscious perceptual control hierarchy, and our first amendment to Kauffman's toy universe is to treat each dimension of his 6×6 patch as a unit with a rule that some patterns of ones and zeros report a one, while others report a zero. Each dimension then acts as a single switch at a higher level. We can play Kauffman's simulation using these composite switches as elements in a higher-level N×N universe, and may be justified in expecting approximately the same 6 items per dimension optimum value for patch size.

The number six is a result Kauffman found when examining the interactions among "square" modules of various sizes, the same number of units in each dimension. We will be using it in the rest of the book, even though in reality, the number of discriminations along a a dimension are likely to have some distribution that is wider than a single peak at six. Garner (reported in Garner, 1962) found that the uncertainty of discrimination on different audio dimensions was about 2.3 bits, representing about five distinctly discriminable values. Kauffman, however, was not looking for clear distinctions, but for a number larger than led to stability by lower than led to chaotic behaviour of the units in his modules. Five precisely distinct units in his toy universe led to stasis in the long run, which one might think to be analogous to Garner's clearly distinct discriminable values.

In Garner's terms, six would lead to a little discriminative uncertainty, while seven would lead to a lot. This parallelism between the toy universe of Kauffman and the human results of Garner encourages me to treat Kauffman's edge-of-chaos value six as a value that is reasonable to treat as a "Magic Number" when we deal later with organizational restructuring that tends to minimize rattling effects between modules in a social structure — or within the perceptual hierarchy of an individual.

We have treated any category that has been reorganized into the perceptual control hierarchy as an element of the soup that is the source of perceptible events — changes in the value of an instance of at least one category — and hence of narrative fragments available to conscious thinking. Now we have argued that rational thought can create a new category to join the hierarchy if the same problem arises sufficiently often. The catalogue of category perceptions is likely to grow rapidly in the beginning, each "same" event in the conscious narrative world enhancing and weakening much the same set of synapses in the same set of neurons, but more slowly with each recurrence of an instance of the category. The trend is like a logarithmic growth function, which from the viewpoint of an observer might well be seen as the person's growing skill at a task, until the task is completed with little or no consciously experienced thought.

All the above refers category catalogue growth due to passive observation, but as we argued back in Section I.11.1, the effects of active control are more important in the healthy survival of the organism — the maintenance of the complex structures of interacting physiological "intrinsic variables" that have their own stability criteria. Accordingly, we expect reorganization consequent on active control to be much faster, but eventually less precise, than the logarithmic growth we expect to depend on passive observation. Since we now should expect around six major inputs to any well-reorganized category function, we can use this limit when discussing reorganization as in Chapter II.8, where no such limit was assumed.

In particular, a phrase or other unitary element of communicative language in any form should be more difficult to understand as intended if it has more than six different active components identified as such in conventional manners of its use — its Grammar.

Chapter II.9. Consciousness and Mechanism

Consciousness is something we all experience, but about which philosophers have argued for centuries, if not millennia. What *is* it, beyond being what we, in a circular definition, consciously experience? Consciousness encompasses everything we (consciously) experience. If we could know what it is, we would have another circular definition, of knowledge. Circular definitions seldom lead to understanding, though they are often useful, since they can serve to energize long-lasting academic arguments, to the financial benefit and personal fame of all engaged parties.

Let us avoid entering the morass into which such arguments lead, and ask instead what functions consciousness performs for the conscious individual. We will propose several different possibilities, beginning with the construction of novel perceptual types that might become reorganized into the non-conscious perceptual control hierarchy. In this respect, it may be worth noting that the entire set of levels in Powers's hierarchy were the result of just this kind of process, building and testing different possibilities in his conscious imagination.

As an example, the hierarchy itself was a new perception that had properties, such as "levels", that permitted the execution of functions testable by experiment. While the perception that is an essential component of a perceptual control hierarchy remains conscious in the minds of most of the followers of Powers, it might well have become part of his own non-conscious hierarchy. I will not follow this illustrative example any further in our exploration of some possible functions of consciousness. Instead, I will suggest how Beach and Wise's (2022) "Theory of Narrative Thought" (Chapter qII.10) and Friston's "Predictive Coding" approach to control may each separately and severally be combined with Powers's Perceptual Control Theory into a single hybrid hierarchic control system that seems to account for much that has been mysterious about consciousness and sleep in living control systems.

Perhaps not coincidentally this hybrid control system has a structure that was used by a powerful artificial intelligence system, AlphaFold2, for deducing the folding shapes of proteins (a difficult problem for previous AI attempts) (Miller, 2021). According to Wikipedia "AlphaFold2's results at CASP were described as 'astounding' and 'transformational' " (CASP is the bi-annual "Critical Assessment of Techniques for Protein Structure Prediction").

Clearly not coincidentally, the perceptual side of this hybrid structure formed the "Bilateral Cooperative" theory of reading (BLC), by the present author (Taylor and Taylor, 1983). The hybrid control model described in this chapter may be considered as an action-side complement of the BLC model as well as an elaboration of the Powers version of PCT.

II.9.1 Consciousness: The Mechanic as "Director"

As has been demonstrated by the success of the psychotherapy known as Method of Levels (MoL), consciousness is an important tool or tool set in the Mechanic's toolbox. In MoL, the patient is guided to "think up a level" to identify a conflict that is always assumed to exist (Carey, 2008) and be the primary reason for the patient's symptoms.

Such a conflict ordinarily arises from wanting to achieve two incompatible ends, both of which may support control of some higher-level perception. This "thinking up a level" is the patient's conscious examination in imagination, thinking about thoughts. When the therapy is successful, as often it very quickly is, to all appearances the patient's perceptual control hierarchy has been reorganized in the area of the conflict, so that the higher-level perception, maybe more than on level higher, may now be adequately controlled.

Consciousness appears also in the teaching and learning of a new skill. To master a skill is to be able to perform it fast and accurately, apparently unconsciously and effortlessly. Many highly skilled performers in various domains do not know how they do what they do. They "will" it, and it happens. The processes involved have become part of the reorganized, non-conscious, perceptual control hierarchy. But they did not start out already not-seeing the techniques they use in their skilled performances. Those techniques were consciously learned or invented, and only after much practice did they become non-conscious. In various TV items, I have seen an interviewer ask a skilled performer, say a solver of Rubik's cube in a couple of second or of very difficult Sudoku puzzles by just writing in the numbers, how they did it. The answer is typically along the lines of: "*I don't know. I just see it.*"

I consciously choose to do something, such as type this paragraph, I can often make my choice happen[32]. I can, for example, swing a golf club and quite often make a little ball go somewhere near where I want it to go, but I would not normally choose to do so indoors or outside in the snow. Even so, I could play and actually have played golf in fairly deep snow to keep up a family New Year's Day tradition, but that, too is a choice consciously made.

The interesting question is the role of consciousness, not in the making of of such choices, but in the acquisition or learning of the non-conscious skill as part of the reorganized control hierarchy, perhaps under the tutelage of a professional. The teacher says things such as "Feel this muscle tense, but not that one", "Hold the club with the hands like this", "Feel your midsection rotate toward where you want the ball to go", "Feel when in the swing your weight shifts to the front foot ", and such-like.

In many skills there are too many of these "feel" perceptions to be controlled consciously all at once, but often by controlling two or three of them at a time, the skill improves over time. One might idly wonder whether this "two or three" is a particular manifestation of the combinatorial issues that led to Kauffman's (1995) "five or six" optimum, and whether perhaps the "two or three" actually link to two or three other "feels" that have already begun to stabilize to make "five or six" in all.

Ignoring this speculation, the incorporation of new controlled perceptions into the hierarchy is reorganization. With practice, the new skill slowly fades from consciousness and becomes an ordinary part of the non-conscious perceptual control hierarchy instead of a novel set of actions performed to

32. This ability is the "Prediction" aspect of "Predictive Coding".

control some previously uncontrollable (and perhaps non-existent) perception.

Assuming the basic distinction between the reorganized non-conscious scalar perceptual control hierarchy and the patterned solution-seeking conscious control processes has some validity, we may look a little deeper into the relationship between the two processes. How does conscious control become non-conscious as skills develop, and could purely conscious control that involves externally observable action ever happen? What is the essential difference between the scalar controlled perceptions of the Powers hierarchy and the pattern perceptions we consciously learn to control?

The essential difference is incorporated into the question. The conscious perception is of an environment, real or imagined, that contains many objects, whereas a non-conscious perception in the hierarchy is a scalar value of some property of an object, the result of applying a perceptual function to possibly many input variables, themselves perceptions. These controlled perceptions are the values of properties of consciously perceptible Objects. The conscious perceptual processes, therefore, can emulate a perceptual function in one respect, by making use of a wide range of already reorganized less complex perceptions.

To repeat, using slightly different words, what we consciously perceive is to a considerable extent a world full of objects, not a world full of properties. Each object, each teacup, each hammer, each computer and each abstract object such as a political party, has many properties. In contrast, any perception controlled in the reorganized hierarchy is a perception of just one of those properties of one object, not a perception of an entire object of which we are consciously aware. When we control consciously, we control some function of several of those property-perceptions belonging to different objects, typically because we see how using this property of that object can help with control of some property of a different object.

For example, we could use the internal strength and hardness of a few moderately sized stones to hold up a flat slab at a certain height, like legs under a table. Since we don't live in the Stone Age, we may have never done such a thing, but right now we want to use such a slab at such a height in the way we have often used a table. We imagine how the stones fit together securely when the slab is placed on top.

Maybe it works, and maybe it doesn't. but if it does, that solution may be remembered in a way that it can be recalled to imagination and re-implemented should a similar problem situation recur. In the case of the golf swing, similar situations recur frequently, and the reorganization process goes beyond problem solution to the optimization of the control and interconnection parameters of the ideal swing.

Powers did not make many conjectures about consciousness, beyond asserting his belief that every conscious perception was built from perceptions already reorganized into the control hierarchy. We, however, are proposing that conscious control of a pattern of existing perceptions as a single unit is a way of creating new conscious perceptions that will either be lost over time or with practice will be built into the hierarchy.

Conscious control also occurs when we want to control an already created single scalar perception using means that we have not previously applied to control of that perception. Maybe the known means is blocked, as, for example, if we ordinarily control for getting to work by one route, but find our road blocked by construction or an accident not yet cleared. We must imagine how to make an effective detour using roads we never used for this purpose. That we do consciously, but if the same part of the same road is frequently blocked, we take the detour "without a second thought", almost non-consciously.

These are by no means all the situations in which we control some perceptions consciously, but since the situations are largely described intuitively rather than as the result of experiment, we go no further here other than to suggest that the phrase "problem-solving" applies to everything we have described.

What does it mean to control even one perception consciously? It means that some sensory data have percolated through various levels of perceptual functions into your consciousness, that you must have an imagined (conscious) reference for what the conscious perception should be, and that your consciousness can force at least one reference value into some appropriate point in the already reorganized control hierarchy, so as to cause actions that would change the perception toward where you would like it to be. That this is possible is one more hypothesis that must be asserted for the general theory of Perceptual Control.

From the perceptions available to conscious experience an enormous number of different complexes could be combined into a perception to be controlled, so why should one be conscious of any particular subset of them other than those that are properties of the same Object? Here we must imagine the reference types or values that apply to perceptions to be controlled.

These reference types and values are conditions in Perceived Reality or (as in the case of MoL) inside oneself that one would like to be true. They probably are not currently true, but they could be. They are perceptions in imagination that do not correspond to any perception yet reorganized into the control hierarchy, because if they were in the control hierarchy and the error was beyond tolerance bounds, they would be being influenced by one's ongoing control behaviour, out of consciousness's purview.

This is not to say that one cannot be conscious of perceptions built within the reorganized control hierarchy. One obviously can, if the conscious environment is built of them, as Powers supposed. It is to say, however, that one does not usually consciously control perceptions that are successfully controlled within the hierarchy. When one is playing a round of golf, for example, one is not usually conscious of feeling "your midsection rotate toward where you want the ball to go". Unless one starts the swing conscious of one's midsection, by the time one is conscious of it, if ever, the swing has been completed and the ball has gone somewhere that you may or may not have wanted it to go.

The references you consciously inject into the control hierarchy might perhaps be in competition with reference values derived from higher-levels in the control hierarchy, especially if you are relearning to swing the golf club effectively after having taught yourself a bad swing. If that happens, the early result is likely not to lead directly to a better swing, but to produce a swing that is worse than your self-taught one before you achieve a serviceable swing that can be further reorganized into a good swing.

Each time some pattern of neural firings occurs, some synapses are strengthened, while others are weakened in the HaH process (Chapter 9). If by conscious thought you identify a pattern you perceive as differing from your reference for what you want that pattern to be, and then inject reference values into the perceptual control hierarchy, the synapses involved in the signals that create that perceptual pattern belong to the "nerves that fire together wire together" class supposed to typify Hebbian learning. This applies to the nerves that produce the reference pattern, the nerves that produce the "conscious comparator" output, and the nerve firings that constitute the injection of values into the already reorganized control hierarchy.

In other words, the nerves that constitute the entire structure of what we might call a "Conscious Elementary Control Unit" (CECU) begin, probabilistically, to form a coherent entity, while anti-Hebbian processes related to nerve firings shortly after the CECU firings on any contributing neuron would tend to be reduced in strength. If a similar coherent pattern recurs, those same synapses would be further strengthened, as would others that happened to have the appropriate timing relationships between their upstream and downstream neurons.

Over recurrences of the pattern — conscious attempts to "*Feel your midsection rotate toward where you want the ball to go*" for example — a Powers type of "neural bundle" will begin to grow, while neurons that fire adventitiously at different times will have the relevant synapses sometimes strengthened and sometimes weakened. The combination of Hebbian and anti-Hebbian processes over time and many repetitions would in this way produce a bundle with a strong core of neurons connected so as to produce the most common pattern. What you are building is a process that forms or modifies part of an intangible Object we can call a "golf swing".

The most common pattern is not necessarily the optimum pattern. The optimum pattern is determined not by the repeatability of the well-learned structure of the CECU, but by where the golf ball goes when you swing at it. That is the Mechanic's feedback. It is also the environmental feedback in a higher-level CECU, that controls for a dynamical perception of the ball flying or rolling true. If this perception is in error by more than its tolerance range, for example if the ball flew into the rough instead of down the fairway, this higher-level CECU will change its output value, and hence the reference values for its contributing component perceptions, whether they are consciously controlled or not.

In the absence of another person as teacher, the reorganization principle would be something like e-coli. Whatever it was, it involved some kind of optimization in a high-dimensional space that had reach or approached a local optimum. Changes in any of the reference values for any of the variables would lead to a result that departed from the optimum, as we noted above. Most changes in the swing would be for the worse, including changes made by injecting new consciously derived reference values into a hierarchy already stabilized by self- teaching.

The role of the teacher, in this case, is to provide a new conscious reference value for a perception that produces a result in conflict with the student's control of the perceived flight of the ball. Other factors, conscious or non-conscious, determine whether the student resolves this conflict by reducing the gain of the "ball flight" conscious control loop or by ignoring the teacher and stopping the lessons.

The "teacher" in this example can be seen as a component of the Mechanic that we might more generally call the "Director". Depending on the scale at which one is thinking (consciously) about it, the Director might be a Mechanic's tool or a component of the Mechanic alongside other components such as the "Stabilizer" (homeostatic loops stabilize the perceptual control hierarchy) and the "Accountant" who checks the feedback through the environment against the stability of the homeostatic loop interactions. These may not be the best names, but they represent functions that the Mechanic will use at all scales from the cellular to the national and international.

The "Director" function is, however, the one of interest here, since much of what we will discuss in the rest of the book involves conscious control. At the level in which the Perceptual Control Hierarchy contains the both the Mechanic and the Structure, the reference value for each ECU comes from a Director, dictating what the perceptual value should be. The Structure is the CEV created by the perceptual function, and the Director uses the tool of telling the different lower-level "Directors" what they must achieve.

At a much larger scale, a government acts as a Director for the social group they are supposed to govern. As Director, the government crates laws and regulations as a Structure in the social part of the environment. This Structure influences what other components of the society can do to influence their Real Reality environment. Just as with the CECU of conscious perceptual control acting as Director with the hierarchy being the Structure, RR determines whether the Director's specifications for the Structure results in a Structure that survives and prospers. At the same time, feedback loops, whether control loops or more general homeostatic loops stabilize and tune the Structure, largely within the design pattern or constraints supplied by the Director.

The word "design" points out that the Director of a Mechanic might be a Designer, not maintaining an existing Structure, but creating a new one that will survive or decay away depending on feedback through RR. The Structure might be physical, such as a highway bridge, or it might exist only in the conscious perceptions of individuals, such as the rules of a new game. Whatever it is, to maintain it will, as McClelland's quote points out, require work, just as does its creation in the first place. If nobody plays the new game, it will not survive the test of RR. It the bridge will not carry the load that traffic imposes on it after its girders and cables corrode without replacement, it will eventually not be a bridge.

Whether Designer or Director, consciousness acts not only as a control system that brings a perceptual value close to its reference value and keeps it there, but also as an organizer of the environment in which this desirable result can be achieved. A Parent cannot tell a child not to want to cross a busy street, but can tell the child not to actually cross without holding the Parent's hand. A Government cannot pass effective laws telling people not to try to get richer, but it can pass laws telling people what actions in support of controlling the "wealth" perception are and are not permissible. Likewise, control in imagination allows consciousness to tell the hierarchy what environmental effects to achieve in order to control a novel kind of perception.

Just as a car mechanic will often test-drive the car on which he has been working before returning it to the customer, so may consciousness actively test a newly constructed perception function by trying to control the perception it produces using means that worked well in imagination. In the process of successfully controlling this novel perception under different disturbance conditions, and perhaps different environmental contexts, various synapses are strengthened and others weakened by the HaH process, until either this perception seems to be not worth controlling or is controlled efficiently enough no longer to need conscious direction.

Reorganization, in a general sense, is simply a communication channel between Real Reality and internal Structure, bringing into being a Structure ever more capable of surviving in RR, or causing some part of the Structure to collapse and dissipate. The Mechanic is the medium for that communication[33].

II.9.2 Consciousness and Genetic Algorithms

Genetic Algorithms are a well-tried way of optimizing systems in a "fitness landscape" (see Kumar et al. (2010) for a short and easily understood review, or see "Genetic Algorithm" in Wikipedia <https://en.wikipedia.org/wiki/Genetic_algorithm>). A system, not necessarily a control system, operates in some kind of environment. Its performance in that environment is given a grade, which is its "fitness" for that environment. The structure of the system is described by a set of variables called, by analogy with organic inheritance, "genes". Sets of genes are combined into ordered sequences called "chromosomes" each of which has some fuzzily defined role in describing the system. That much is the "Genetic" part of a "Genetic Algorithm".

The "Algorithm" part is the particular method of discovering a system description or construction plan that has the best possible fitness. The basic idea is to start with a "population" consisting of a large number of candidates described by random values for the "gene" variables, and determine their fitnesses. With probabilities correlated with their fitness values, choose a number of candidates to "mate" and produce one or more "children". The higher the fitness the more likely is a candidate to mate. Low fitness candidates are not usually prohibited from mating, but their probability of doing so is relatively low. How low depends on the designer of the algorithm in the specific case.

33. Thinking of a masseur as a mechanic of bones and muscles, this comment is simply Marshall McCluhan's "The Medium is the Massage" (not "message") (McCluhan, 1967)

Each mating creates one or more children who join the population. Again depending on the designer of the algorithm, there may be exactly one child from each mating, or matings of two high-fitness parents may generate more than one child. To produce a child the chromosomes from each parent are cut into complementary sections, so that Part A of a chromosome from one parent reconnect with part B from the other parent, in what is called a “crossover” to create the corresponding chromosome in the child. The split point is randomly chosen for each child, and often but not always for each chromosome in a specific child. At this stage, a mutation — a random change in some gene — may be allowed, but some genetic algorithms do not incorporate mutations.

Finally, the fitnesses of the new children are determined, and from the augmented population (the old population plus the new children) the members with lowest fitness are dropped — they “die out”. The pruned new population is the same size as the old one, but contains only the fitter members of the augmented population that included the new children. Some versions add new randomly created members to the population before pruning the population size. Most often, a randomly created member will have low fitness, but it is always possible that one will happen to be found in an as yet unexplored high-fitness regions of the available space, extending the range of possibility for the population as a whole.

Once the next generation population has been generated, the algorithm repeats the procedure as before, by probabilistically mating members to create the next generation of children. After a few such generations, the fittest members usually are so much fitter than any of the randomly produced members of the original population that few, if any, of the originals remain “alive”.

Genetic Algorithms rarely produce populations that include members with truly optimum fitness, but in most fitness landscapes, they usually come quite close rather quickly, and improve slowly thereafter. Ordinarily, Genetic Algorithms take no account of the possibility that interactions among the different members of a population may occur and may influence the fitness of those who interact, competitively, cooperatively or by way of side-effects.

One question should be in your mind if you have not previously encountered Genetic Algorithms: what does a gene represent? The answer is that there is no one-size-fits-all answer. The gene itself may be a single bit within a chromosome that is a string of ones and zeros. On the other hand, a gene could be a very complex Object in Object Oriented Programming (OOP), and in the crossover that occurs in mating, an Object from one parent is swapped for a corresponding Object from the other.

Functionally, if two swapped Objects have input and output terminals that accept and deliver the same kinds of data, they are indistinguishable before they are tested by applying data to their input terminals and observing the dat emitted from the output terminals. The two Objects might contain very different function processes and do very different things with the data, most of which would probably lead to very low fitness of the child, but occasionally might produce a child with much enhanced fitness. Since the next generation will probably not include children from matings of very low fitness parents, these lethal or near-lethal variations will not recur, whereas matings of the enhanced fitness children will probably be found in the next generation. In organic populations such as humans, this effect is sometimes called “hybrid vigour”.

That genes might represent Objects leads to another possibility, since Objects can incorporate Objects internally. The functionality of the more complex “envelope” object is entirely produced by the functions and interconnections of the contained simpler Objects. Each of the inner objects might be represented by a gene, while the envelop Object is represented by a gene that invokes the genes for the inner Objects. In other words, the Genetic Algorithm can incorporate “nesting”, in which the mating not only involves crossover swapping of envelope genes, but also swapping of “inner” genes without necessarily swapping the corresponding envelope genes.

In a recent (December 2020) privately circulated working paper, Rupert Young used just this nested Genetic Algorithm technique to discover an improved hierarchic control structure that solved the classic “inverted pendulum” problem with only three ECUs rather than the four ECUs required by intellectually created hierarchies to solve the same problem equally well.

In one unpublished experiment I did on parameter optimization for perceptual control many years ago, some genes represented rotations in a five-dimensional parameter space, some represented scaling factors, and so forth. They were packaged as simple numbers, packages larger than binary bits, but smaller than function-containing OOP Objects[34]. The best answer I can come up with for my question “what does a gene represent” is that a gene represents whatever is most suitable for the problem at hand.

What does all this have to do with perceptual control, let alone with consciousness and the Mechanic as Director? Rather a lot, I think, if you look at a Genetic Algorithm as a problem-solving technique that usually puts together existing components (chunks of chromosomes, a.k.a perceptions whose control is already incorporated in the reorganized hierarchy) that each might help solve a part of the problem of achieving high fitness in the existing environment.

In control, high fitness implies rapid and accurate control with few side-effects that cause difficulties for other control loops. The problem to be solved may be that the current environment does not permit accurate and fast control, or perhaps any control at all, of some perception by use of the reorganized hierarchy as so far developed. New perceptions with new methods of control must be produced, or invented, as we will discuss in various places in Volume 2 such as the parts of Chapter III.6 on “The Commons of Ideas”.

From the Genetic Algorithm point of view, the “population” of the control hierarchy is the set of perceptual control loops already reorganized to use one another effectively. In humans older than a few months, how to perceive oneself to be standing upright would be one such. How to perceive oneself holding a toy might be another. The genetic-algorithm-like equivalent would be how to put together a selection of object properties (a population) that fit together functionally and that as atenfels would solve the problem effectively and efficiently.

Is this not exactly what we do when we try to solve a problem in conscious imagination? We try connecting different things we know how to do, taking into account any tools that might or might not be available, putting together ever better approximations to a solution and discarding what seem to be dead-ends. Looked at that way, consciousness is not so much a tool the Mechanic uses as it is itself the Mechanic for an individual organism, building and rebuilding structures until one novel structure serves the purpose of the moment, and perhaps of other moments that the organism may encounter later.

The more imagined future moments a problem solution will serve, the wider its “range” and the “fitter” it is in the Genetic Algorithm sense. It is in that sense that Ockham’s Razor (Annex to Chapter 1) considers the range of application of a theory, and it is in that sense that one role of this book is to illustrate the wide range of applicability of W.T.Powers’s Perceptual Control Theory. The evolution of Powers’s thought over the decades of writing about and experimenting with PCT must have involved a lot of conscious imagination work in reconfiguring what serves to improve the range and accuracy of the theory, just as we describe occurring among populations during the evolution of Genetic Algorithm solutions to a problem.

As noted above, Rupert Young[35] has used Genetic Algorithms to produce an efficient stable structural

34. Unpublished, but described at the 2005 annual meeting of the Control Systems Group <https://www.mmtaylor.net/PCT/CSG2005/CSG2005bFittingData.ppt>

35. Unreviewed privately circulated work, December 2020.

solution with optimized parameter values, to address the "inverted pendulum" problem. His nested Genetic Algorithm solution uses only three control loops on two hierarchic levels rather than the four loops that most thought-out human solutions use (e.g. Johnson et al. 2020). This might be something of a surprise, since much of the argument of this Section is that the Genetic Algorithm approach closely emulates Conscious problem-solving in imagination. Why, then, should Young's procedure have found a better solution than was found by human intellectual work?

The relation of Young's work to the human imaginative processes is something like the relation between the analysis of huge amounts of data in the 2020 versions of Artificial Intelligence and human statistical analysis using pencil and paper. The computer can do very many similar computations orders of magnitude faster than can any human, whereas the human can, using massively parallel neural processing, "fuzzily intuit" much that (so far) no computer could. But in conscious human imagination, only a few elements can be treated at a time, trying to imagine what they might do when combined and used in the real world.

Young's Genetic Algorithm processor worked with a population of 100, and for even one generation the fitness values would need to be determined for each of those hundred individuals, plus the children from many matings. A human brain would take days to do this for one generation, whereas the computer, even a low-powered one by 2020 standards, would take milliseconds or seconds.

The human may "see how it works" in imagination when combining different already reorganized atenfels, without being very exact, and may be able to invent a novel tool built from existing components. The computerized Genetic Algorithm might find that invention, given the same problem, but starting from first principles could try out far more possible combinations in the same time including linkages the human would never think of because they would not be incorporated in the perceptual control processes already present for use as component atenfels.

If you have an interest in cosmology, you probably know that, in the Universe as a whole, the galaxies are concentrated in clusters linked by tenuous streams of galaxies like nodes in a network, while most of the space of the Universe is a void containing almost no galaxies. I think of the universe of possible imagination as being something like that, in which reorganized perceptual controls are like galaxies, the clusters being perceptions built from groups of other (lower-level) perceptions, and most of the space of possible perceptions and possible actions being essentially void and inaccessible to imaginative thought. That universe of possibility is, by contrast, all equally accessible to a Genetic Algorithm that starts with no preconceptions about what is worth controlling at low levels and what should be ignored.

Despite this kind of cluster-and-void structure of imaginable possibility, the human is likely to have a wider view of "fitness" than is the Genetic Algorithm, and to have a fuzzy rather than precise fitness measure. For example, the human solving the inverted pendulum problem might not want to invent novel functions and linkages because of the efficiency of re-using old stock, which would be included in the concept of fitness.

This same consideration might well apply to the social problem of "Us and Them" that we address in Volume 2, especially in Chapter III.7. To understand what members of "Us" may do is much easier than to predict the effects of what a member of the unknown alien "Them" might do. Human imaginative intuition and machine-based Genetic Algorithms may perform similar operations, but they are likely to use different raw materials and to have differently-bounded theatres of operation.

The Genetic-Algorithm-like work of conscious imagination described in this section requires the concept of Qualia, as we now discuss.

II.9.3 Consciousness and Qualia

It is often stated that it will forever be impossible to know and to experience what someone else experiences when they say that something is “red” or that it “smells of roses”. By interacting with them we can, however, determine what objects they label “red”, “cold”, or “hard”, and for the most part we will agree that we would not object to those labels for those objects. But when they use those labels are they perceiving what we perceive when we use the same labels? Are our qualia their qualia as well, or are we just labelling properties of objects that we both consciously perceive in our environments? How could we ever know?

Ramachandran and Hirstein (1998) use a couple of thought experiments that they claim would offer a future researcher a way of answering this question: directly couple corresponding neurons in one person’s brain to those of another, so that both people have the same neural firings in the parts of their brains where the future researcher knows qualia to be represented. I do not accept that this device would resolve the question at all, because it does not address what the brain-coupled people would *experience* consciously. What either would experience is very likely to depend strongly on their history, not least the requirements imposed by the different environments in which their neural systems matured through individual reorganization.

Qualia are useful for communication. If I ask you to give me the red box, you are unlikely to give me what I experience as the blue box (unless you are colourblind). Red and Blue are qualia of which we are conscious when we observe appropriately coloured boxes, and you and I are unlikely to disagree as to which box is red and which blue. There are, of course, many context-based colour illusions, some of which depend on unseen lighting variation (e.g. Wikipedia article “The Dress”[36]), others on the colours of objects in the local environment, but we can postulate that when I ask for the red box, both boxes are in similar lighting and local context conditions. Under those conditions, we usually agree on which box is red and which is blue.

Functionally, then, we can agree on our qualia to the extent that we can use differences among them as elements of our conscious perceptual control, even though we may not have any idea how each other *experiences* “red” or “blue”. In this respect, qualia are like Black Boxes (Chapter 11). We can discover what they can do in relation to other Black Boxes and to ourselves, but we cannot discover what they are internally.

So what can qualia do, functionally, inside us and ignoring their potential use in communication with other people (a topic important in Volume 2, especially Chapter II.9 and Chapter xII.11). Before attempting to answer that question, let us consider the construction of a simple perceptual function. Like a Black Box or a White Box, a perceptual function has a set of input terminals and (in the Powers hierarchy) a single output terminal, though we can easily imagine perceptual functions with multiple output terminals.

A Black or White Box, or the equivalent OOP Object functions properly only if each input terminal is fed with the right kind of data. All the engineer’s Boxes of either colour are supplied with varying voltages of electric current (or of some fluid in a hydraulic system). Electric or fluid currents are not distinguished by their nature. Just looking at the input would not allow you to determine that it is nonsensical to provide one input terminal of an adder process with a current signifying 2, and the other with a current signifying blueness. For an adder, the two inputs must signify the same kind of data. The only way to ensure this is by requiring the input to be supplied by a connector that connects only to sources that always and only supply the kind of data suited to the function.

36. Retrieved 2020/12/31

The same is true of perceptual functions in the non-conscious reorganized hierarchy. As Powers often reminded us on CSGnet, the signals passing around the neural system are created only by the firings of nerves, which for simplification of the analysis he chose to collect into neural bundles that cary neural currents. One neural current is indistinguishable from another, and one cannot do what the White Box builder can do by ensuring that the input lines he connects come from the right places. Reorganization can do that job of ensuring the correct kinds of connection, by pruning and building neural synaptic connections according to whether controlling the resulting perceptions can be done with adequate Quality of Control and whether controlling that perceptual output benefits the intrinsic variables on which survival depends.

A well-reorganized perceptual function, then, has a two-dimensional pattern of inputs. The source that connects to an input terminal determines "what" the input is, and the quantity of neural current indicates how much of it there is. Informationally, those distinctions mirror MacKay's (1953) partition of uncertainty and information into "metron" information (What is it?) and "logon" information (How much of it is there?) discussed in Chapter I.10 (see Section I.10.3 or the illustrations in Figure I.10.11).

The perceptual function is created by reorganization to produce a perceptual value worth controlling (or to be used at a higher level in a perceptual function that produces such a perception). Each level of perceptual function in the hierarchy has a fairly stable set of inputs identifiable by their particular input connection lines (less stable as we go to higher levels in the hierarchy).

Conscious control has no such way of labelling its information sources. Consciousness perceives *objects* that are packages of properties — functions that the object might perform. The properties that may be used in controlling some value consciously cannot be tied to specific input lines by prior reorganization, because if they were, that control would already be incorporated in the non-conscious hierarchy, and it is not. So conscious control cannot use source-line identity to label the metron content of the available information. Without that labelling no temporary function simulating a new perceptual function could be constructed.

This is where I think qualia come in. We think of qualia as properties of an object. Shine white light on the box and we will perceive "red", and a certain kind and amount of redness at that, perhaps a pale slightly bluish shade of red, perhaps a strong carmine. "Red" is a category that includes a wide variety of reds, so as suggested in Figure 9.10 or Figure 9.15, "red" can at the same time be a continuum (analogue) variable. Press a finger onto an object and we perceive "hardness" and just how hard it is, as well as a few more qualia such as coldness or roughness. The qualia associated with an object identify properties that we might use to control something not controlled (yet) or controlled differently in the hitherto reorganized hierarchy. Qualia are property labels. We can't use the quantity of "redness" for controlling how to make a paper aeroplane that flies a long distance. We can, however, use the "stiffness" and "weight" qualia (properties) to choose an appropriate piece of paper for the construction.

When we discussed the Mechanic as Director, the Mechanic used qualia in producing a wide variety of skilled performances out of an initial hodgepodge of "instructions". Every repetition of those instructions to produce a similar skilled performance using already constructed perceptual control loops as atenfels builds synaptic trails by the HaH process (Chapter I.9). As this happens, the labelling produced by qualia are being replicated by metron labelling based on source identity, and a reorganized new perceptual function is being gradually grown. The conscious qualia used to build the initial version of the new perceptual function slowly fade out of the picture as the new function becomes stabilized and takes its place within the reorganized hierarchy.

The emplacement of a new function into the non-conscious hierarchy in no way prevents the person from perceiving the qualia associated with the source-line-based signals used by older established perceptual functions, though the older and better-established the perceptual function, the harder it seems to be to single out any one input and determine what object property corresponds to it. A highly skilled performer need not be able to teach the skill. Learning to be a teacher entails conscious rediscovery of the qualia intrinsic to the skill, so that their functional forms can be described as “feel that muscle tension” or “watch the angle of the front wheel of the car ahead of you in the other lane”.

The experience of qualia cannot be transmitted from one brain to another, but the functions of qualia can be and often must be.

II.9.4 Consciousness and Emotion

This section consists of speculations even less supported than those of the previous sections, leading to suggestions that might turn out to be wildly wrong. I do not consider this to be misguided, but the reader should be warned. Powers might have disapproved of such arrant speculation, but as Chapter xII.11 of Powers (2005)[37] shows, he was not averse to speculating on his own account.

> *"The whole system is utterly fascinating, a multileveled system that begins in the brain and continues down -- who knows how far? Perhaps the first-order systems in the biochemical chain are inside the cells themselves, throughout the body."*
>
> ...
>
> *“Not too sharply, I am afraid, but most tantalizingly, a picture begins to form of a second hierarchy of control that splits off from the behavioral control systems at about order two or three; this other branch is concerned with the sensing and control of quantities derived from sensors and from chemical messengers throughout the body. ... the effect is to produce patterns of feeling states that arise as the biochemical balances in the body change in response to the commands.”*
>
> *Powers (2005) p258-9*

I might add to the end of the first quoted passage "*as well as in the microbiome, the assemblage of micro-organisms that outnumber our cells by an order of magnitude and that produce and use a flood of chemicals in which our cells are bathed.*"

In the last few years much research has shown how the microbiome may influence our mental and physical well-being in many ways (e.g., Deans 2016, Flowers and Ellingrod 2015, Kaplan Rucklidge and Rolijn, 2015, Sharon et al. 2016), up to the level of mental disorders such as autism (e.g. Kang et al. 2103, Mulle Sharpe and Cubells 2013, Vuoung and Hsiao 2017), schizophrenia (Dickerson, Severance and Yolken 2017, Kanji et al. 2018, Castro-Nallar et al. 2015), depression and bipolar disorder (Dickerson, Severance and Yolken 2017), and perhaps Alzheimer’s (Kohler et al. 2016, Shoemark and Allen, 2015)These mental “disorders” are distressing, but to conduct a PCT analysis of them is far beyond the competence of the present author. Nevertheless, I can follow Powers in suggesting a possible place for emotion in PCT.

We might start with the fact that we explored above, that there may well exist homeostatic loops that incorporate both neural-perceptual and biochemical components. Collecting these individual “membrane-crossing” loops into a single concept, there exists a loop in the generic sense between the neural system and the complicated system of interacting hormones, enzymes, neurotransmitters and other biochemicals in our bodies, and between those and our various microbiome ecologies. Synaptic activity releases neurotransmitters into the surrounding medium, though most of them are usually resorbed. Glands secrete and release their various biochemicals on neural command, and the biochemical environment strongly

37. Powers states that he wanted his chapter on emotion to be in the 1973 edition of B:CP, but his editor insisted on its removal.

influences how neurons act and interact.

There is an interior biochemical environment for neural output and input just as there is an exterior (outside the skin) environment for physical output by way of muscles and input by way of sensors. The biochemicals in the interior environment might, for example, globally change the gains of control loops and the interconnection strengths among them, whether the biochemicals are hormones released to the bloodstream in the operations of synapses or the waste products of microbial communities. They might locally affect the performance of specific types of control, given that there are many different types of neurotransmitter.

I will not explore these possibilities, but will take a more global, functional view. I think there is much PCT research to be done in this area, but my suspicion is that in this, as in so much else, Powers had a correct fundamental insight, even if it was too much for his original editor to accept in 1973.

Here's one observation: emotion and mental health problems need not be associated with error in biochemical control systems, which is what Powers suggested. Emotions are perceptions, and perceptions are functions of states. Contentment or aesthetic pleasure is as much an emotion as is anger or frustration. So I would suppose that emotional perceptions are some function of the values of biochemical concentrations at different places in the system, particularly in the neighbourhood of synapses (which might well link emotions to reorganization).

We have "dissociated" or "free-floating" emotions like chronic depression or mania of bipolar disorder, which are easily distinguished from the depression that is related to an uncontrollable perceived state of the external environment such as the death of a loved one or the excitement that is easily related to a conscious mental or problem solution. The biochemical states may be the same whether or not an environmental associate is evident, but when they are associated with consciously perceived external states, they cease being "disorders" and become "normal responses to environmental stimuli".

Both may well be the product of closely related processes. When the environmental states produce error and the biochemical states produce, say, the excitement of the chase, success in the chase reduces the environmentally perceived error and changes the effective references down the biochemical side of the control structure to produce an emotional state we might call "happiness", which is then associated with the *reduction* of error in the neural part of the loop. "Happiness" is often associated with solving difficult problems, of which the chase is an exemplar.

When we perceive an emotional state that we don't want, we might ask ourselves whether there is an external environmental reason for feeling that way, or whether it is a purely internal error that needs correction ("planning in imagination"). Whether we can do anything about it depends on whether our reorganization and the current environment provide us with the means to change the amount of error. Perhaps waiting for a transient effect to dissipate, perhaps through biochemical processes, might be sufficient.

Since the nervous system works faster than the biochemical loops as a rule, waiting might be a generic way of avoiding introducing error into a system by too rapid action to fix a problem that might fix itself. We have here just another way of saying “If it ain’t broke. don’t fix it”, even though initially the perceived emotional feeling might suggest that it is “broke”. According to Powers, the rate of e-coli reorganization events depends not only on the magnitude of error, but also on its trend. If the error magnitude is trending lower, reorganize more slowly. If there does not seem to be an environmental reason for feeling an emotion we don't want to feel, and the error does not seem to be dissipating by waiting, then we might ask whether some action on internal variables by way of diet, exercise, or as a last resort, drugs, might work.

II.9.5 Music and Emotion

Duke Orsino: *If music be the food of love, play on,*
Give me excess of it; that surfeiting,
The appetite may sicken, and so die.
Shakespear, Twelfth Night Act 1 Scene 1

These are the very first lines of Shakespear's comedy "Twelfth Night". Duke Orsino asserts that music is the food of love, in that listening to music enhances his feeling of love for the unattainable object of that love. Shakespear presumably assumed that his audience would not consider that assertion ridiculous (though they might not extend the same courtesy to Duke Orsino). The Duke knows from experience that enjoying too much of a delightful food reduces the appetite for the food, and he hopes that by continually experiencing the qualia of the love emotion, he may likewise cease to feel it so deeply. It's a logical conundrum why the Duke feels that way rather than that too much musical food will diminish his liking for music, but we are not interested in that conundrum here.

Rather, we ask whether science, including PCT science, agrees that music is the food of love, its only food if the Duke is correct, or has any association at all with love and other emotions. Since we have not addressed either music or emotions such as love so far in this book, we must introduce some new material which will necessarily be very skimpy in its coverage of either of the two topics, each of which has a history probably far longer than recorded language. Music, or at least deliberately chosen notes, is far older than history, as we know from several bone flutes around 40,000 years old[38], and I think it likely that love, or something very like it, goes as far back as the origin of mammals.

There is a field of research called "archeoacoustics", the study of acoustic phenomena in prehistoric and modern constructions, of which the Maltese underground complex known as the Hal Saflieni hypogeum[39] is a prime example, Stonehenge being another among many. What these ancient religious places have in common is an acoustical resonance in the region of 110 Hz. The same is true of at least some of the places in caves that were chosen for beautifully rendered images of animals such as horses and buffalo, often with cartoon-like hunting scenes (e.g. Fazenda, 2017)

> *Within the field of archaeoacoustics a new arguement is starting to emerge:that within prehistoric monuments or temples that are devoid of an acoustically induced resonance phenomena, other phenomena more natural in origin are frequently found. Such natural phenomena include infrasonic vibrations and magnetic fields both of which can influence the human mind, expanding perception or sensitivity or even enducing altered or 'mystical' states of consciousness. After seven years of studying archaeoacoustic phenomena at over fifty sites, Super Brain Research Group (SBRG)1has discovered that they share the similar properties; natural vibrations or magnetic fields that have a physiological influence on brain waves and consequently on ones emotional state… (Debertolis, Tarabella, & Marcuccetti, 2018)*

Music has a long evolutionary history, as does love. Birds use music, some of it remarkably akin to human musical forms, as I can attest from personal experience. Walking along a trail in New Zealand a few years ago, I heard a bird sing several repetitions of a four-line stanza with an AABA rhyme form and about four or five syllables to the line (my memory fails on that point), rather like a human folk song that might have been heard in the countryside before the invention of radio. The very intelligent corvids would

38. Wikipedia article "Paleolithic flutes", retrieved 2020.12.20.
39. This hypogeum is one of two places in my life where I have felt "reverential awe", the other being Hagia Sofia in Istanbul when I was there almost alone.

perhaps not do that, but there is a reason we call some birds "songbirds". Their songs supposedly are used to attract mates — as atenfels for the bird to correct error in a perception of their bachelorhood.

Whether in the singing bird that perception is conscious is not easily determined, but in the human, it often is. The concept of a "serenade" is of a musical offering to a desired mating prospect, typically from a man to a woman. Mediaeval balladeers performed this role for a client to woo the woman, though the less well-off suitor might have serenaded her himself. Wagner composed and had played on his new wife's birthday his "Siegfried Idyll", the orchestra reportedly being arrayed on the stairs leading to her bedroom. All these common actions seem to be for control of the same perception. That perception is very likely to be induced as a component of a homeostatic loop that includes concentrations of the sex hormones testosterone and estrogen, presumably mainly testosterone in the serenader, estrogen in the love object.

Plato distinguishes several different versions of what translates into English as "love", but we are interested only in erotic love for another person, which is Duke Orsino's problem. That kind of love intrinsically involves control for receiving feedback from the other person, without which the lovelorn experiences substantial error in that controlled perception. The error is not decreased simply upon the receipt of feedback, since the object of the love may indicate displeasure. The lovelorn Duke wants the target to express pleasure at his expressed desire. We go deeply into dyadic feedback processes in Chapter xII.14 and Chapter xII.15, where we discuss protocols, so we will not pursue them further here.

Is there any evidence that music can influence, as part of the reorganization process, a homeostatic loop such as we propose in Section II.8.2? Indeed there is, both from modern science and from considering various religious practices old and new, but all of it is more suggestive than definitive. For example, from a brief consideration of a few religions, my impression, and it is little more than an impression, is that puritanical religions that more strongly suppress sexual expression by women also are more likely than others to suppress musical expression as well. I don't know whether there is sociological data to support such an apparent correlation, but the examples of music-banning in Puritan England during the Commonwealth and by contemporary Islamic extremist sects are illustrative.

Evidence from modern science is rather more persuasive, though not conclusive. Reviews, such as those by Gangrade (2012), Cheng (2013), or Fukui and Toyoshima (2014), all agree that numerous studies have shown that at least some kinds of music can have a strong effect on many endocrine hormones. Here, for instance, is the abstract from Fukui and Toyoshima (2014).

> *Music is well known for its effect on human behavior especially of their bonding and empathy towards others. Music provokes one's emotion and activates mirror neurons and reward system. It also regulates social hormones such as steroid hormones or peptides, and increases empathy, pro-sociality and altruism. As a result, it improves one's reproductive success.*

Different kinds of music have different effects. Music that might be called "jagged", such as "techno-Rock" tend to activate or inhibit pleasure-enhancing chemicals such as dopamine oppositely from the effects of gentle classical music and what we might generically group with serenades and ballads. None of this, however, directly addresses the question of whether homeostatic physiological-neural loops are involved. From the number of "*neurotransmitters, hormones, cytokines, and peptides*" (from the abstract of Gangrade, 2012) affected by music, it might be surprising if such homeostatic loops were not involved, but direct evidence that they are seems to be lacking.

The effects of particular music also presumably differ among people. Anecdotally, I once listened in Germany to a broadcast of what I consider to be one of most calming and beautiful pieces of music ever composed (Brahms Clarinet Quintet). I asked a young girl sitting next to me what she thought of it. Her answer was “Schrecklich” (frightful or horrible). Presumably it would have had very different effects on her hormones and on mine!

Music also seems usefully to affect the immune system, and to have some analgesic value. Cheng, in particular, concentrates on the therapeutic value of music, while Kim et al. (2015) find that Korean Buk drum music can alleviate anaphylactic shock, albeit in mice. Popular books such as Sack’s “Musicophilia” (Sacks. 2007/8), or Levitin’s “This your Brain on Music” (Levitin, 2007) go into much more detail about the effects of music and what music may influence in the brain

From the PCT/homeostasis viewpoint, these reviews suggest that little or none of the research, whether in humans or in other species, is more than suggestive, apart from demonstrating that the detailed analysis is difficult, no matter what the underlying theory may suggest. None of the books or reviews hint at the involvement of disturbances created by music on the values of our biochemical agents as participants in homeostatic loops.

We may, however, guess that when so many biochemicals that ordinarily interact in the body are affected, at least some of them will involve loops back into the perceptual control hierarchy or the reorganizing processes, but I find no direct evidence that this is the case. What we can say is that though the emotional effects of music may average out in certain directions, different people are influenced differently both in the kind of influence and on its strength.

The same is true about what different people control for and with what gains and tolerance levels they control, though the underlying functional processing may be the same for all of them. If we treat music as a disturbance to some controlled perceptions, what perceptions are disturbed and how the resulting error is countered is likely to be highly variable, including influences both on the external environment and on the interior biochemistry. Statistical studies of the effects of music on endocrine function and the reverse are unlikely to be easily interpreted, unless the relevant homeostatic loops are taken as a starting point for the interpretation.

Although several studies mention enhanced sexual desire and activity induced by suitable music, even in mice, we cannot say from a PCT viewpoint that Duke Orsino was correct to think that music is the “food of love”. He might perceive it as such, whereas another person might perceive it as an annoying noise, or as a distraction at one moment or at a different time as a pleasure-enhancer. Perhaps this Section may serve as a pointer for some future PCT-sensitive researcher with physiological expertise to unravel the apparently tangled web.

II.9.6 Emotion and the “stiff” Personality

Having introduced the possibility that hormones, enzymes, neurotransmitters, and the whole biochemical panoply of the body and its various microbiome ecologies might form part of one large homeostatic network, we now return to the emergent property, “stiffness” (Section 8.1). Stiffness is irrelevant to control of the perception in any one loop. The force opposed to the disturbance is proportional to the difference between opposed outputs, regardless of their sum. Stiffness depends on that sum, and applies orthogonally to the axis along which the outputs are opposed. It is an emergent property of a structure that includes opposed forces, as tensegrity structures must do.

Since we are concerned with control rather than with mechanical stiffness, it is an emergent property of a structure that requires either one control loop that has two opposed outputs, as so most control loops, because neurons cannot have negative firing rates (Figure 4.7 and Figure 4.8) or of two opposed control loops such as exist in a standard conflict situation. It is not hard to imagine in this the genesis of the personality quality sometimes called “stiffness” or “rigidity”.

In some cultures, children, or at least boys, are discouraged from showing visible emotion. They learn to control an image of "self as seen by others" that incorporates the ideas of Rudyard Kipling in his poem "If":

If you can keep your head when all about you
Are losing theirs and blaming it on you,
If you can trust yourself when all men doubt you,
But make allowance for their doubting too;
...
Yours is the Earth and everything that's in it,
And—which is more—you'll be a Man, my son!

Kipling's "Man" is resilient and therefore strong, in the tensegrity sense, but the same control of self-image can result in a personality we call "stiff", apparently emotionless and rigid in the application of rules to the actions of others, and in many cases to his own actions. How might we reconcile these different possibilities, which may arise from similar childhood environments? Kipling was a strong supporter of British colonialism, and the character that he described is the ideal British colonial officer who would agree that his "is the Earth and everything that's in it", at least in the domain he is sent to rule.

In English fiction, some colonial officers are shown to be rigid and rule-bound, some to be strong and dedicated to the welfare of their subjects, and some to be stupid and weak, but only the last category is shown as expressing their emotions easily. Why should such different personalities come from such similar environmental conditions during reorganization? To say "genetic differences" is easy, but to explain what the genetic differences, if any, influence is not. Here, I propose a completely speculative possibility, based around the concept of the parallel internal biochemical and external environmental fields of action for control.

If some controlled perception is blocked, perhaps by conflict, perhaps because of environmental facts such as a simple wall or as complicated as a Kakfa-esque tangle of bureaucratic red tape, the effect of the continued failure to control will quite probably be an escalation of output, perhaps leading to the explosion of a "Bomb in the Hierarchy" (Section 6.5) or a conscious set of changes of action method. The former may correspond to a "rigid" rule-bound personality, the latter to a more flexible one.

In our first discussion of the "stiffness" emergent property of control structures, the key point was that tension between two opposed, conflicted, control systems affects the ability of systems to control apparently unrelated properties. In a tensegrity model, the tensioned opposition can, in part, take the place of a rod (Figure 8.15). The substitution is not exact, because the two ends of a physical rod are always a fixed distance apart, whereas changing tensions in the all-wire meeting points that are the ends of a virtual rod can move them independently. But for many purposes the oppositely directed tensions can serve the same purpose as a solid rod in compression.

Elsewhere we make the case that reference values derived from the outputs of higher-level structured perceptual control units can serve as rods that do not involve opposed tensions, but here we are considering the stiffness that is produced by perceptual controls acting in opposed directions — conflict. The stability that is introduced by this conflict may ease control of something else entirely that uses the conflicted CEV in some manner, as does the conflict in Figure 8.16 that improves control of the components of two complete loops of units linked only through their side-effects. Not all conflict is bad, but less energy might be used in stabilizing the variable in question by other means.

If we consider a physical tensegrity structure, the wires will break if the tension is too high, or the rods may buckle under excessive compressive stress. If the wires are too slack, the structure will not easily support itself, and will not distribute very well any applied stresses. The structure performs best when the internal tensile and compressive stresses are intermediate.

The control analogy to the tensile stress on a wire is the error, and to the compressive stress on a rod the rigidity of the relations among the reference values of the various control units that support a higher-level ECU. Too tight a specification of the reference values, too high the loop gain, or too low the tolerance for differences between the perceptions and their reference values are the analogues of excessive tensions and compressions in the physical structure. Too tight tolerances induce conflict that lead to stiffness, and the entire structure becomes rigid, inflexible, and liable to collapse if one of the wires "breaks" under the tension — control fails or is given up.

When we go back to the possible biochemical side of the hierarchy with the concomitant emotions that depend on what is circulating in the neighbourhood of the neural system, the analogues to the physical tension and compression stresses have a namesake in the form of "stress hormones" such as cortisol, cortisone and adrenalin. Psychological stress is the inability to bring perceptions to their reference values, we have just argued is a component of rigidity, with its inherent incidence of conflicts.

Putting all these notions (for that is all they are) together, we get a picture of a stiff, rule-bound person, full of conflicts that make it difficult to control perceptions effectively. This "stiff" person will probably be chronically angry or likely to have outburst of temper, and to have very little resilience against small disturbances, especially at the higher levels of the hierarchy such as the controlled "self-image" perceptions by self and by others.

We are moving into the realm of interpersonal interaction. Interactions between pairs of people or among many people form the core of Volume 2, where we will refer back to many of the constructs developed in this Volume 1, as we discover how PCT might be useful in thinking about Culture, Language, Power, and Politics.

Chapter II.10 Thinking and Reorganizing

Most of this Chapter is concerned with consciousness and control within the broad context of Perceptual Reality (PR) rather than Real Reality (RR). It might well have been titled "A Tale of Three Theories", a title I give to the first Section of the Chapter. The three theories are PCT, Predictive Coding, and a Theory of Narrative Thought (TNT, Beach and Wise, 2022). All deal with the interaction of the individual with the outer world, though our conscious thinking often does not. In this chapter, I propose a theory of Narrative Thought (NT), different from TNT that deals with imagination and rumination (such as solving mathematical or artistic problems, devising new product designs or programming techniques). I suggest that the human mind, and perhaps the mind of a few other species, employs a common structure that combines into a single cooperative unity these four apparently different approaches to thinking and behaviour.

In considering Predictive Coding, which is a popular theory of behaviour, and the Theory of Narrative Thought (Beach and Wise, 2022)[40], which both deal only with conscious processes, I will use in NT only their most basic concepts, not their further elaborations. In, the basic concept of Predictive Coding is that if one can perceive the current state of the environment, one can predict the effect on one's perception if one manipulates the environment in certain ways. Beach and Wise use a more stochastic approach, founded on "rules" that represent the probable result of prior "causes". Thinking about the sequences of events likely to happen as a result of applying one's rules to causes that might happen or have happened in the current environment results in a "narrative", a structure unique to each thinking individual.

In what follows, I assume that the reader is familiar with the basic PCT hierarchy of perceptual control and the basic nature of the perceptual input function. This hierarchy is built by a process of reorganization that improves its functioning. Its operations are generally non-conscious, computationally efficient in the sense of energy efficiency, and fast. This "Tale of Three Theories" will argue that reorganization occurs in feedback between the PCT hierarchy and related hierarchical structures implied for the other two. Some of these feedback loops influence reorganization not only of the PCT non-conscious hierarchy but also of the hierarchies of narrative thought and of predictive coding.

I start by outlining my own interpretations of both of TNT and of Predictive Coding, interpretations which may well diverge from those of the proponents of either theory.

II.10.1 A Tale of Three Theories

In our Perceptual Reality (PR), we see a "landscape" that contains myriads of different but inter-related objects, each of which has many different kinds of properties, such as softness, location, odour, colour, and so forth. In the reorganized perceptual control hierarchy, the scalar variables controlled are perceptions of the degree to which the pattern of inputs to the perceptual function matches the pattern of inputs — the category — selected by the perceptual function.

The inputs to a perceptual function form a vector of scalar values, but none of them are "values" of the objects, Objects, or White Boxes discussed in Chapter I.11. The objects we experience are components of *conscious* perception, not necessarily components of the non-conscious single-valued perceptions of object properties that are reorganized into the perceptual control hierarchy.

At this point I will seem to contradict myself, but I believe I do not, when I say that each perceptual

40. Everything about Narrative thinking throughout this book was inspired by an e-mail exchange with James Wise (initiated 2022.01.14), but is otherwise independent of the material in Beach and Wise (2022).

function in the PCT hierarchy is actually a category recognizer, and one possible category that could be perceived is a category of object. For example, consider the "chair" perception whose locations in three dimensions we discussed back in Section I.5.4 as being built from four perceptions of chair legs, one of a chair seat, and one of a chair back.

When these four leg perceptions, one seat perception, and one back perception are configured in a certain way, to form a perception of "chairness", the resulting "chair" constrains the relations among the six different location perceptions and the six different motion perceptions in such a way that the 36 possible pairwise relationships have only six degrees of freedom among them. Those constraints define the object category, at least in the location and motion domain. The relevant perceptual value controlled in the hierarchy is the degree to which the configuration moves like a single entity rather than like a collection of six independently moving objects.

If this seems to contradict the concept of there being a separate category interface (or "level" in Powers's version of a perceptual control hierarchy) that lies between the analogue and logical components of perceptual control, the contradiction is only in that the Powers "category" is discrete and can be verbally defined, as, for example, by adding it to a dictionary, whereas the "category nature" of a perceptual function has no well-defined edges, but rather a distribution of the strength of different properties in determining how well a pattern of lower-level inputs matches the perceptual function.

On the other side, PCT does encompass conscious perception, which is the only kind of perception we experience and can talk about. In principle, we might be conscious of any property of an object whose changes alter how we perceive the object as a category, such as of "things I could use to crack a nut". "Things I could use to crack a nut" are potential atenfels when you do want to crack a nut, but are simply conscious observations of the world otherwise. The more we learn of the world around us, the more such perceptions do we control non-consciously in the reorganized hierarchy as a "matter of course" when they are needed. Later in this Chapter, we will refer to such "knowledge" as functional components of a "narrative soup".

For example if I want to perceive a nutshell opened so that I can eat the nut inside I may just pick up an object I might automatically label as a "nutcracker" and use it unthinkingly in a well-reorganized way. Or, if I am a chimpanzee who has watched an older chimp crack a nut, I may put the nut on a stone and hit it with another stone. These learned actions in control of a perception of the nut with a reference of "cracked" may eventually be reorganized into the control hierarchy, but they do not start there. They start with the growth and decay of synapses that connect one neuron to another (Chapter I.9).

The view of a perceptual input function as a category recognizer offers a direct connection from and into Predictive Coding, a two-way connection by means of which a predictive coder might invent a nutcracker, might add the category "nutcracker" as an object whose learned use turns it into a PCT atenfel, or in which PCT discovery might converge onto a standardized means for cracking nuts on demand and be used by a thoughtful process in which a predictive coder might predict a perception of a cracked nut by selecting a "nutcracker" process (use of a "nutcracker object" or hammering the nut between two stones).

This applies even to very complex perceptions. Incorporation of a complex function into the non-conscious part of the hierarchy of perceptual control is not something that happens suddenly. It develops over time from being something consciously worked out into eventually being understood automatically and without conscious effort, just as a child learns to work out and refine the meanings of written or spoken words. In my own experience, I once thought that "fatal", as in "a fatal accident" meant simply "serious", and only later was the meaning refined in my mind to refer only to an accident in which someone was killed.

Using another personal anecdotal example, at one time in my engineering physics studies, I could read as a unit like a single word the frequency integral form of a Fourier Transform ($S(f) = \int s(t)e^{-2\pi ift}dt$). Just as effortlessly, its meaning and import in its mathematical context made good sense in the problem or illustration in which it was used. I still easily perceive in imagination what it does in a space in which half the dimensions are time and half are frequency (it performs an orthogonal rotation of the space).

To write this once familiar formula 65 years later, I had to read it from a source, in the same consciously laboured way as a child might read a word letter by letter, but I still visualize its meaning when I see it, and know consciously that in radian measure rather than in cycles the 2π in the exponent is eliminated. I might use the entire formula within the brackets in a Predictive Coding way of solving some mathematical problem after a lot of conscious thought that used this formula the same way one uses the movement of ones legs to walk to the local shops, totally non-consciously. In a narrative, the entire formula has an effect with much the same relation to its elements as a molecule of salt (Sodium Chloride) has to the component atoms of Sodium and Chlorine.

The "Free Energy" aspect of Predictive Coding is a technical property that applies identically to both PCT and Predictive Coding. There's no need to go into the technical details, since what the "principle" means is simply that the processes tend to minimize the mutual uncertainty (Chapter I.10) between, in PCT the reference and the controlled perception, and in Predictive coding the mutual uncertainty between prediction and result. The latter is related to the concept of "surprise", the deviation of the result from the prediction, which is the information gained about the world from observing the result. Surprise corresponds to the PCT construct of "error", and is related to the "rattling" measure (Chapter II.5). The more the surprise, the greater the rattling.

"Narrative Thought" as I understand it joins this picture by concentrating on the passage of time. Narrative thought occurs in conscious imagination. Powers linked imagination to each individual ECU by treating its perceptual input as though its reference value had been achieved by some process that occurred in imagination. We do not do that. Instead, we argue that in narrative thought one imagines that certain events can happen, and each event alters the stage set in which other events are likely. For example, if one event is that the refrigerator power plug is pulled out, then a sequence of events in which ice turned into water in the refrigerator, followed by an event in which water would be perceived on the floor, would have an increased likelihood of happening.

Beach and Wise (2022) call such a highly probable sequence observed in Perceptual Reality (PR) a "rule", a word which seems inappropriate because it implies subservience of the rule follower to the rule maker, whereas the sequence is just an instance from category of event successions we might call "melting" because we have seen what happens when ice gets warm. In reality, it might never happen that ice in the unplugged refrigerator would fail to melt and flood the floor, but the narrative thinker might have a train of thought in which the unplugging event would not be followed by the melting event.

Trains of thought are sequences of imagined events, with no necessary relation to anything in the Perceptual Reality (PR) of the thinker. An imagined event may take a few milliseconds to imagine, but if the imagined events were to happen in Real Reality (RR), they might take years. Total melting of the Antarctic ice cover is easy to imagine in a fraction of a second, but would in reality take tens of millennia. The melting event would have that long duration, during which many related events would occur, sequentially or in parallel.

In a very brief outline form, here is my concept of how three theories (including NT but not TNT) work together as a unit in the context of Education. This outline (here lightly edited) was originally drafted (2022.02.06) as an e-mail message to Eetu Pikkarainen after having read a draft of his Chapter for the forthcoming second Handbook of Perceptual Control.

"Education", as partly opposed to "education", "teaching", and "learning", refers entirely to conscious perception, to concepts that can be described in curricula, in words. Lower-case "education" is the larger set of what a student might learn during maturation, in school and out. Teaching is a part of education, but encompasses more than Education, because it also includes teaching by example (as various non-human mammals do), teaching by asking the learner to do, while learning includes learning by modelling what others do, by thinking about problems, or by incorporating word-based stories (which I tend to think of as large-scale narratives, much as Powers thought of systems as networks of controllable perceptions).

What this implies to me is that "Education", "education"[41]*, teaching, and learning all suddenly fit within the vision I have of Narrative Thought, which refers to sequences of events that change the world in trivial or dramatic ways, Predictive Coding, which is based on a hierarchy of conscious thoughts about how to create a wanted event, and PCT, in which the base of the Predictive Coding hierarchy takes for granted the ability to bring about certain variations so that they conform to these basic requirements. Predictive Coding says "If I put X on top of Y I'll get what I want, and I can do that", and the "I can do that (put X on top of Y)" bit is true if "that" is a reference state of a perceptual variable controllable in the PCT hierarchy.*

Narrative Thought and Predictive Coding both rely on the individual having learned to "do that", but as far as I know, they no obvious way on doing the necessary precursor learning. Narrative Thought (always conscious, in imagination) just assumes it. Predictive Coding, also conscious, could use memories of successes and failures to learn by reinforcement, without determining how reinforcement works, whereas learning how to bring about desired states is at the heart of PCT and the reorganization of its non-conscious perceptual control hierarchy.

Looking upward in the hierarchy, Predictive Coding seems to have no obvious way to determine what should be predicted. It has no reference state. But that reference state could be supplied directly by PCT or by Narrative Thought. Narrative Thought says "for this event to happen, thus and so must become true", and Predictive Coding says "I can do that for you", but when it gets down the hierarchy to needing simple things to be done in order to bring about predicted states, PCT says "don't worry about that. I can do it without thinking".

Lower-case "education" and "learning" apply to all of that, at different levels and by different means.

So let us consider time and my version of narrative thinking (NT) before we follow up more carefully how TNT, NT, Predictive Coding, and PCT might work together in mutual support of the health and survival of an individual of whatever species.

II.10.2 Time and Conscious Thinking

Many of the distinct theories of *conscious* thinking appear to have little in common but some seem to mesh with PCT in similar ways. One is labelled "Predictive Coding" or "The Free Energy Principle" mainly influenced by the work of Karl Friston (e.g. Friston, 2010; Seth and Friston, 2016). We will discuss how Predictive Coding ties in with the reorganized and reorganizing Perceptual Control Hierarchy

41. "Education" with an upper-case "E", unilaterally deciding what a student ought to learn, versus "education", setting a stage within which a student is likely to learn what the teacher wants. The word is based on Latin "*ex ducere*" "to lead out".

later in this Chapter. Another, called the "Theory of Narrative Thought" (TNT), was developed in Beach and Wise (2022). TNT is based on what they call "rules", regularities in the way the world works, many of which represent apparent causal influences. We will discuss TNT first.

TNT assumes the thinker is thinking about how to act on the current Perceptual Reality to produce a desired effect. In everyday language, it deals with causes and effects as primary processes. TNT deals with conscious thinking, whether or not the thinker is interested in the current environment. Thinking in TNT includes thinking about the current environment, but also about matters unconcerned with matters unrelated to the current environment, such as thinking about the development of some mathematical proposition, or about the possible status of an old friend one has not seen or heard from since childhood.

The TNT "rules" are structures that inherently incorporate a time dimension. Rules stand out as perceived structures because they are repeated more often than would be expected from a random series of changes in perceptual variables, just as is true of the patterns that allow synaptic growth and decay to form perceptual functions. The levels of perceptual control Powers developed intuitively incorporate time-based structures at every level from his "Event" level upward, but Beach and Wise treat events differently, partly as controllable perceptions, partly as highly likely consequences of causes perceived earlier. These frequently observed sequences of events are "rules" that can be used in conscious thought to make things happen in Perceptual Reality.

What is an "event"? That depends in your point of view. To Beach and Wise, an event is primarily a cause of some consequent event. By changing something in the environment, an event causes some subsequent change in the environment — another event. As I see an event, it is a perceptual phenomenon, a change in the state of the universe of thought in the perceiver's mind. The changing data need not derive from the state of the perceived environment, since all thinking is done in the thinker's imagination, and the event might occur in a fantasy universe inaccessible to an outside observer.

Either way, a "perceived event" and a "narrative fragment" are functional synonyms, the event not necessarily being consciously perceived, the narrative always being its conscious representation. This includes the incorporation of imagined events along with events perceived as occurring in the environment. When I talk about "a perceived event", I am also talking about a "narrative fragment", and vice-versa.

Predictive Coding, TNT, and NT all depend on some kind of conscious understanding of how some aspects of the perceived environment are likely to change what our senses will tell us when we take some specified action. This understanding develops over time, as the "thinker" matures and experiences a greater variety of types of event.

Quite early in life, we learn that if we drop something, it will fall (unless, as we discover from experience, it is a helium-filled balloon). Infants often delight in discovering that they can control an object's location (with a reference value for it being on the floor) by performing a sequence of actions, starting by controlling for perceiving that they are holding the object and ending by releasing it in a spatial context that leaves the object some distance to fall.

Such understandings do not come cheaply, at least not within the domains of the individual theories of conscious thought. They depend on material learned in some other way. Prediction cannot be very useful as a guide to action if you have no idea what your action might do, nor would be a narrative in which what happens depends on knowledge you don't have about what might happen. I argue that PCT reorganization, such as we have discussed in earlier chapters and that requires no prior knowledge other than what has allowed your ancestors to survive, is a saviour for all theories of control by conscious (rational) thought.

If what we have developed so far in this book by using PCT is anywhere near a correct White-box emulation of the Black Box of Real Reality, that "other way" is the development of the perceptual control hierarchy, which we have suggested is built on the growth and reduction of synaptic strengths in the neural network of the body, and the development of biochemical homeostatic loops that interact with the electrical functioning of that network. Predictive Coding and TNT are not concerned with this level of detail, both being based on the assumption that already learned processes will be consciously available that will lead to the expected result. They do not include how these processes might be learned, and I will suggest that they aren't. Instead, I suggest that they use the abilities previously developed in PCT by reorganization.

Time is central to TNT, is not directly involved with Predictive Coding Theory, is important to the stability of perceptual control loops in PCT and is critical to our survival. Every control loop has a loop transport lag that consists of two components, a base lag and a processing lag. The precision of control depends on how fast the processing around the loop works, and if it produces approximate values very quickly, those approximations might well save your life.

In everyday life, there is an adage "*The best is the enemy of the good*", which often applies to the processing in a single control loop. If you need to *think* about, or consciously dredge out of memory, what to do to avoid the rock hurtling toward your head, you are likely to be seriously hurt before you decide on what to do. If, on the other hand, you act to control an already reorganized perception with a reference value for where the rock should go relative to your head, you quickly dodge without thinking.

In a quite different way, time is at the heart of Narrative. Narrative thinking necessarily considers the flow of time, and not simply time in the form of sequence ordering. A narrative about perceptual control might, for example, stretch out the functioning of a control loop from the two-dimensional image we have been using, into a helix with its axis along the locally perceived time dimension, making explicit that the effect of a change in the disturbance propagates over time around the loop, and that the disturbance value might have changed again by the time the first effects of that change return to begin to oppose it, as in the often mentioned. That long sentence is itself a narrative fragment, intended to bring your concept of narrative thinking closer to mine.

Time, to repeat myself, is at the heart of Narrative. A narrative thought, whether or not it is ever expressed in words, consists of imagined events, changes in the imagined Perceptual Reality (PR) of the thinker's Universe of the moment. The thinker's PR may represent their perception of the external environment of the moment, or their perception of some completely imaginary world. A narrative fragment situates an event in its own local time, which can, in some novels, jump back and forth across centuries between fragments of the overall narrative. The local time within the narrative is not the time within which the narrative itself plays out in the thinker's mind, nor is it the time of the hierarchy of perceived "Now" durations of Section I.7.2.

Indeed, when we leave our usual haunts, perhaps on holiday, after a while the time stream in which the remembered events of the holiday exist may be perceived as decoupled from the time stream in which events have happened in our home location One can imagine a narrative of either normal life or holiday life, but the events that couple them get lost. The narrative of the holiday gets cut off from the stream of the narrative of the everyday, much as an ox-bow lake gets cut off from the stream of a meandering river over fairly flat terrain. I call this phenomenon an "Ox-Bow of Time".

The ox-bow is not the only structure that departs from a linear flow in the thought narrative of time. We can imagine event streams happening in parallel until they meet, if they ever do, a situation often mentioned in news reports as someone being in the right place at the right time (or wrong place at the wrong time). Later, when we deal with autocatalytic processes in narrative, we will talk about merging and splitting streams of events in much the same way.

All these different flows of time must be kept distinct if, as an Analyst, one is to think of narrative thinking. In the internal time of a narrative, at whatever sequence they occur within the narrative representation, some events happen before or and some after others, while other events may be happening at the same time in a different narrative stream. The ox-bow is by no means the only occasion on which one chunk of a narrative sequence of events is separates in memory from another, but it is one with which many people are familiar. We will look more closely at narrative and time perception later, but now we turn to the unbroken flow of a unitary narrative stream considered as a sequence of events.

A narrative fragment consists of one unique event that occurs over time. During that event, something happens in the Universe of the narrative to change in the thinker some perception of the state of that Universe, the thinker's Perceptual Reality. That changing perception may be as simple as the brightness of a light, or as complex as the devastating social and physical effects of a war. Within the narrative, any single event has a duration within a narrative hierarchy built from a series of shorter and simpler events, just as the perceptual control hierarchy reorganizes to build complex perceptual functions that are built on simpler ones that produce simpler perceptions, quicker to control.

We will associate this hierarchy of simpler events within longer and more complex events with the inherent hierarchy of "now" in the perceptual control hierarchy (Section I.7.2). It is not that hierarchy, but is closely allied with it. It is not the same, because the perceptual control hierarchy of "Now" durations includes only events caused by the actor's perceptual control, whereas most events have other external causes and effects. For example, if a visitor rings my doorbell, the "visit" event includes the time of the event of the doorbell ringing, the event of my controlling to answer the door, the event of the visitor being in the house, the event of the visitor leaving the house. Most of those events are likely to contain shorter sub-events, only some of which are events of my perceptual controlling.

We will follow this narrative of thought that I am putting into words, words that have a very different sequence from that of the interior narrative whence they came, but only after we provide a simple outline of some salient points of another important theory of conscious thought, Predictive Coding.

II.10.3 Thinking as Narrative

What do I mean by "Narrative" and "Narrative Thought"? In common speech, "a narrative" usually refers to a story told by one person to one or more other people, either orally or in writing. This book is such a narrative. I may call this use "Communicative Narrative" if the distinction is important between that and an "Interior Narrative", a coherent braided stream of conscious thought that the thinker might not even be able to put into words. This interior narrative is based on conceptual frames such as "if this happens, what then?", "what might have caused this to happen?" or "what else was going on while this was happening?". It is this internal network of thinking, in which one imagined event leads to another or prevents another from happening, that I call Narrative Thought.

An internal narrative might be composed entirely of imagery, non-verbal sound patterns, sequences of touches, smells and tastes, or any kind of quasi-sensory imagination, among other, more abstract, possibilities, or any mixture of them. The key, though, is that interior narrative occurs in the imagination, not as part of active controlling. I will later argue that Predictive Coding is required for *effective* narrative thought, while PCT is required for *effective* Predictive Coding.

For now, however, I want to consider Narrative Thought in isolation, as interior narrative, not communicative narrative, nor narrative whose component events might be tested against the perceived results of similar events in the external environment. A trivial example of such a test might be a narrative like "If I pull this block out that I can see, the tower of blocks will fall down", which could be tested by actually pulling the block out and seeing what happens to the tower. This kind of testable narrative thought is not the topic of the next few paragraphs. A narrative in which it is possible that the tower would remain floating in the air is testable in the Perceptual Reality of the thinker's external environment, but in the Thinker's imagined Universe for this narrative, the tower might be able to float comfortably after all the supporting blocks were removed.

Sometimes an interior narrative becomes partly observable, such as when we might say someone is talking to himself, but these, usually fragmentary, sound productions seldom tell us much about his ongoing internal narrative. As like as not, if "talking to himself" persists, we might consider him mentally deranged (of course, women do it, too, but the usual image, at least in my mind, is of an unkempt, unshaven man in ragged clothes wandering apparently aimlessly and perhaps drunkenly along the street).

My idea of "Narrative Thinking" is based on the concept that all (conscious) thinking takes the form of narrative, in the sense that one conscious thought depends on an earlier thought and is likely to lead to more future thoughts. Neither PCT nor Predictive Coding place much emphasis on this apparent fact, though the eleven level perceptual control hierarchy intuited by Powers includes a "sequence level" for control of the order of "events" which are contributory to "sequence".

We will soon discuss causes and consequences in that kind of context, but first we look a little closer at simple word-based communicative narrative by using as an example this sentence fragment: "*As she was walking along the path, well trodden by long-ago ancestors, ...*". It is not so simple as it may appear.

There's much to consider, even in this story fragment that I just invented. Let's think of (imagine) walking and separately think of the implications in the fragment of the word "walking" for the larger narrative. The central woman's walking action is dynamic, slower than running, but faster than "leisurely strolling". Since the action is dynamic, it is about time, but only so in part, because walking, like running and leisurely strolling, is a sequential series of repeated events in which one foot is lifted and replaced on the ground in front of the actor. According to PCT, that sequence is a non-consciously controlled perception, as are the perceptions of muscle tensions needing control for the action to succeed as an element of controlling for perceiving oneself to be "walking".

If you are a novice race-walker, to avoid being accused of running, at each step you have to be aware of how you place your foot and when you lift again it in temporal relation to the lifting of the other foot. By the time you are in a position to enter in a race-walking event, however, you have ceased to be conscious of correctly performing the stepping cycle. Walking can be something of which you can be aware and can include in a narrative, but when you actually do it, it is not something of which you are ordinarily conscious. In the fragment, the word "walking" is an element of the narrative that is presumably intended by the writer to allow the reader to visualize the dynamic scene.

The initial word "*As*" in the fragment tells the reader that while in the story she continues to walk, something else whose nature we are ignorant is happening at the same time that will contribute to a more complex, higher-level narrative of which this fragment is a part. It tells us that her walking is set in an extended moment in story time. Later in the sentence the fragment is also set in a different and longer time-frame over which the path has existed, because of the reference to how the path came into being, which may well turn out to be relevant to other parts of the story.

The reader might be induced to imagine centuries of parties hunting prey in the area, later convening with their spoils at some frequent gathering place. Perhaps generations of traders used the path to travel between two settlements, or the reader might imagine some quite different internal narrative about the history of the path. If "long-ago" does not induce you to think of some such side-narrative that might possibly become important later, so be it. You at least will probably remember that the path is old if that fact becomes relevant later.

What of "*she*"? This little word also spawns a historically branching narrative part of the story. "She" might be someone socially important to other people in the implied higher-level narrative. If she is, there's an offshoot narrative in those relationships that might be pursued or may have already been pursued. "She" might be on the run, walking now because she has run far enough to be safe. Equally probably, she might be on a pleasant afternoon stroll, looking for mushrooms or picking wildflowers in a forest.

If this fragment forms the very first words of the story, all these possibilities and many others might spring to the reader's consciousness in the form of questions like those that surround a basic control loop — Why? What? and "How". As in a control loop, "Why" is likely to be the most salient of the three. "What" happened earlier that led her perceived time to this, her "now" (Section 1.7.2). though earlier in her time is not necessarily earlier in the linear sequence of the narrative.

Even this trivial fragment exists in a forest[42] of possible relationships among fragments of narrative, not all of which are necessarily mutually consistent, but which would be consistent if the narrative is about a Perceptual Reality (PR) version of Real Reality (RR), rather than about an imaginary world that might or might not coincide with some historical ("was walking") fact of the real world. If it is, the fragment asks the question not only of what is simultaneously happening in the story-world, but also how the narrator (for the fragment's form tells us there is one) came to know that "she" was walking specifically on a long-trodden path.

Why, for example, does the writer bring up the concept of "ancestors". Are they ancestors of "She" who is walking, and if they matter to the story, does that imply she is of noble birth, whether or not the other characters in the story perceive her as noble? Does "She" apparently thinks about older users of the path, but does she think of her own ancestry, or people of earlier generations in general? Are "ancestors" required in order to let us assume that her family has lived in this area for many generations? Will that matter later in the story narrative? Are we going to be told next about the spatial context of the path, whether it leads to or from some place significant to the story? Does it matter whether she is walking through a pleasant wood or a tangled jungle?

The word string "well trodden by long-ago" is more consistent with an environment generally and historically obstructive to control of walking direction than would be an open field, and "walking" suggests little or no obstruction on the "well trodden path", though the fragment could easily be continued with "… her way was blocked by …". How much of that list of questions does the hypothetical novelist author of the fragment expect or hope that the reader will entertain?

II.10.4 The Narrative Setting

Earlier (Chapter II.2), we based our discussion of autocatalytic loops on Kauffman (1995), who used as a basic medium a "soup" of chemical elements from which a creative autocatalytic loop would almost inevitably emerge. The components of this soup were the 90+ elements that would have been found in Nature when life on Earth began, together with any molecules that formed easily from the elements, such as H_2O (water) together possibly with more complex molecules known to exist in inter-stellar space.

As we progress through this book, however, we will find creative autocatalytic loops constructed from very different kinds of ingredients. One that we explore here is the "narrative soup", which contains all the the ingredients of which a narrative network would be composed. We treat them as being available in a "narrative soup" analogous to the chemical "soup" in which chemical creative autocatalytic loops and networks almost inevitably form and from which stable homeostatic loops and control loops emerge in

42. Did my use of the word "forest" link in your mind to the idea that the path in the fragment might pass through a forest? Keep this kind of possible link between parts of a narrative when we begin to discuss possible stabilities in the narrative "soup" in the form of creative autocatalytic loops.

structures of ever-increasing complexity (Chapter II.2 and Chapter II.3).

What are the “states” that might change to create a “narrative soup”? And where did they come from initially? In Kauffman’s chemical “soup” from which (with the addition of some energy flow) all complex molecules develop and form creative autocatalytic loops and networks, the elementary components consist of the 90+ chemical elements and the molecules that tend to form easily between pairs of them. The “states” in the narrative soup that correspond to Kauffman’s chemical elements must be perceptible if they are to be elementary units that can interact in a mind. They are the patterns that produce outputs from perceptual input functions, category recognizers. Perceptual functions become connected into a perceptual control hierarchy in ways we discussed in Volume I. The “narrative soup” is a structured network, not a medium in which perceptual function float freely.

In Kauffman’s chemical soup, a reaction consists of a molecule of one available type combining with another of the same or different type to form a molecule. In a narrative soup, an event consists of change in at least one perceptible state as a consequence of some event in RealReality (RR). Two elements combining to form a molecule in Kauffman’s soup is an event in a narrative soup that treats the relationships of the elements as states. That particular kind of event does not involve perceptible states, and therefore cannot qualify as initially available for the formation of founding narratives in the mind, conscious or non-conscious.

The issue of the origin of the initial components of the narrative soup therefore demands further clarification of what those components might be. One could take “initial” to be any starting point back to the beginning of life billions of years ago, which would be ridiculous, so what we treat as “initial components” of a narrative soup are to some extent arbitrary. They could be, for example, the narrative fragments of the events of yesterday, but those components were generated yesterday, and are about only the events that were perceived at a relatively high level, such as breaking off a meeting to have lunch.

One possibility is the birth of the individual whose personal “narrative soup” (or perceptual control hierarchy) will be what that person builds on, though it may differ from the initial narrative soup on which another newborn person builds. Whether the newborn, who we will call “Alex”, perceives distinct perceptual categories at birth is unknown and immaterial. A short time later, however, Alex will show signs of distinguishing one perceptual category from another, by, for example, crying or smiling, two categories of behaviour that adults distinguish as different, and requiring different actions if Alex’s behaviour disturbs a controlled perception. Most adults perceive these different actions as related to whether Alex perceives his or her state as requiring or not requiring some action from the local environment (which is likely to include Mother). If Mother acts, the result is an “event” that changes the state of Alex’s narrative soup.

Here we come to the same question of the likelihood of developing autocatalytic loops in the “soup” that will firm up into the “setting” of Alex’s narrative possibilities in later life. If Alex has developed N perceptual categories, s/he has approximately N^2 possible relations among them (exactly N^2 if we include the possibility of events that affect only one kind of perception). As things change in Alex’s Real Reality environment, very few of the potential events that relate perceptions of two categories Alex has already developed will ever occur so long as N remains small.

As Alex constructs more and more Perceptual Input Functions and a larger perceptual control hierarchy involving different control actions, the probability of Alex perceiving events involving perceptions of two particular categories increases until it becomes almost inevitable, and soon thereafter, some of the changes created by these events are likely to increase the probability that particular other category pairs will react to create a perceived event — catalyzing the creation of a narrative fragment in the soup of memory.

The same argument as was used by Kauffman (1995) applies. With enough variety of perceptual categories in the "soup", the pairwise relations between them will almost inevitably develop into at least one autocatalytic, and hence stable, loop. Furthermore, that loop will be creative, in Kauffman's sense of creating molecules, to which narrative events are analogous. An event that catalyzes a different event is often perceived as a "cause" of the subsequent event, in that the occurrence of an instance of the catalytic narrative fragment (the event) greatly increases the chance that an instance of the catalyzed event will occur.

For example, an event that might be described as "a key turning in a lock" when the perceived state of a door has been "locked" greatly increases the likelihood that a subsequent event would be "the door containing that lock will open". The key turning does not cause the door to open, but a narrative sequence of "key turning" → "door opening" is much more likely than before the "key turning" event, which, in itself may or not have been very likely to occur. Indeed, without a prior catalyzing event such as "key insertion into lock" or the existence of a perceived state of "key in lock", the "key turning" event is highly unlikely.

The existence of such chains of catalyzing events is very reminiscent of the precursor chains of molecule-producing reactions in Kauffman's chemical soup, from which the emergence of creative autocatalytic loops becomes almost inevitable when there is sufficient variety of components. Our "narrative components" initially are states observed by the baby Alex. Alex senses changes in the environment, which we argued in Chapter I.11 will lead to patterns whose variable frequency of occurrence affect the synaptic strengths relating patterns of nerve firings, enhancing some and weakening others. These changes in synaptic strengths are the beginning of reorganization of Alex's perceptual control hierarchy, in which Alex has yet to include purposive action categories other than those genetically "installed" such as crying, kicking, arm-waving, and so forth. Actions in these categories, though perhaps initially random, are likely to disturb Alex's mother in ways that reduce error in some variable Alex does control, such a hunger-level.

Most of the ever-changing sensory patterns Alex does perceive do not repeat, at least if they do, they are obscured by the great variety of non-repeating patterns of coincidences among the different sensor firing rates, even including as "sensors" those genetically constructed complex patterns of relationships such as detectors for moving edges in the visual field discovered by Hubel and Wiesel (1962). Some patterns, however, do repeat sufficiently often that their influence in strengthening certain synaptic connections outweighs the weakening influence of random changes in the sensed environment.

As sensory patterns repeat over time, Alex develops an increasing variety of patterns that induce different sets of strengthened and weakened synapses. Internally, these patterns represent what we call categories, sets of sensory patterns that have much in common, though never exactly repeating an earlier pattern exactly. These distinct categories are states in our "Analyst's view" of the growing narrative soup.

The more instances of a particular category Alex experiences, the more likely it is that Alex has also sensed an instance of that category having been involved in an event. If instances of the same two categories of perceptual states occur in another event, the event will enhance and weaken synaptic connections in similar ways among neurons whose synaptic connections often produce similar outputs. In other words, Alex will begin to construct categories of events as well as of perceptual states.

Just as with Kauffman's chemical elements and simple molecules, every event (analogous to a chemical reaction) produces changes in the narrative soup, and occasionally those changes may increase the likelihood of a particular category of event occurring — catalyzing it. Kauffman's mathematics apply, showing that when enough states that Alex perceives as of distinct kinds have developed into distinct perceptual categories, some will have participated in events (narrative fragments) that consistently catalyze events of a different category. After Alex has developed different categories for enough states (probably in the low hundreds of perceptual states) the catalyzing effects of the events involving those states are almost certain to have formed into an autocatalytic loop.

When narrative catalytic actions have formed into a creative autocatalytic loop, the sequence of events forming links in the loop will be perceived as what we might call Alex's basic truths about the world, although Alex would not. For Alex, all there is is a stable pattern in what we have called "the way the world works". Alex may or may not consciously experience any of this, but in the same way as occurs with chemical autocatalytic loops, the variety of new states generated by events is likely to grow exponentially, and along with it the complexity of autocatalytic structures of narrative — the inventiveness of the actions that Alex might use in discovering new ways to control perceptual states in the narrative soup.

The events that enter into loop creation may be passively observed, or may be the results of the maturing infant's random or purposeful (perception-controlling) activity. What matters is that there be a sufficient variety of possible event types. The formation of autocatalytic (or homeostatic) creative loops is very unlikely before there is a sufficient variety of possible perceptual event categories, and very likely once the variety passes a critical threshold (Figure II.2.3).

From that point, as the infant matures into a child and into an adult, the inherent creativity of an autocatalytic loop means that the "narrative soup" contains an ever-increasing number of potential event types (relationships in the "narrative setting" network), and the base network's initial autocatalytic loop becomes an ever more complex stable autocatalytic structure, that is likely to form higher-level loops of loops in a hierarchy of stable structural support (very much as Scarr (2014) claims is true of the physical tensegrity structure of a vertebrate body).

The setting for generating the stability of autocatalytic loops is the Real Reality that affect our sensory apparatus. But what of the imaginative reality of "planning" or of dreaming? Have we reached a level of catalytic loop stability that could be used in "Predictive Coding" or the Theory of Narrative Thought? Let us imagine planning and dreaming.

II.10.5 Dreams and Reality in the Narrative Soup

Here, from an article by Charlie Foran in the Globe and Mail newspaper 2022.02.04, is a reasonable description of a more mature "narrative soup", from which the components of events and their catalysts must be drawn.

> *[James Joyce and Virginia Woolf] were preoccupied with a then radical new project: how to capture the grinding interiority of the human mind. Humans have always been aware that we spend most of our time, and live most of our lives, inside our own heads – talking, arguing, disagreeing and commiserating with ourselves. Until Joyce and Woolf, writers hadn't figured how to render that incessant, often unflattering chatter on the page.*

All this internal chatter is of narrative fragments, some long, some short, but mostly of longish fragments that could sustain an argument or a complaint. These are the contents of a much matured narrative soup, evolved far from the initial components that might have reacted together to form a person's "first narrative", a tiny fragment that presumably is constructed because of an early event in the infant thinker's Perceived Reality (PR).

One important consideration about the chemical "soup" out of which life might have arisen was that elements would be more likely to meet and to react in a small volume, and particularly on surfaces over which they might move, than they would in an open ocean. Several authors (e.g. Patel et al (2015), Pearce et al. (2017)) have argued that life quite probably began in shallow tidal or warm rainwater pools. We will make the same kind of argument in respect to communicative narrative development and stability, treating family and the friendship networks that tend to build on social media as the small accommodating pools in which autocatalytic networks may easily develop. These somewhat isolated networks must create idiosyncratic narrative constructs (Volume IV). Some of those constructs create narrative "truths" of the world that are not truths to other people.

The Narrative setting takes a place analogous to the self-generated Perceptual Reality (PR) of the perceptual control hierarchy. The setting is the PR of the narrative at the time of any event in the narrative, of the environment in which the events of the narrative occur, including the perceived Laws of Nature and the entities in the perceptual environment. The Laws of Nature in a narrative need not be those of our perceptual control hierarchy, which are constrained by the Laws of Nature of the unknowable Real Reality (RR) in which we seem to live. The Narrative need have no similar ties to RR unless it is used to create perceptual effects on the sense organs of the individual or some other individual, such as an animal being hunted.

As an example of the non-necessity of consistency in the Laws of Nature of a narrative, I offer this fragment of one of my recent dreams. I was in a speeding train, and we passed a farmyard in which were two hens, one white, one black, neither pecking at the grain strewn under their feet. I wondered why they were not, and continued to watch them as the black hen wandered ever closer to the train.

In the dream I was concerned that this black hen might be killed if its wandering led it to the train, as the train was travelling very fast past the yard in which I was watching it. Clearly these observations, simultaneous in narrative time, of the train in which I was sitting rushing past the farmyard and my observation of the meandering of the black hen could not be consistent with the Laws of Nature in Real Reality, but they were consistent with the Laws of Nature in the setting at that moment of the dream narrative.

In my own waking imagination I find it difficult intentionally to imagine a setting with such internally inconsistent Laws of Nature. It is easy, however, for me to imagine narrative setting with Laws of Nature that deviate from those I have learned from everyday experience and from textbooks. Those appear to be perceived by other people very much as they are perceived by me. The Laws of Nature in my imagined or dreamed reality are mine, and unless I use communicative narrative to describe them, they are mine alone. In a waking imagined narrative, they usually are the Laws of a stable setting, and events in the narrative are lawful. In a dream, they are not.

A fiction writer is free to assume any Laws of Nature at all, be they stable or fluid, and to allow the reader to infer them only from the flow of events. Such a writer is likely to control a perception of wanting the reader to continue reading rather than at some point having to decide that understanding the writing was either too easy or too hard to be useful.

Most perceptible events in the setting involve the participation of more than one entity, typically an actor and an acted-upon when there are exactly two, corresponding to the subject and the object of a transitive verb in a communicative narrative expressed in English. If we again use the analogy to Kauffman's chemical soup and the network of potential reactions that might occur in it, we have in the setting a network of possible narrative fragments, each corresponding to an event. The narrative fragment "As she was walking…" describes a rather high-level event that extends over an undefined length of time, but its component parts imply something about the structure of the narrative setting that allowed this event to be possible.

An infant's first control loops formed in the womb, whether they were genetically determined or discovered from the experience of moving that the mother often feels to be "kicking". A healthy baby is born able to move its limbs, make sounds emitted from its mouth, and so forth. It can control perceptions it needs for finding and sucking on its mother's nipple, but not much else. Nevertheless, these minimal events, and repeating changes in the environment are what a life-experience must be built on.

I have used the word "event" many times in this section, but what *is* an event in the context of narrative, and of "narrative soup"? Is it necessarily something of which one must be conscious? Clearly not, because the event might be something like someone in another place turning on a light. Usually, that act would have little immediate effect on a distant person, whereas in this hypothetical case the actor became able to read a book, which was not possible before the light was turned on. On occasion, however, to start a light where there had been none might be an event that had widespread and long-lasting changes to the narrative of many lives.

One thinks of, say. Cornish smugglers avoiding punitive tax rates on imports, who might signal by a light to a ship offshore just where to land the goods, with luck out of the sight of the customs police. Most people along that coast would not perceive any happenings intrinsic to that event itself — the lighting of the fire or candle signal. the arrival of the boat from the ship, the offloading of the illegal produce (e.g. French Wine) and so forth. Nevertheless, the event changed the state of the environment in which they lived. Afterwards, trusted buyers could indulge in a bottle of the smuggled wine, which, before the event of starting the light, they could not.

These examples illustrate a "narrative event". A perceived event involves perceived changes in the state of at least one component of the perceiver's Perceived Reality (PR) environment, though the event itself is not necessarily perceived. An event is a dynamic concept, involving not states, but state changes. An instance of a perceptual category has, alone or together with a partner or perhaps two, changed the thinker's perception of some aspect of the narrative soup. As with a description of a control loop, the actual states that change are irrelevant, the "event motif" is that the changes occur and alter the possibilities and requirement for control action if those changes occur in perceptual values as well as in RR. Events in PR may act as disturbances to controlled perceptions at any level of the hierarchy, requiring action to reduce their current errors, if the errors are beyond tolerance limits.

As we noted when discussing the nature of a perceptual function, what the perceptual function describes (or defines) is a category, not an individual instance of the category. When an event happens, however, that event is itself an instance of a category, as is each perceptual participant in the event. A person crossing a road is an individual, as is the road in question, but a category of event that we might call "a crossing of a road" cannot be developed until many persons have been observed crossing many roads. Furthermore, "the same" kind of event cannot occur until the categories of persons and of roads have both been developed.

We must make a new point about our analogy to Kauffman's chemical development of reactions that produce catalysts for other reactions (Chapter II.2). When two atoms react to create a molecule, say a sodium atom with a chlorine atom to create a molecule of salt, the reaction would have had an exceedingly low probability of ever occurring in a multi-billion year time-span if there had been only one sodium and one chlorine atom in a wide and deep ocean. And after it had happened, there would remain only one salt molecule in the ocean. What makes the probability of making a salt molecule in the next second worth considering is if there are many sodium and many chlorine atoms in a small volume of fluid (say, water) within which the individual atoms could move and perhaps meet one of the other kind. The concentration of the reacting atoms is important. So it is with narrative events.

When a perceptual function is created, it defines a category, but to do so, many input patterns similar enough to induce the creation of one of Powers's "neural bundles" must have occurred in the patterns of excitation and inhibition of neural synapses in different neurons. The resulting category recognizer perceptual function would be one among many, were it not for the existence of anti-Hebbian learning and the consequent lateral inhibition (Chapter I.9, or Figure I.12.8, repeated here).

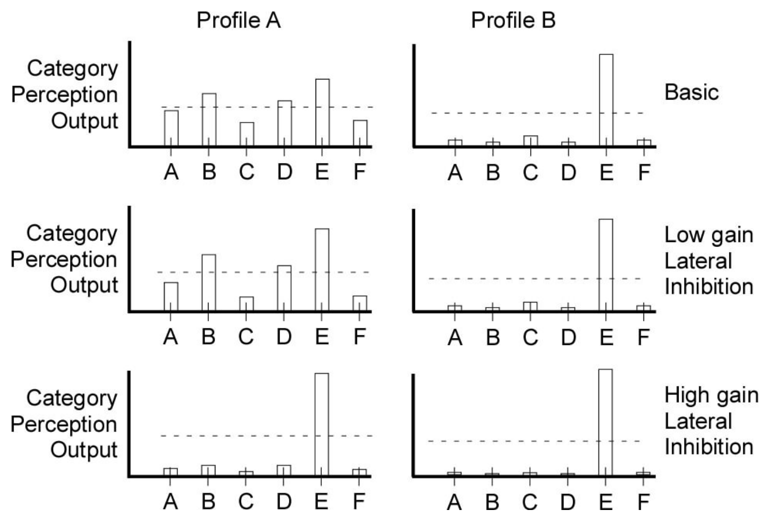

Figure II.10.x (Figure I.12.8 repeated). The effect of lateral inhibition on the outputs of neighbouring category recognizers. If category recognizer "E" is more often than its neighbours the one with the highest output, eventually the suppression of the others will tend to reduce the excitatory strength of their input synapses, and will eventually eradicate the others as effective category recognizers. Overall, this enhances the perceived category variability of Perceptual Reality, though not affecting that of Real Reality.

Lateral inhibition reduces the excitation of categories similar to the most excited one, and if there is one that is more often excited than its near neighbours, their suppression will enhance anti-Hebbian over Hebbian learning in their incoming excitatory synapses, over time "dissolving" them as effective category functions. In the "narrative soup" the effect analogous to a concentration of particular categories of reactions that produce potentially catalyzing molecules is the recurrence of similar events involving individual categories of perceptual states that have been isolated into discrete types by lateral inhibition.

II.10.6 Changing Components of the Narrative Soup

How can new components of a narrative soup be built? In Chapter I.9, we talked about how repeated patterns of sensory experience lead by the HaH (Hebbian-antiHebbian) process to new motifs of neural interactions. In Chapter I.11, we went up to the level of White Boxes as emulators of Black Boxes, starting with emulators of simple relationships among the output terminals of the Black Box that for living systems is Real Reality (RR). The "simple relationships" are firstly among the Black Box's output signals that affect the sensory inputs of the living individual, and then among the effects on the sensory inputs of the activities of the individual that affect something in RR. Consistent patterns of these relationships between action and sensation lead to the development by the HaH process of structures we call control loops, which incorporate novel functions we call perceptual input functions.

I would now argue that since each event does something to change at least one instance of at least one existing category into something else, or at least to give one of its properties a new value, events always introduce new elements into the soup. Some, or even most, of these "new elements" may be instances of existing categories, but occasionally something new may be produced by the event, something that does not seem to belong to an existing category.

Perceptual input functions — category defining function — produce signals we call perceptual values of instances of their types, which we call simply “perceptions”. These functions correspond to environmental complexes, though early in life the “complexes” are probably much simpler than some that become stabilized in later life, especially at the higher levels of the eleven Powers discovered in his intuitive self-analysis.

I am going to assert that every category function at every hierarchical level produces as its output a perceptual “state” of an instance of its type. This state is not in RR, but in PR. State changes occur because of something that happens in RR, but only in PR does a change have an influence on the values produced in the perceptual control hierarchy. At low hierarchical levels, that influence is more commonly in the value of an instance rather than in the category of what a perception represents in RR as detected by the senses. Only at higher levels does consciousness or imagination come into play as inputs to perceptual functions or, as discussed in II.10.6, as guides to reorganization of the hierarchy.

A category function incorporates a record of past experiences that have recurred in some form. In fact, any functioning control loop is, as a whole, such a record. It incorporates what actions have had what effects on what is perceived in certain contexts, something that none of its constituent parts do. The record is an emergent property of an evolved or reorganized control loop that controls consistently well. The enhanced rattling caused by what Predictive Coding would call “surprise” is an indicator that the local reorganization rate is likely to increase.

The existence of a well-functioning control loop indicates that the living control system has experienced a variety of situations, at least one of which has recurred as a type (such as controlling by crying for perceiving mother) relatively often. In Chapter I.11 we demonstrated the extreme improbability of any particular situation exactly recurring, though it may share enough properties with other situations to generate or augment a category function. If it did, and was remembered as having occurred before, there would be no need for the fine-tuning performed by every control loop. The agreement of its membership in a category rather than as a precise instance is what requires the fine-tuning of the instance property values. If Each instance had to be treated as something new, the cost in neural energy, with the accompanying problem of heat dissipation, would be extreme, as opposed to the existing fact of evolution that a healthy living organism can and does dissipate its heat of operation.

The brain-heating cost of adjusting the value of a type (logon information, Figure I.10.11) is presumably rather lower than the cost of determining how to act so as to bring the perception to be a precisely determined type (metron information), especially if the disturbance or reference value keeps changing at a rate comparable to the information rate available for control actions. In practice, the perceptual control hierarchy balances both, trading the number of clearly distinct perceptual functions against the precision of the value reported by the varying “neural current” associated with one bundle of neurons that contribute to the neural current.

This is all very well, but in what way does it imply that the “existence of a well-functioning control loop indicates that the living control system has experienced a variety of situations”. To function well, the action output of an Elementary Control Unit (ECU, the internal functioning component of a control loop) must be a close match to the RR disturbance that influences the PR correlate of the perception. The complex set of inputs to the Perceptual Function must convey enough information to distinguish what is happening “out there” and is to be acted upon from other similar possibilities that are not.

A newborn infant has very few, if any, controllable perceptions of any precision as compared with the repertoire available to the adult it will become. If a narrative is to be construed as an event that involves perceptible changes in the value or nature of some perception or perceptions, the narrator must have a sufficient variety of different perception to work with. These are the contents of the “narrative soup” that correspond to the elements of Kauffman’s chemical soup. Kauffman’s soup included rules based on physical laws that determined what molecular structures formed easily, needed to be built in a catalyzed reaction, or were very unstable if they formed at all.

The narrative soup requires a similar non-narrative background of what combinations of perceptual changes tended to happen easily, only in catalyzing circumstances, or only with difficulty survive quick decay as components of the soup. Furthermore, if the analogy with the chemical soup holds, the result of some events should be able to increase the likelihood of some entirely different events, in the way that the event of perceiving a change between a door being closed and the same door being open significantly increases the likelihood of an event in which a person is perceived as changing their location from one side to the other. If that event is also preceded by an event in which a key is inserted into the door-lock, the likelihood of the "passing through the portal" event is even more increased.

None of these apparently catalytic results of events could occur unless the necessary prior events had built appropriate perceptual control functions into the perceptual control hierarchy. Perceptual control of events such as "door opening" implies control of at least those perceptions that alter their value in the event. The construction of a particular event perceptual function implies that a particular class of perceptual changes, in this case from "aperture blocked" to "aperture open" at an early stage of crumpling, and from "my front door closed" to "my front door open" at a more refined stage of crumpling has happened often enough to build the synaptic connections into a perceptual function for the event describable in a narrative as "[I] opened my front door".

If the door in question was as frequently opened and closed by me as my own front door, the perception of the event might well have become a perception controlled in the non-conscious control hierarchy, but that might not be true of opening the front door of a newly rented apartment. The latter probably requires me to search for a place to insert the key provided by the renter, among several other points in which the action of the door might differ from that of my own front door, such as being hung on the left or the right. Under these circumstances, the category narrative component "unlocking, opening, and passing through this door" is almost certain to be performed consciously.

Consciously, one can imagine controlling the perceptions involved in the same narrative fragment, just as it is possible to imagine a woman walking along a well-trodden ancient path, either in general or in detail, clothing her, giving her hair a specific colour, making her be of a certain age, and so forth. All of this is involved with a narrative question "Why is she walking at that place and time?" "Why" is one of the three questions of a perceptual control function "What", "Why", and "How". A narrative event is its own direct answer to "What", as in "What is happening in this event". "How" also is often incorporated in the event narrative. In the case of that woman (identified only as "She"), "How" is "by walking", and as a narrative, it is subsumed under "What".

We have not considered a perceptual instance situation as a part of the inputs to a perceptual category function, and have not needed to do so yet. But any crumpling refinement of the category of a perception must add some more information about the category that distinguishes it from another of the same ancestry. That need for refinement depends on refining the situation, adding information about just what is being perceived. The more situation refinement is incorporated into the "narrative soup", the more discrimination is available among the possible and probable events, which are actions that affect the instances in the narrative fragment, and how they may be catalyzed.

Looking through time in the other direction, we come back to the newborn baby, who has very few experiences it can perceive as new event types from which narrative fragments could be built. With an insufficient variety of narrative components, the probability is very low that there would exist in the soup many catalysts for imagined events. Nothing perceived, and nothing that changes is particularly surprising, or, as a corollary, particular able to distinguish the likelihood of one possible future event from another.

Catalysis among narrative fragments is extremely unlikely until there are enough distinctive perceptual functions to permit changes among then to appear in characteristic, categorical, form. By using the same analysis as we did for the side-effects of controlling in Chapter II.2 to produce Figure II.2.3, it is likely that the critical number is in the low hundreds. Indeed, one might see a catalytic effect of one narrative fragment's production on the likelihood of another specific possible fragment as being an instance of side-effects that we will describe as types 2, 3, or 4 when we discuss in Chapter III.2 the ways the actions of one control system influence another's quality of control (Figure III.2.2). Type 2 enhances (or reduces) the speed or precision of the perception by the affected controller, Type 3 similarly affects the speed or precision of the effect of the output on the CEV, and Type 4 affects the influence of the disturbance.

Narrative thought, however, does not depend on events created by active perceptual control. Indeed, though among the first events in the soup for a newborn baby may be the effects of crying in bringing about a perception of mother's presence, almost all are the kind of repetitive sensory patterns and sequences over which baby's actions have no perceptible influence. We could call such sequences of perceptual changes "passive events" as opposed to "active events" that occur because of baby's control actions, whether or not baby can reproduce them (control the perceptions involved). If baby can intentionally change some perception in a desired way more often than by chance, we could go so far as to call the resulting event a "controlled event".

Each individual event offers the potential of adding its effects into the baby's "narrative soup". Individual events, however, are just that. If later, an adult would see another event as belonging to "the same" event class, the baby might not, because the situational context is almost certain to be different in detail. This is the same situation Powers described as a reason why a perceptual control hierarchy is likely to exist: that you can, for example, treat "going to the kitchen" as being the same every time you do it, even though your actual foot placement and body attitude and travel speed are never repeated exactly between times you create that "same" event of going to the kitchen.

In Kauffman's chemical soup, an individual reaction combines *this* "A" with *this* "B" to form *this* molecule of "AB". However, when Kauffman talks about an autocatalytic loop being formed, those individual A and B are no longer in the soup. A single AB is there in their place. What matters is the concentration in the soup of A's, B's, and AB's. The concentration determines the likelihood of "the same" reaction occurring. "The same" has to mean that a different A reacts with a different B to form another AB, changing the concentrations of all three types in the soup.

The possibility of saying that "this" A is the same as "that" A provides the necessary link between the chemical soup and the narrative soup. In the narrative soup, changes in perceptions cannot be perceived before perceptual functions have been constructed, which we are assuming happens by the growth and decay of neural synapses in the HaH process. The HaH process inherently involves events in which input neuron P fires just before or just after receiving neuron Q more often than would happen by chance, given the firing rates of the two neurons. Each "P just before Q" event tends to strengthen the connecting synapse, whereas each "P just after Q" weakens it, as to a lesser extend does a firing of Q with no firing of P nearby in time. Much repetition of the "P-Q" event distinguishes it, at least at that synapse, from "Q-P" events or "Q-alone" events.

Similar patterns at different synapses between incoming neurons and neuron "Q" enhance a whole collection of "X-Q" relationships while differentiating this pattern from a very large number of different patterns of firing involving Q. Neural Q becomes a pattern recognizer for that pattern of firings, and Q becomes a member of a group that have collections of synaptic inputs correlated with its own. Such a group is what Powers called a "neural bundle". What matters at the moment is that what we might call a "PQ" class of event is a member of a category distinct from "non-PQ" events, of which there will be many more.

For the purposes of analyzing the narrative soup, these temporal categorical patterns that form a "neural bundle" take the place of Kauffman's "A" and "B" classes of molecules. The neurons like Q that are in the bundle will all respond to the "P-Q" temporal pattern, but in other respects, that "P-Q" pattern occurs in different contexts of what and when other incoming nerves fire when the "in bundle" neuron fires. Each "P-Q" type of event becomes, over time, an "experience", and each nerve firing context implies the possibility of becoming a member of a different experience, the joining of these two experience now representing a "crumpling bend" that develops into a crease with the increasing precision of both the "P-Q" class pattern and of the precision of different contextual experiences.

Moving ahead a long way, a conscious experience has the same developmental pattern as the neural HaH produces. Repetitive experience sequences — event sequences, now — are likely to produce pattern recognizers for those sequences, which become refined or lost in time.

An event is of no consequence for perceptual control if does not produce a change in a controllable perception. At least one perceptual value produced by a perceptual function must change if it to be part of a narrative, internal or communicative. Here we come to a rather provocative point that up to this point in the book we have skated around, which is that nerves respond primarily to changes rather than to intensities, and the more abrupt the changes, the more strongly they respond, but, in contrast, the perceptions we consciously experience and we control are of the differences between the current values of perceptions and their references, not the changes in these critical values. To perceive the actual values cannot depend on neural firing rates as such.

Over time, then, the firing rate of a nerve varies in much the same way as lateral inhibition exaggerates firing rate changes across locations. A Narrative Event, then, is signalled by temporally correlated deviations of neuron firing rates across perceptual function outputs sensitive to properties of objects in the Perceptual Reality (PR) environment that were involved in the event. The consciously perceived values of the object must use some form of integration to smooth out these change effects, as must the comparison between the perceived value and the reference value in the comparator of a simple control loop.

An evolutionary rationale for emphasizing the changes rather than the state values is the same as it is for lateral inhibition, regularizing the use of energy in nerve firings that would otherwise heat the brain more strongly. In I.9.1, the rationale is put this way:

> *One reason why lateral inhibition is not merely plausible, but necessary is that the brain needs a way to keep its energy usage as low as is compatible with effective operation. Every nerve firing dissipates some energy in the form of heat, and this heat must be dissipated outside the brain that holds the nerve that fired. When the brain is as convoluted and tightly packed as is the human brain, dissipation of heat is a major problem.*

Consciously, it takes careful arrangements of changes between light and dark for us to perceive edge-change effects such as visual "Mach Bands" first described in a scholarly publication by Mach (1865, according to Pessoa, 1996), though they were apparently used by artists as early as the 15th century. Indeed some thousand-year-old Korean pottery arranges balanced changes between light and dark so that the viewer sees changes between light and dark areas even though the parts of the two areas away from the edges are objectively identical. The pattern uses the integrating function that normally compensates for the edge effects to create an illusory edge effect, allowing the viewer to see, for example, a bright moon in a dark sky although the pottery glaze is the same in flat areas of both moon and sky.

On the other hand, an evolutionary rationale for internalizing a representation of state values in some way is that there is a big difference between a tiger crouched nearby and a tiger moving into a crouch at some unknown distance. Your survival might depend the distance of the tiger and your ability to perceive its readiness to spring at you. Being able to perceive the changes in its location and in its attitude is important, but so is being able to perceive perceptible values that correspond to where and when those changes ceased. Powers apparently never resolved this issue in his mind, to judge from his actual simulations of on-screen tracking behaviour, in which he sometimes tried using velocity and/or acceleration as the lowest level, and sometimes as a level above the current value of the tracking error variable.

The brain heat dissipation rationale argues that the representation of a steady state should not be in a neural firing rate, at least not directly, or not if all the states available to conscious experience are simultaneously represented that way. However, neurons are bathed in an extracellular fluid full of a chemical "soup" that includes not only complex molecules but also elements such as sodium, potassium, and calcium among others. The local variation in the concentrations of significant components of this liquid "soup" depends on the recent firing history of a neuron. Furthermore, it affects the readiness of the neuron to fire again, the immediate determinant of its firing rate. The soup has memory in its chemical state, as opposed to in its changes of state.

Since I am more or less ignorant of the relevant physiology, and because this is a book more on the applications of PCT than on its mechanisms, I will not venture further into this complicated field of study. All that needs to be said is that there must exist a mechanism for storing stable perceptual states other than in neural firing rates, which respond to rapid changes in their inputs more than to the steady levels of the same inputs. My guess is that the mechanism for storing the stable or slowly varying value of a perceptual variable is likely to be found in the homeostatic properties of the fluid medium of the neurons (Chapter II.3). If this is true, then it forms an even more intimate link between the biochemical homeostatic structures and the perceptual control hierarchy than is described in Section II.8.2.

II.10.7 Narrative thought and Reorganization

In earlier discussions of reorganizing the perceptual control hierarchy, all the reorganizing processes amounted metaphorically to a tuning procedure, improving control of instances of categories that were built by changes in the synaptic connections among nerves. New categories were created from frequent recurrences of closely similar patterns of synaptic excitations and inhibition. Now we look at a different process based on *conscious* narrative thought rather than on the passive effects on synapses during ongoing perceptual controlling.

During reorganization, not only are the connection strengths among control loops changed, but so also are the very categories among which these connections are made. How can this be? Among the students of PCT, including Powers, it has long been acknowledged that with more experience, a living thing will encounter more new and different conditions, some of which it will not previously have encountered. Above, we argued that if it re-encounters sufficiently similar patterns of states often enough, a "bundle" of neurons will produce a state that we might consciously describe as "I've experience that before".

A "seen that before" may be the only label that particular pattern ever acquires in any particular language. Even between closely related languages some well understood patterns of language in one language may have no words in the other. English sometimes appropriates a foreign word, such as "gemütlichkeit". An English speaker might use the word as though they expected to have "the same" experience as a German speaker who used it. Indeed, a fluent German-English bilingual might be at a loss to translate it other than by listing a string of qualia that the experience includes and does not include, such as tension and nervousness. The nearest I can come in English to what I have been told it means by German speakers is perhaps "feeling snug, comfortable, and at home".

This example is just one of many I might have chosen from among the languages spoken around the coasts that face Britain across the sea. The point is that some people have an experience that seems in their inter-communication to refer to a common experience. They could be wrong, but if they are very wrong, at some point a misunderstanding on some communication might very well occur. We will not pursue here the correction of the misunderstanding, but instead note that if either party is unfamiliar with the word and the communicative syntactic, semantic, pragmatic, or situational context fails to clarify which components are missing or wrong (see, for example Section I.12.7 and especially Figure I.12.9) the error is most likely to be ascribed to a component of the conscious experience rather than to a malapropism.

Coming at the problem of integrating conscious and non-conscious modes of reorganization from a different angle, a "seen that before" is a collection of perceptual states that is consciously one that has recurred, for which at that moment one has none of the components of an ECU, namely a category recognizer/perceptual function, a reference value, a comparator, or an output function in the perceptual control hierarchy. Consciously, using rational thought, those components may exist and be used as they would be if they were embedded in the hierarchy.

In a narrative set in a novel, the reader might at some point perceive that the hero is in a difficult situation, and wonder how the novelist would get him out of it. One is the hero in one's own narrative, and might find oneself similarly situated unpleasantly — the conscious equivalent of a normally non-consciously controlled perception, already reorganized into the perceptual control hierarchy, deviating from its reference value. One might well then seek a rational plan of action, which might well take advantage of components that already exist in the already reorganized perceptual control hierarchy[43].

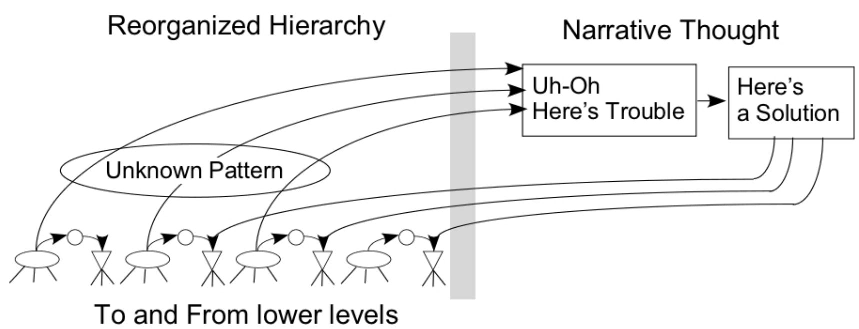

Figure II.10.10 A skeleton of an interface possibility between narrative thought and the already reorganized perceptual control hierarchy.

In the interface between the non-conscious reorganized hierarchy and conscious thought what Figure II.10.10 shows is a skeletal possible rational-thought analogue for the perceptual function (combined with the reference input and the comparator) and the distribution of the output to lower levels of a control hierarchy. In conversation, this is sometimes called "conscious control" because the analogy is so close. It is skeletal because the figure shows only the major highlights, the linked "bones" of conscious control connected to the more pliable "flesh" of the non-conscious hierarchy.

In the figure, "Here's a Solution" implies invention, because that solution is not one yet incorporated into the reorganized perceptual control hierarchy. Invention implies doing one thing so that you can do another. Considering the conscious or non-conscious control hierarchy as a way of altering the detailed contents of the narrative soup — or rather, as a variety of ways — "invention" must describe a sequence

43. The mathematician solving a novel problem will likewise seek to use theorems and lemmas generally accepted as proven in the new solution, rather than reprising the whole development of the field from basic axioms.

of “events”, narrative fragments.

At the very least, the invention suggests the possibility of refining perceptual categories into sub-categories, in that some instances of a category might not be similarly controlled in a context that otherwise is “the same”. If so, the instances for which this “solution” works will form a sub-category, different from that of the set of instances for which it will not work. What works, according to the inventor’s imagination, will determine the inventor’s conscious distinctions between the subcategories.

Figure II.10.10 might suggest that the previously unknown pattern and the solution connect with the conscious narrative thought all at the same level of the hierarchy, but this will be true only by happenstance. Figure II.10.11 might be a little more realistic, as it shows the problematic situation pattern and the actions that might resolve the problem linking at different levels with the existing already reorganized perceptual control hierarchy. Nevertheless, the box labelled “Uh-Oh Here’s Trouble” accepts inputs that define the problematic pattern just as would the perceptual input function of an ordinary non-conscious Elementary Control Unit (ECU).

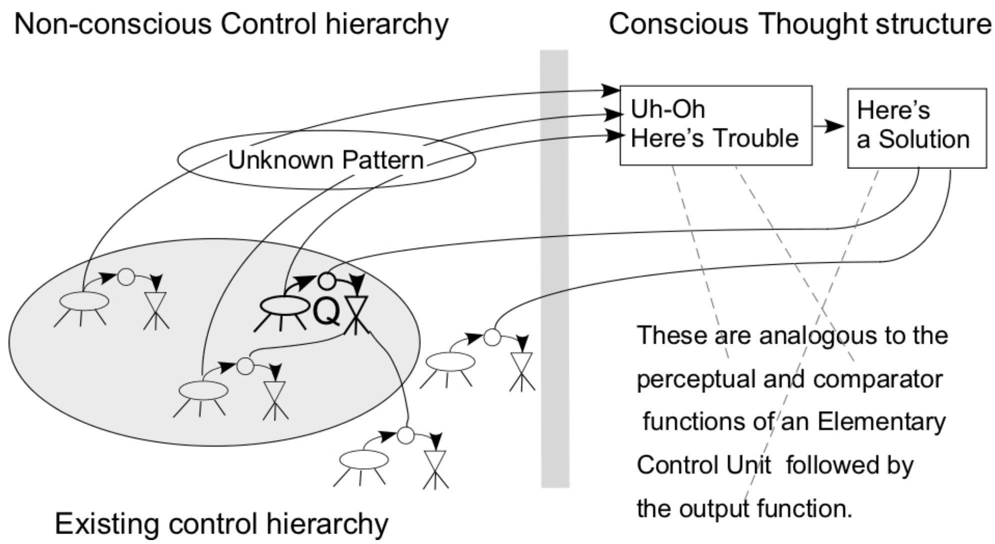

*Figure II.10.11 Figure II.10.10 modified to show that the pattern and the solution might each require connections from and to different parts of the already reorganized perceptual control hierarchy. The figure may be read from the shaded oval that includes a current pattern of perceptual kinds not (yet?) brought together as a known perceptual category. Conscious thought, however, determines that this pattern represents a problematic situation that demands action to resolve, and a probable resolution might be achieved by doing “this” then “that”, a sequence that can be controlled using already reorganized sequence controller “*Q*”, while another property represented by the controller outside the oval is being controlled to some rationally selected reference value that would be needed if the sequence were to work.*

Having identified the problem situation and having thought out a solution using existing properties of controllers within the non-conscious hierarchy in the form “I know how to do this bit and that bit, so I can solve the problem”, the “Here’s a Solution” box sends reference values to the appropriate ECUs in the existing hierarchy, and if the consciously imagined solution performs as anticipated when enacted, the problem is solved and the unknown perceptual pattern changes to a perhaps equally unknown but more benign perceptual pattern.

The reader may have noticed that the conscious thought part of this process is exactly that of the free-Energy-Predictive-Coding approach, except for the incorporation in Figure II.10.11 of a reference value at the top of the hierarchy. Where the Free Energy design incorporates a "free-floating" prediction of "what would happen if", we replace the prediction with the TNT (Beech and Wise, 2022) narrative approach of wanting and causing something to happen. In other words, we "cause an event". As always with so-called "causes", the result is occasionally not what is expected, either because of a faulty analysis of the situation or because of a disturbance to the consciously controlled "Unknown Pattern", which no controller in the reorganized hierarchy acts against.

Earlier, when we discussed building category recognizers, we relied on repetitions of patterns to create "bundles" of neurons in which the participating neurons developed by Hebbian-anti-Hebbian (HaH) similar synaptic excitatory and inhibitory input structures. We do the same now, but rather than relying on environmental variation to produce these bundles, we use conscious control to develop them.

Figure II.10.11 shows conscious thought using previously reorganized — "overlearned" — control units to perform a collection of tasks that would solve a problematic situation, a situation not previously encountered. That solution is a one-shot event, a narrative constructed of narrative fragments represented by the connections shown in the figure, and in words by this passage in the figure caption: "*doing "this" then "that", ... using already reorganized sequence controller "Q", while another property ... is being controlled to some rationally selected reference value*".

As a one-shot event, this pattern of perceptual states and corrective control actions will induce small Hebbian-anti-Hebbian changes in synaptic sensitivities among the neurons involved in the operations of the already reorganized control units, but unless a similar situation arises again from the environment or is imagined (perhaps in a dream) and is resolved the same way again, these synaptic changes are likely to dissipate in the noise of other influences on the same synapses.

On the other hand, however, the consciously thought-out problem recognition and solution defines a narrative fragment akin to a user-manual entry for what to do in such-and-such a situation. If the thinker later encounters, or imagines, a similar problem and uses a sufficiently similar solution again, remembering the earlier narrative fragment, the relevant synapses in the non-conscious hierarchy will be altered, again slightly, in the same direction as before. The entire pattern begins to define a category, together with its solution. "*Last time I searched for hidden Easter Eggs, I found one behind a candlestick. I'll look behind candlesticks this time, too*." "*Sometimes when a battery-powered machine won't start, the battery needs charging. I'll try that now*."

Every time a problem the thinker consciously perceives as "much the same" as one previously solved, and the same imagined narrative fragment works to solve it, that fragment becomes more like what is consciously called a "method" of addressing a class — category — of problem. Non-consciously, a set of synaptic connections begin to define an increasingly stable neural bundle or set of bundles that together specify an ECU (perceptual function→comparator→output function) that begins to become a rapid and non-conscious way of defining (categorizing) and controlling the perceived state of the world that when first encountered required conscious thought and invention to control.

In the narrative fragment instance contained in the last four paragraphs, a narrative fragment category is described as having constructed a corresponding event category in the course of its repeated enactment with multiple instances of the pattern types involved. This did not happen all at once, but over time, and the skill implied by the slow evolution of the new event category develops slowly, perhaps under the influence of a teacher. In the case of a manual skill such as hitting a golf ball or playing the right notes at the right time on a piano, the teacher may ask if the student can feel this or that muscular sensation in this or that order.

Without the teacher asking whether the student felt a particular muscle, the student might never have been conscious of an action controlling such a perception. Even after the teacher draws attention to the muscle, it might take training to become conscious of it, even though the relevant muscle tensions have been performed hundreds of times a day, in a variety of contexts. What is being taught is not how to create the muscle movements that the student probably learned to control while an infant, but how to sequence the well-learned movements in a new action pattern.

In Section I.5.6 we called such serial-parallel results *atenfels* and *molenfels* in the control of perceptions. The student has an existing repertoire of them for controlling a variety of perceptions, not yet including the new skill being taught. The new skill involves creating a new perceptual category by a conscious process that invokes the procedure of Figure II.10.11. This new process is highly likely to create a new high-level atenfel or molenfel that combines the control of existing perceptual types. For visual simplicity, Figure II.10.11 shows only two, but more would probably be invoked in problem solving by rational thought. Of how many of them might the teacher try to get the student to be simultaneously conscious when controlling for the student's skill to keep increasing?

II.10.8 Causes and Consequences in the Narrative Soup

In Section II.7.4 we pointed out that the basic neural mechanisms respond to rapid changes — events — rather than to steady states. Earlier we suggested that an event is to the soup of possible changes as a reaction in a chemical soup is to the atomic and molecular chemical components of that soup. In II.10.5 we argued that events in the Narrative Soup have intrinsic time-scales, and that events form a hierarchy, an event that takes a certain time to complete being formed of a sequence of smaller events that are quicker to complete.

One point worth repetition is that for an event to be perceived, it must involve at least one changed perceptible property in PR. However, not unusually in Real Reality (RR) the perceived property that changed will entail changes in other, unperceived, properties. For example, a lethal dose of radiation is completely unperceived by the victim, while the consequent event of failing health is clearly perceived, in the form of slowly evolving and changing components of the narrative soup available for perceptual control.

Perceptible events often appear to have no "cause". The natural reason for this is that the observer did not perceive a preceding event that changes the content of the narrative soup in a way that might lead to the observed even in the longer narrative. Thousands of years ago people could perceive the actions of other people and the changes they create in the probabilities of what might happen next, but other events might occur without the intervention of any apparent sentient purposeful actor, and were often attributed to the actions of the spirits of the material elements involved. These old beliefs are sometimes called a coherent "animism" religion.

For example, in a placid pond the probability of perceiving a splash some time in the next half second would ordinarily be extremely low, but after perceiving someone lobbing a pebble toward the pond, that probability suddenly changes to very high. The observer perceives the act of lobbing the pebble, and event, as having "caused" another event, the splash. If, on the other hand, the pond happened to lie over a geyser the observer had neither previously observed nor been told about, the splash of the geyser would suddenly occur with no prior perceptual changes that could alter the very low probability that such a thing might occur. The naïve observer might perceive it as caused by a capricious action of a hidden actor, such as a water-nymph, a naiad, or a deity such as Vulcan.

We might call such an unprecedented event a "miracle" caused by the presumed nymph or deity. Most of us usually don't, because we have become accustomed to discovering or being told about relationships, conditions and prior events in the natural world that could have caused the improbable event, and imagine that some as yet not understood natural structure or variation rather than an unobservable sentient entity contributed to the perceived event. Only Ockham's Razor (Working Paper W1 in Volume IV) could distinguish between the two possibilities.

Returning to the presumed ancient naiad-deity perception of the geyser event, such an interpretation could not be perceived unless the content of PR included such a purposeful, if invisible, living being capable of acting on the perceptible world. Whence might the construct of naiads, dryads, nereids, gods, goddesses, avatars of deities, and so forth arise? And how could it remain, in essence, stable over millennia before (in the minds of many people) dissipating over centuries as a structure, as it did beginning with the establishment of monotheistic religions in Mesopotamia and then with the Enlightenment in Europe? Why should every event have a visible cause?

The concept of sentient invisible actors inhabiting or controlling all aspects of nature that people experienced was stable in part because narrative events apparently caused by such beings supported each other as concepts in multiple levels of homeostatic loops.

In the chemical "soup" we have been using as an illustration, the reaction between two components of the soup produces something new, a creation. This is true whether the reaction is catalyzed or not. Even uncatalyzed, some reactions between some relatively concentrated elements proceed easily enough that their creations are also reasonably common and available not only as catalysts but also as reactors that produce creations that are even more complex. Even in the deep space among the stars, molecules as complex as some animo acids, the building blocks of DNA, can be found by observations from earthbound telescopes. Such molecules can be destroyed by being hit sufficiently strongly by photons, other atoms or molecules, but in principle could be built into even more complex molecules by less violent encounters.

Our problem is whether analogous effects can be found in the narrative soup environment. In that environment, two narrative elements come together to form a longer, more coherent narrative fragment. This longer fragment is a creation that describes a more complex pattern of changes in the environment, just as simpler patterns built from sensory events are collected together by a perceptual function into a higher-level perception. The event itself, as a description of the before-to-after change in the availability of the soup components (analogues to the chemical soup components that might participate in future reactions as reactors or catalysts), takes the place of the chemical reaction.

In the narrative event, a catalyst would be a unchanged state of PR in the presence of which the probability of that event happening is greater than if this state were otherwise. An example we have used elsewhere is the semantic and pragmatic nonsense sentence: "The cow ate glass", which would normally be a very low probability event, but if the narrating observer has established that there were shattered glass particles where the cow often ate, the event would become much more likely to occur than it ordinarily would. It takes the place of a chemical catalyst. If spoken to a veterinarian attending a sick cow, the sentence is by no means semantic or pragmatic nonsense.

There is, however, a big difference between a narrative catalyst that could be created by an event and a chemical catalyst that could be created by a reaction. Specific events cannot recur. As Heraclitus is said to have observed, "*No man ever steps in the same river twice, for it's not the same river and he's not the same man*."[44] Events of a similar pattern may occur frequently, just as atoms of the same kinds may often react to form more of the same kind of resulting molecule. If, however, a particular event happened, it happened and the world has moved on. It can never happen again.

No matter the truth of that oft-repeated saying, many men have and will step into many rivers, with a

44. https://www.brainyquote.com/quotes/heraclitus_107157

limited number of different kinds of results that depend on other situational results. The narrative might continue with the man dead, refreshed, pleasantly cooled, less thirsty, and a few other continuations that are sufficiently unlikely to cause surprise. For the man to fly across the surface of the river would indeed be surprising, unless you perceived that he carried a jet-pack, a perception of a state resulting from earlier events that probably would have needed to be deduced as a necessary narrative history if it had not been observed or imagined before the "across the river" event. Only these "surprising" events contribute significantly to the rattling of a perceiving entity sufficiently to appreciably enhance the likelihood of reorganization.

What this argues is that narrative events can recur in a largely conscious discrete category of the perceptual control structure, though they can occur exactly only once as an instance of a category perception in the analogue, largely non-conscious, part. In the category representation of the narrative — or other perception — the contextual details usually do not change sufficiently sufficiently to alter which category is perceived. For example, at a very high level of literary Narrative, the label "Romance Novel" describes many thousands of different books, and it takes many "crumpling" levels of refinement before a label identifies a unique book. The same is true of a narrative event at even a very high level, such as "boy meets girl".

Even if "no man ever steps into the same river twice", does this imply that narrative stability cannot be due to an autocatalytic loop structure? The answer is not obvious, so let us look a little closer at the chemical autocatalytic loop and the narrative analogy.

When describing the chemical autocatalytic loop, the actual soup contains many components, none of which are individually tagged as different from another of the same kind. For example, one oxygen atom or one carbon atom is, for the purposes of the loop, the same as another. If any oxygen atom meets any carbon atom they may react to form a carbon monoxide molecule. Only one carbon monoxide (CO) molecule is formed in this reaction, and what happens to it later depends entirely on its circumstances, such as how fast it moves through the soup, what other specific atoms or molecules it meets, whether it is knocked apart by a stray high-energy photon, and so forth. That CO molecule exists until it dissociates into the same individual carbon atom and oxygen atom that came together in the original reaction, or until it participates in a more complex reaction.

None of the three histories of that particular carbon atom, that particular oxygen atom, or the CO molecule that binds part of their two narratives together for a while between the reaction event and the dissociation event, is of any interest when we look at whether an autocatalytic loop is likely to form. All that is interesting is that to label an atom as being an atom of carbon is to lose the individuality of the atom as a member of a category that has been labelled "carbon". Any time any member of the "carbon" category meets and reacts with a member of the "oxygen" category, a member of the "carbon monoxide" category joins the count of how many CO molecules exist in the soup.

The point is that individual narratives, whether of atoms or of a boy and a girl in a novel, are of interest only insofar as the individuals are members of their categories. In the chemical soup, what matters is the concentration of members of a category, such as how many "oxygen" atoms are in a litre of gas, or the concentrations of boys and girls in the population. If the "boy" and "girl" interact to form a "couple", for the analysis the interactions would be just the same as of the "oxygen" and "carbon" in the chemical soup.

That it does not feel the same depends on the fact that this boy and this girl are identified. They have individual personalities and necessarily so if they are characters in a novel, if the novel is to describe a long narrative. The novel itself is a category, in total it creates in the imagination of the reader a long-drawn out event that consists of a hierarchy of shorter and shorter events, the "reaction products" of which change the probabilities of other classes of events, reactions in the narrative soup by changing the concentrations of possible catalysts or reaction participants.

Along with all this, in the background of all this is the effect of a changing physical background. In the chemical soup, the reacting components and their created products might be distributed thinly through a broad and deep cold ocean, or in more concentrated form in a small shallow warm pool where they can't move far away. In the Romance Novel, the boy and girl might live in a grand mansion, among a small tribe of hunter-gatherers in a forest, in ancient Rome, or in some other context that remains stable for a long time compared to the shorter narrative events in the novel, or over the entire novel in many cases.

The background matters. A romance Novel featuring a rich girl and a poor boy, or vice-versa, is of a different category (after only a few crumpling episodes) than one featuring two people living in the same background, whether it be opulent and imperial or impoverished and filthy. Cinderella stories would be surprising in real life, but as the basis of a novel, they are not. Too many stories worldwide feature the same "reaction" between rich boy and poor girl. A growing child can hear them for the first time and perhaps reorganize because of the resulting rattling, but for most adults they are quite familiar.

For stability, the story background changes either slowly or, if suddenly, only very occasionally. The information-theoretic uncertainty of the background remains low, the same criterion that allows control loops to function effectively. The so-called "information generation rate" (rate of increase in uncertainty) of the background is low enough at each level of the hierarchy (Chapter I.10, especially Section I.10.4) to allow the tensegrity effect to work, different specific perceptual control loops supporting ones that interact through mutual side-effects.

The same must be true for a coherent narrative event hierarchy. In the perceptual control hierarchy, coherence is in the control of the highest perceptual level currently being controlled, which often in this book we have taken to be one of the two self-image types. People who have what is called "split personality" presumably have more than one control system based on different reference self-images. Their behaviour lacks coherence when they switch from controlling one self-image to controlling another. Their clinical (externally observable) symptoms differ from one physical body to another, according to the Wikipedia article on Dissociative Identity Disorder[45]. However, apparently they do have in common that one personality behaves sufficiently differently from another that to call them distinct seems justified to the observer.

In a living control system, a major symptom of Dissociative Identity Disorder (DID) is mutual amnesia between the personalities. Their internally perceived narratives are discontinuous, as reported by them to someone else. Everyone, or nearly everyone, forgets some things, and memory researchers find that people recreate the narrative of their own history to incorporate events that did not happen. The researchers can create memories that their subjects appear to believe are of historical events[46] (***REFs***).

I introduce DID here to illustrate how the perceived events in a person's own history tie into perceptions controlled in their own control hierarchy. A person's narrative may not be stable as seen by someone who reads a person's life-long daily diary, but the person controls the recall of perceptual states so that they remain consistent with, and are part of, their historical narrative at the "now" moment of ongoing events, however short or long the "now" moment — the duration of the current relevant event — might be (Section I.7.2).

45. <https://en.wikipedia.org/wiki/Dissociative_identity_disorder> retrieved 2022.01.24

46. From personal experience, quite often when I talk with my brother about an event I believe we shared, he says that it never happened or it happened in a way very different from my own narrative of the sub-events involved. Which of us, if either, has a narrative that conforms to our shared history is usually impossible to determine.

II.10.9 Reorganization by Conscious Thought

Returning to the basic event structures that might culminate in narratives that form creative autocatalytic loops, we are starting to make fundamental discriminations among kinds of narrative events. We are tying them to types of object in PR, perhaps not changing the underlying "reality" of the narrative world, but changing how likely one event rather than another comes next as a consequence of an earlier event, as in the example of meeting a dog that might be docile or aggressive. The narrative soup might include fragments with previous encounters with either this particular dog or with others of its breed. These perceived types of dog require different actions to control perceptions that might be disturbed by the dog's actions.

In our use of Kauffman's probabilistic analysis procedure, if we acquire a sufficient variety of soup components, at least one stable loop is likely to be formed. How might such a variety be initially produced? It does not occur in the syntactic variety of most linguistic analysis of a human language, let alone of non-human language insofar as we humans have been able to analyze it. We should therefore expect the requisite variety to be in the experienced in the environment perceived during an infant-child-adult maturation process.

Keeping in mind the one-way flow of time, it is clear that no loop will contain specific examples of a narrative fragment. the loop members must be types, categories, of narrative fragment. It cannot contain "Rob meets Lucy" but it can contain "Rob meets a girl" and "A boy meets a girl". Frequent encounters with instances of the types and their characteristic differences in observable behaviour generate the crumpling splits between the subcategories of the larger category "person". These categories are likely to have been reorganized into the non-conscious perceptual control hierarchy.

A newborn infant of any mammal species has never experienced any events in the external environment, except what it experiences while inside its mother. Those pre-birth experiences include sounds, and tactile sensations including pressures of the mother's breathing, blood flow, and muscular actions, some tastes but probably not smells, and not much else. Changes in any or all of these are events from the Analyst's viewpoint, and might possibly be perceived in the fetus's growing brain. Whether they are or not, the fetus has little or no ability to act so as to influence them.

What it might be able to control are kinaesthetic perceptions such as those involved in kicking, and the tactile perceptions that change as it kicks. A human mother also feels the effects of kicks, and if the father is nearby, may say something to the father and ask him to feel the kicks through her stomach wall. The sequence of kick, perhaps sound of talking, and changing pressure of hands on the stomach form an event or a series of events. The fetus is born with the possibility of perceiving events in the form of changing sensory input sequences. As we will argue in Chapter II.11, this ability lies at the basis of a baby's first language, a language co-invented by mother and baby together in a two-person form of control loop.

At this very early time of independent life, the newborn has a very extensive and diffuse network of neural connections that it will prune depending on how the effects it produces on the external and internal environments are consistently changed by its actions. This pruning, along with new synaptic development, enhances the consistency of the effects some actions have on changes in its sensory inputs. It develops perceptual functions, some of which it was born with, some that it develops according to a genetic program inherited from its parents (and they from theirs). All of its actions alter the perceived state of its environment, and are therefore events perceptible in principle, though perhaps not in practice.

However it comes about, according to PCT the infant develops a hierarchy of perceptual types, each of which has perceptual functions that distinguish particular perceptions of different kinds with individual relationships to the outside world by way of a succession of simpler perceptual functions. Some types often considered as levels in the hierarchy may consist of sub-levels. For example, relationships may be of relationships, as in A-B>C-D at the same time as A>B and C<D, all of which may simultaneously be perceived. Event categories, likewise, may consist of event types that consistently follow one another. For example, heavy rain is consistently followed by the wearing of waterproof footwear by pedestrians, and often also followed by a period of sunshine. Events that follow one another constitute a basic form of narrative.

As the infant matures, it experiences an ever-growing repertoire of events both in its body and in the external environment, most of which do not include events created by its own actions or that the infant fails to perceive as having been caused by its actions. We talk, of course, of consciously perceived relationships between its own actions and changes in its perceptions. Non-consciously, some synapses strengthen while others reduce in strength through Hebbian-anti-Hebbian (HaH) learning (Chapter I.9), generating ever-improving perceptual control in the non-conscious hierarchy.

Only when the maturing child has developed a sufficient variety of perceptual functions whose perceptual values can be consistently changed by its actions can there be a sufficient variety of perceptible events in the "narrative soup" to create even short narrative chains in which one event produces a state that renders likely (not definitively "causes") a new state. As with the Kauffman (1995) chemical soup, when enough ephemeral chains get produced and stay long enough before they decay away, they will almost certainly form into narrative loops, autocatalytic creative loops, just like any other self-referential structure for which the self-reference returns through identifiably distinct stages[47].

It may be clear that if the perceptions are categorical, such a condition is actually a control tensegrity structure in which the effects distribute themselves through the structure. In Chapter I.8, we described tensegrity structures in the analogue hierarchy. Here, we see them emerging naturally from narratives involving events in the categorical, conscious, structure developed (metaphorically) by crumpling. The point in each case is that an appropriate set of "wires and rods" are connected appropriately.

In the narrative, the events take the place of wires that bring two states together to be perceived as a new complex state (as in an upper level of a two-level control hierarchy) whereas the changed states form the rods that maintain their distinct differences. If those "rod-like" differences did not exist, no event would have occurred. The newly created complex state, at the same time exists in the soup in a way that may catalyze another event or may be a state changed by the next event in the narrative.

II.10.10 The Narrative Fragment Hierarchy

Let's start by repeating a basic statement: narratives are not composed of words, except as needed for communication. A narrative may be communicated in words from one individual to another, along with gestures and acoustic modulations of the way the words are spoken, if they are spoken rather than being written. That communicative narrative structure is at best an approximate translation of an internal narrative structure within the transmitting individual.

So what is an internal narrative composed of? A narrative is a series of events perceived to have occurred in a Universe that constitutes a consciously imagined version of Perceptual Reality (PR). That consciously imagined PR might be deliberately fictitious, or the imaginer might believed it to correspond to Real Reality (RR). We will discuss the political implications of this latter possibility in Volume IV. Here we need note only that for our purposes in the Chapter, the two possibilities can be treated as essentially the same.

Using an analogy to Kauffman's chemical "soup", the imagined dynamic PR in which the narrative

47. We may suggest at this point that these creative loops constitute the re-creative aspect of memory for events, in their conscious form as interlocking narrative fragments that produce a remembered story.

fragments reside is made from what we might call a “soup” of interconnected perceptual constructs of different complexity. A chemical reaction in a chemical soup alters the availability, the concentration in the soup, of the components involved in the reaction. One unit less of each reacting component and one more unit of the reaction product exist in the soup after the reaction than existed before. The change in the concentrations of the components is an “event”.

In any abstract version of a “soup” of constituent components, of which we will claim a Narrative soup is an example, events are sometimes treated as describing the analogue of a reaction and the changes in the soup as the result of the reaction and sometimes as the reaction itself, as in “this event had that result”, analogous to the chemical “this reaction produced a molecule of that compound”. Both reconfigure the soup in which they lived, if only slightly if the soup initially contained many copies of the components, of which one was used in the reaction.

In a narrative, it is harder to say that a component has been depleted by being used, but from an information-theoretic viewpoint, it is. The reason for the depletion is that little or no extra information about a given subject is transmitted by repetition. What may be transmitted by repetition in a suitable context is information about the differing relationships of the various repeated elements, including how important the creator thinks the explicit facts are. This is as true of interior narrative as it is of communicative narrative.

In the chemical soup, molecules created in a reaction may be destroyed in interactions with other soup components, just as a brick wall may be destroyed by being hit with a fast-moving projectile. Molecules decay over time at rates that are inversely related to what we call their “stability”, which depends on how much energy is needed for the molecules to overtop an energy barrier and dissipate into their constituent components (atoms or stable sub-parts of the more complex molecule). Likewise information created in an event may be lost by interaction with the results of other events. We call this loss by decay, “forgetting”, elimination from the “soup” a component molecule, a narrative fragment that is the instantiation of an event, that originally changed specific components of the soup. It corresponds to the entropy increase that is almost inevitable in an isolated system.

Events are changes in the contents of the soup. In Narrative, the contents of the soup are the states and relationships among the contents of Perceived Reality. In a narrative, some perceived states may come together to create an event that changes some other perceived state or relationships. For example, one might perceive oneself to be sitting in a stopped car facing a red traffic light, and imagining that the light will turn green, at which moment the perceived dangerousness of proceeding to drive into the intersection will be sharply reduced, just as though a dog perceived to be “dangerous” might be perceived as less dangerous if its owner were to say “Sit” and the dog obeyed.

A perceived event affects the contents of the perceived reality (PR), whether in imagination or the external environment. This is exactly what a reaction event does in Kauffman’s chemical ”soup”. It creates a new component in the soup, as a perceived event alters the possibilities and relative likelihoods of specific kinds of new events in a narrative soup. An event, then, is anything that changes the consciously perceived state of PR. This would be just as true of a fictional world in which the characters in a novel act as if the world of their Perceptual Reality maps onto the Real Reality in which they live just as our individual Perceptual Realities map onto the Real Reality in which we live.

In this sense, among words that communicate a narrative through effects on the senses of another individual (word sequences, in human languages) verbs are not the only kind that change the environmental soup. Consider the probable different narrative consequences between “… met a dog …” and “… met a big growling dog …”. The adjective changes the probabilities of the possible following narrative structures, between, say, the protagonist stroking a friendly dog and protecting against the apparently dangerous dog that may bite.

The adjective signals an event in PR, a change between the perception of a generic kind of dog and a specific category of dog that we might call "threatening". In Section II.6.4ff, we saw such a class fragmentation as a crumpling event that splits a facet into two, and the object class likewise split into two. The distinction across a facet boundary is the value range of a perception produced by a perceptual function that determines to which fragment a current observation belongs.

In the "dog" case, one might ask "how dangerous", a type of question that indicates a fuzzy "bend" boundary exists between "cuddly" and "lethally fierce", a boundary that turns into a crease if the difference between the facet fragments determines how to act — to pet and stroke the dog or to protect against a possible attack. This development of a bend into a crease is called "task stress", the perceived importance of the difference between the likely effects of two possible actions (Figure I.9.14). Such a possible sharp or fuzzy value-determined distinction is signalled in at least the English language by an adjective or an adverb, a modulator.

An event is, by definition, any change in the state of the Real Reality Universe, but perceptible events are very much more restricted. Perceived events are changes in PR that occur over time, between past and future, an interval called "now" that is longer the higher in the hierarchy are the perceptual states changed by the event. This includes reorganization events that, while not themselves perceptible, alter perceptual states in PR, creating new kinds of events that might be perceived if they were to occur. Since, according to PCT, controllable perceptions at any level not directly connected to the sensors (i.e. all perceptions) are based on simpler perceptions more quickly controllable, we must accept that events in a narrative also have a corresponding hierarchy of complexity and time-scale of change.

Individual events in this event hierarchy are conceptually linked individually to the perceptions whose values they change. The implication is that event types exist in a completely parallel hierarchy to the perceptual control hierarchy, exactly as does a category hierarchy of perceptual states (Chapter I.11), or so we have asserted. The event hierarchy then will likewise exist in both analogue and discrete forms. Categorical forms are open to linguistic description using a finite uncertainty space (Chapter I.10). Analogue forms can be described linguistically only when the uncertainty space is artificially divided into distinct categorical macro- and micro-states.

The essential point here is that an internal narrative may incorporate changes in the consciously perceived continuum values of analogue perceptual variables, whereas a communicative narrative incorporates events that alter the perceptions of identifiable categorical labels — facets of separable objects in PR. Changes from one facet of a category to another implies changes in the class of an object (or an Object or a White Box). Analogue values, even those of perceptions controlled in the reorganized hierarchy, normally non-consciously, can usually be made conscious and available to internal narrative thought by approximating them to some discrete category value.

We have now linked individual events to changes in perceptual states, whether those changes were initiated by changes in the corresponding reference values or by changes in disturbance values. These linkages imply that the duration of an event corresponds to the rate at which actions on the environment (including the internal environment) control the relevant perceptions. These durations might be very different among the interacting changed perceptual states, whether they are controlled or not.

In our example of crossing an intersection when the traffic light changed from red to green, the change is essentially instantaneous, but the control actions that result in the driver perceiving the car to be on the "stop-go" side of the intersection are far from instantaneous. The event duration is determined by the slowest component, the fastest in this case being only a modulator of the relationships among the perceived states of the environment, including in this example, the velocities of the vehicles using the road one might want to cross.

What else is happening at the same time that also depends on the colour of this traffic light? Someone may be teaching a child how to interpret the colours so that it should be safe to cross if you see green, but not if you see red (or a don't walk" sign). In the parallel but interacting narrative of the pedestrian, there are prior events experienced and future events expected, both of which correspond to possible conscious experiences, the whole forming a fragment of a larger narrative fragment such as "I walked to the store". The fragmentary narrative components of before and after, together with parallel narratives such as that of the driver who stopped at the red light thus enabling the pedestrian to continue the walk to the store combine into the overall "walking to the store" narrative. Just as in the perceptual control hierarchy, a narrative hierarchy builds larger scale narratives on shorter components.

What we have not yet done is to treat the narrative structure on its own terms. We do so now, recognizing the series-parallel nature of a narrative such as the above "walk to the store" that extends over time. A narrative consists of a linear sequence of events, each of which results in changes in the possibilities contained in the relationships of states of the contents of Perceptual Reality, changes in what events could happen next with reasonable probability.

For example, if a person wanted to take an elevator and happens to be facing a nearby open elevator door, the event we might call "taking the elevator" is rather likely. If, however, before the person reaches the elevator door it closes and the number displayed above it changes, then the "taking the elevator" event incorporating that person is unlikely to occur until after some waiting time has elapsed.

A complex narrative does not consist of a single event, because that event consists of several shorter events, each of which represents something changing in the soup of Perceptual Reality. One event changes one or more perceived states of the external or imagined environment, or both. Those changes alter the relative likelihoods of other possible events, one or more of which may be perceived to occur. Events exist in three possible time states, some have happened, some are happening, and some are yet to happen. Imagine the event that is happening as a tick of a clock, which distinguishes events whose "now periods are before and after the tick.

When the event is a literal tick heard from a physical clock, we say that the ticks indicate the passage of time. Why do we perceive it to be so? Largely because a flow of time past one is what it feels like, but also because it is possible to count the ticks between, say, when start to pick up a water glass and when we put it down after the event of having taken a drink. There is a "now" moment while we are drinking, a moment that may contain other events, such as the "now" event moment during which we are taking the first sip or gulp, the "now" event moment while we are taking the second gulp, and so on. Each "now" event can be split into a sequence of shorter events. Likewise, each "now" event is likely to be a sub-event of a larger narrative, such as a "now" of eating a meal.

Each event in the tree of ever faster event sequences packaged under the "now" event of eating the meal exists in time, either before or after the 21,653rd clock tick since sunrise, but it is unlikely that the count of ticks is perceived other than fuzzily, such as might be expressed by the thought "It's getting late". Without having kept track of the ticks (or consulting a physical device that keeps count automatically, a device we call a clock) why would we have a perception of "getting late"? We do have a clock, several of them, in our physiology, though we perceive only effects such as getting hungry or sleepy. We can keep time on much shorter scales, as in the rhythms of music, for example.

Music is a useful example of how the narrative structure of events does not need to depend on a story that can be put into words, though some Western music does exactly that. Western music of the last few hundred years can be suggested by the equivalent of written words, a score that describes how the timings of parallel events can be related to each other, and often (in the form of bar lines) to a mechanical tick-tock-clock time that novice pianists might hear from a metronome. The pianist or conductor/leader has to make sure that the acoustic events notated in the score start and stop in the order that is shown in the score. In the language we have been using, each event, a long-held note or a staccato blip, defines its own "now", as does any event in the music formed when these notes are heard consecutively or together while the music is performed or imagined.

A tune is a long event in music that is composed of a sequence of shorter events such as phrases, each containing yet another level of shorter events, the sequence of successive notes, some of which are played in parallel, in the form of chords. Just as with any narrative, each note changes the probabilities that the next one in the sequence will be a particular one. After a few notes, in most European-based so-called "classical" music, a set of probabilities called a "key" will be established. The tune being in that key is an event with usually a slow start and a slow transition into another key, perceptually though perhaps not in the mind of the composer. Within that key, the tune may contain unexpected events both in performance (wrong notes) and in the score (moments that the composer wants to mark by their unexpectedness in context).

In much such music either there will be more than one tune-event occurring at the same time as another in a relationship called "counterpoint". On a shorter time-scale, more than one note may occur at the same time as another, in a relationship we call "harmony". Composers may emphasize one kind of temporal relationship over the other, but always, what is happening "now" at any level of "now" alters what may happen next, and depends on what happened previously.

The same kind of within "now" relationships across separate but simultaneous strands of events have combined effects on the world they share in the same tick-tock time. Their internal event durations, their "nows" at the different levels of "now" are related only insofar as their effects on their common "narrative soups" are independent.

A crude way of distinguishing classes of events is by the classes of words used to describe them. Verbs in English, for example, may have one, two, or three states associated with the event the verb describes. For example, "the snow fell" describes an event in which one entity changed, the snow, though the narrative thought implies at least one other also changed, the ground on which the snow fell. "The cow ate grass" involves two changed entities directly, the actor cow and the eaten grass, while "John gave Jane the book" involves three. It seems reasonable at a minimum to allocate events into classes based on the number of static structures (states, entities, objects, etc.) implicated as being directly changed in the event.

We are not really talking about the varieties of language structures; we are just using them as possibly useful guides onto classes of processes in narrative thought that change the contents of the soup in different ways. If you think of language as one way of allowing another person to perceive a piece of your ongoing narrative, it seems reasonable that the forms of grammar that analysts over the centuries have applied to language might have some relationship to the kinds of narrative thought relationships imagined by the speaker/writer who attempts to communicate them to another.

Any change in the narrative soup implies other possibilities for change. If John did give the book away, before the event John could have read it, whereas after the event, he could not. The book might have been safe from damage when owned by John, but perhaps from previous narrative threads in our imaginative mind our thoughts include that Jane is not very careful with her books, so in possible narratives of the future, a variety of possible events may befall the book. The "now" of the transfer probably is not very long, but the event stream that follows would probably define a narrative very different from any event stream that would have followed if John had kept the book or given it to Daphne.

II.10.11 The Formal Grammar of a Language

At school, a child may be taught "English Grammar" (or the Grammar of French, Russian, Chinese, or whatever the school curriculum demands). If this is their native language, the child already speaks it fluently and with conscious thought only to the meaning to the recipient of what is spoken. What, then, is actually being taught, and why?

One might ask the same question about what is taught and why, of any topic that is not requested by the child. One likely underlying answer from a teacher or administrative creator of a curriculum nearly always is the same "*In my opinion, it is a tool they will need when they are older if they are to think as I would like them to do*", though who is the "me" who holds that opinion is undefined, and the answer is often framed very differently by a Department of Education supervisor as compared to a hands-on teacher. All the teacher's actions are to control errors in their own controlled perceptions, some of which may be perceptions of the knowledge state of the students. The teacher may be controlling for following instructions contained in a syllabus produced by higher authorities for purposes supporting perceptions controlled by someone else (See Section IV.5.3 for some relevant material).

This Section, however, is not a perceptual control analysis of education, but an analysis of a possible linkage between conscious and non-conscious perceptual control in the use of a language. In the next Chapter, which is entitled "The Private Language of Mother and Baby" we will look at the development of a language between a mother and her newborn baby long before baby learns any sound strings we might liken to words. For now, we will presume that this baby has grown much older, and can both understand and express much in ways that unrelated children and adults can use to interact with him or her. The child is old enough to go to school and be taught the "proper" use of the language they have been using happily for some years without having had to be taught how to do it.

What is the point of teaching or learning a "grammar" of the language one already speaks fluently? Indeed, if one does speak "correctly", one may be perceived as stilted, speaking that way only to create some effect.. After all, this does not happen to children who speak any of most of the world's languages, few of which have had their formal grammars studied, and these children seem to get on quite well without it. However, most languages of the world are spoken only by small numbers of people who reside in areas small enough that those speakers regularly contact each other through only a few intermediate speakers of the same language. This does not mean that they are geographically close, since trade languages called pidgins and creoles, mixes from different place-based languages, can span oceans.

Formal grammars that can be found in textbooks tend to be those spoken by many people over a wide range of territory. Chinese, Russian, and most European "National" languages are of this kind. The "correct" version of the language is spoken in politically important parts of the country, and different versions spoken elsewhere within the country are called "dialects". In Chinese, for example, the "correct" language is Mandarin, and other, perhaps mutually unintelligible, versions are "dialects".

When I was young in England, "correct" English was called "The King's English". Originally, "The King's English" was the English spoken in the court of George II, and many of its pronunciations were varied from standard English to match those of the King, who had been brought up speaking German. "Correct" English a few decades ago maintained those originally Germanic pronunciations, which were not used by most of the "lower classes", who considered them affected and "hoity-toity".

At least one rationale for teaching the formal grammar of a language is to enable the dominant people capable of controlling more kinds of perception (the "bosses") to communicate with the submissive people they command throughout the political region over which they assert control (see Volume IV, especially Section IV.5.3, for most of our discussion on such political topics).

This approach treats a Formal Grammar as "prescriptive", laying down a law that a child should learn to follow when speaking or writing for others to understand. In the other direction, an analyst of the grammar of a language "as she is spoke" is constructing a "descriptive" grammar, in which there are no wrong ways of speaking or writing, but there are many wrong ways of describing how people actually do speak and write. Many of those wrong ways find their ways into teaching material, especially if they are not very wrong.

The “Grammar of English on Mathematical Principles”, by Zelig Harris (1982), is neither prescriptive nor descriptive at its core. In one sense it could be described as a narrative analogue of the development of the event structure that permits communications in the form of the General Protocol Grammar to proceed, with English as an example of the method. The “Mathematical Principles” are those we have described in the development of event categories within the perceptual control “soup”.

Harris himself writes from within a self-contained domain that consists of the language being described, but we take a view on his work from outside the language. Where Harris starts with the concept of a sentence as a unit of communication, we also want to know where the sentence comes from, treating it as a unit of meaning rather than as an actual word string. This may seem an unwarranted departure from what Harris says about his own work, but I do not believe it is. Whether or not it is, we can at least treat his important work as a guide, as we do for that of Powers.

Central are types of patterns that describe static perceptual categories that we can observe more often than we would expect to do in a random universe. These we consciously label “noun”. Dynamic patterns that are perceived as changes in the values of instances of static categories, events, we label “verb”. In the realm of narrative and event, a verb is the language equivalent of the dynamic component of an event or a narrative fragment, the accompanying noun or nouns the instances of perceptual categories that are altered by the event or mentioned in the narrative fragment. In my childhood, labels such as “noun”, “verb” and various modifying operators such as “adjective” or “adverb”, together with relationships of nouns to verbs such as “subject, ”object”, and “indirect object” were what we were taught. We had to “parse” screeds of written texts, explaining the kinds and relations of every word, in sequence.

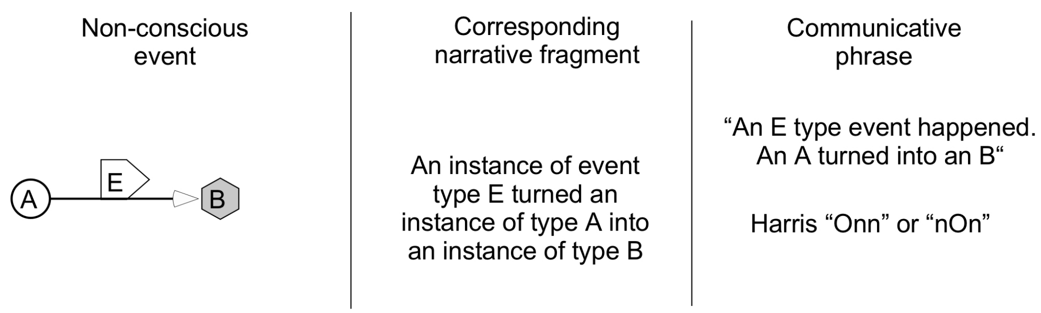

Figure II.10.12 The equivalence between a non-conscious event occurrence, a narrative fragment, and telling someone else about what happened.

How these hard-edged, consciously distinguished, categories were derived was never explained to us as school students. They were as arbitrary as the Ten Commandments handed directly to Moses by God. They “just were”. Now we, in describing the development of the perceptual control hierarchy, and Harris, in describing the types of interactions among word types, both claim to describe how those linguistic commandments came to be as they are. Harris deals largely in conscious thought, we in the non-conscious perceptual control hierarchy. The two approaches, if they are both correct, should conform across the conscious to non-conscious boundary.

Harris uses a symbol O for operator, which operates on “n” type units of language. A group such as “Onn” indicates an operator acting on two n-class units, analogous to a non-conscious “event” or a conscious “narrative fragment” that describes a change in the states of an instance of each of two perceptual categories. These are typically what we often call “nouns” (hence the symbol “n”). The symbol

“O” corresponds to the category “event”, not to the particular instance category of event, which would correspond to the actual way of expressing the verb or equivalent Operator.

The Harris view starts from a very different place than does either perceptual control or the narrative structure of rational thought, and if they all are clear views on the same structure, they necessarily emphasize different aspects of the structure. Harris looks from outside, using what we called “The Observer Viewpoint” (Section I.2.5), as does, say, a person using the Test for the Controlled Variable (TCV) in an experiment. Harris starts from the statistical relationships within samples of a language, English, and finds in those relationships a variety of categories of words.

Both PCT and Narrative start with relationships in the perceptual world based on sensory data from the environment, PCT with the effects of repeating patterns on synaptic connections, Narrative Thought with how changes in the consciously observed world relate to what might possibly be observed in the world. PCT and Narrative Thought refer to what happens within an individual, whereas language refers to the effects one individual’s actions may have on another. It may therefore be surprising how closely related to each other all three approaches seem to be. Should it be so surprising?

It should not. If the basic thermodynamic underpinning of PCT is correct, then the effects, deliberate effects or side-effects, the actions of Alice have on Bob will alter Bob’s perception, non-conscious or conscious, of some state in the perceptible environment. If Bob is old enough to have developed categories of static types or event types, he will control to reduce the error in any perception that Alice’s action has disturbed beyond its tolerance zone (if doing so does not conflict with something else he is controlling). Bob’s action will complete a feedback loop between them as Alice perceives the effects of his action as a perceptually disturbing event.

Whether this feedback loop between Alice and Bob has positive or negative loop gain depends on what perceptions each disturbs in the other and how strong that disturbance was. What perceptions are disturbed, the logon information each acquires from the actions of the other, is exactly what language is evolved to supply. Alice can use language to allow Bob to determine which of his perceptual categories may have a perceptual value of an instance she controls for him to act on.

More subtle aspects of language, including what we know as modifiers of verbs and nouns — adverbs and adjectives, but also including gesture and intonation patterns — largely deal with logon rather than metron information. Metron information is hard to transmit without language, but logon information, such as how strongly the speaker, Alice in this case, feels about what she wants to convey, can be modulated by her facial expressions and her gestures, as well as in how she pronounces the string of words. This is difficult to do in writing, and it is primarily in writing that words are clearly separated, something that is not clear to a child learning to speak. Written German, for instance, often packages as one word what might well be treated as two or more words in English, though their packaging in speech might be the same in both languages.

None of this is particular to any of the three views of communication function we have been discussing. The same inter-individual feedback loops apply whether or not language is used, whether or not the interactions involve conscious rational thought, and even whether or not the individuals are capable of language or rational thought. In the next Chapter, we explore the earliest development of language between a mother and her newborn baby. Our example is human, but the earliest stages of it should apply equally to species that nurture their offspring for some time after birth, such as great apes, dolphins, elephants, or ravens, though not to species such as turtles or frogs that produce many offspring of which very few survive to adulthood.

Do these non-human species that seem to use a form of language also have a grammar that can be formalized? Presumably they do, because if not, their sounds would be disorganized, and patterns are clearly discernible in sounds such as birdsong[48] or whale-song, and these patterns often can be identified not only with the species but also with the region or even an extended family such as a pod of dolphins. One question is the degree to which the structures of these sound patterns influence behavioural changes and identifiable patterns of response. Another is whether an new addition to a family simply learns a sound pattern by imitation or learns when and how to perform it from the return feedback from adults. The short answer to the question asked at the start of this paragraph is “It’s possible”.

Since we know more about the human use of language than we do about how, if at all, other species use language in their communications, we will continue referring to language as though it is used only by humans, recognizing that this is probably not true. Other species clearly communicate some things by patterns of sound, perhaps together with gesture. This may be especially true of herd animals (such as ourselves), including elephants, wolves, dolphins, and ravens. Even animals isolated by preference may communicate with each other, as, for instance do orangutang who are said to signal their intended direction of travel for the next day so that they would have little chance of meeting one another.

II.10.12 Braided Streams of History and Prediction

Here’s another narrative fragment in a string of words: “*John gave Jane the book*”. Like the earlier fragment example, this fragment by itself raises many possible questions about possible related earlier and future events. Among them, How did the book come into John’s possession? What did John expect to happen as a consequence of his giving it to Jane? What would a later historian perceive as consequences of the fact that the book happened to be in Jane’s possession when some event occurred? These questions demand answers in the form of narratives, and let us suppose that part of the answer about possible consequences is to be found in the fragment with which we began this chapter: "*As she was walking along the path, well trodden by long-ago ancestors, ...*".

Suppose “she” was the Jane to whom John gave the book, and the reference to the long-ago ancestors had been in the book. Possibly the fragment earlier isolated in the narrative flow of this Chapter but now set in its own narrative might have continued by describing some thoughts she was having during the “now” of the walking event about an ancestor about whom she had read, and of whom she had heard in family discussions but about whom she had known very little. The two events, though probably not contiguous in time, would form an event with a longer “now”, a “now” in which Jane clarifies her thoughts about this previously nebulous ancestral presence in her mind.

Equally possibly, the book donation event might not precede the path-walk fragment, which might continue with something like “*a tiger was stalking her, hidden by the dense underbrush*”. The fragment only tells us that a parallel narrative of thought exists in the tale-teller’s imagination. There are enormous differences among the many strong possibilities, to the extent that they might be in separate domains of thought. Such a narrative continuation would in the same “now”, or a portion of it, as the walk. If the book donation is an event in the narrative, it might occur later after John learned about Jane’s escape from the tiger.

All of these possibilities are just that — possible events that may occur before or after, but are in some way related to the walk. We readers of this Chapter may take it as very likely that “She” is Jane, and has probably been so identified earlier in the story narrative, but that is uncertain. The same is true of the

48. As a personal anecdote, I once was walking though some scrubland in the south island of New Zealand, and passed a small bird that burst into song. I wished I had been in a position to transcribe it as a musical verse, because the song was a beautiful well-formed four-line stanza of (as I remember it) ABBA structure with the internal rhythm of an English folk song. At the end of the verse, the bird flew away. I did not hear any response from other birds.

narrative fragments that describe events that might have preceded or that might follow the walk along the path. “She” might pick some flowers, fill a basket with berries, find a tree with her parents’ initials carved into it, or get attacked by a tiger.

From just what we are told in this fragment there are also unlikely precursor events. “She” had to reach the path from somewhere, and that somewhere is probably not the car of someone who kidnapped her, nor a central railway station in a big metropolis. We cannot tell very much, since we do not know what components are in the “narrative soup” out of which both fragments grew, and in particular which soup components have changed recently, both in story time and in the time of the author imagining a narrative of a reader of the story.

Just as the continuum of possible values of a perception in the analogue hierarchy is simplified by its fragmentation into categories distinguished by the possible relations to actions that could serve to control related analogue perceptions — hardness, for example, can be distinguished into hard enough to scratch glass and too soft to be used for that purpose. Knowing what is already in the soup, the haze of possible immediate predecessor or following events can be separated into event kinds that are or are not reasonably likely.

Any event has a corresponding narrative fragment. That fragment changes the contents of the narrative soup whence it emerged. Like an action in the control of perceptions through the environment, an event has consequent effects such as ripples on a pond after an event that caused a splash. Side effects of perceptual control are the consequent effects of the events — the changes in the environment — caused by control actions. They change how other controlled perceptions and their control loops may be affected. They are the elementary units of narrative when and if they become part of conscious experience. Previously implausible consequences may become more likely. If John gives Jane a book, a previously impossible consequence such as Jane using the book to swat a fly, becomes possible. Other consequences, such as Jane puts the book on a shelf with other books, may be of higher probability.

No matter what the possible consequences might be, the set of them has an uncertainty value. If, in the soup, there is a narrative event that Jane had asked John if he had that book because she wanted it, then a high probability consequence is that she might in the future read at least some part of it.

Part 6: Dyads and the Creation of Meaning

In Volume I, most of the discussion was about control effects that occur within a single organism. In Volume II, the focus changes to what happens when two or more independently functioning organisms share an environmental space. The book has the subtitle: "an Enquiry into Language, Culture, Power, and Politics". In this Part, we begin in earnest with language and its main medium, the "protocol".

The next four Chapters develop concepts of control when only two entities interact. In example of Figure II.2.9, we built a dyad consisting of two autocatalytic loops from the products of one such loop. Now we consider interactions in dyads in which the entities are independent control systems. The individual control systems are now very complex, far above the level of interacting autocatalytic loops. Nevertheless, the dyads we mostly discuss consist only of two people, though it could be a person and a computer, two wolves, or a porpoise and a human trainer. Since we are working towards an understanding of society, language, and culture, we concentrate on people while keeping in mind that many of the same principles apply much more widely.

In Chapter II.9, we meet a mother called "Cora" who has a new baby called "Ivan". Cora perceives some small degree of consistency in the variations in Ivan's activities when he is hungry or has a pin sticking into him. When Ivan uses those particular variations of movement and sound his hunger or pin-prick perceptions approach their reference values (not hungry, no pin-stick) through the actions of Cora, his reorganization slows for those movement patterns, which become methods of controlling those perceptions rather than random movements and sounds. As for Cora, she is beginning to learn a private "Ivan language". His sounds and movements begin to have meaning for her, in that she can do things for Ivan that produce sound and movement patterns that she perceives as "Ivan is feeling good". Ivan and Cora are developing basic interaction protocols.

Chapter II.12 "A Geometric Interlude" offers a short introduction to two invented geometric concepts, "syncon" and "synx", which we will use extensively in a synthetic environment that will ease our investigation of language production and understanding in the following Chapter and in many other places in the book. A "syncon" is related to a reference value, whereas a "synx" is related to a perception or an action. They are not simply other words for the same things: syncons are relatively fixed target locations in spaces of various dimensionalities, whereas a synx can move around freely through the space.

Chapter II.13 introduces a family of simplified beings to whom we give human names, beginning with Len, who soon meets Sophie, and together they have a child we call Dan. They are able to perceive the values of perceptual variables in the three sonic dimensions that define syncons. These syncons represent the entire pseudo-phoneme complement of what will become their language. The family members can perceive and emit sounds that range over all possible values of the variables, not only with the particular combinations of values that define the fixed syncons.

The chapter begins with a robot called "Rob" that can perform certain actions and can produce a synx in the form of a sound having defined values of the three parameters or sonic dimensions. Rob's synx productions are restricted to the values of his fixed set of syncons. Rob can interact with the initial member of the family, Len. Len's synx can range freely through the available space of variation of the three component values, but he learns first to produce a synx with values near those of any selected one of Rob's syncons, and then to produce a synx trajectory that approximates the sequences of values of syncons used by Rob. Because Len used perceptual control in bringing his synx to its syncon reference values, his approaches to syncons is smooth and imperfect, as opposed to Rob's which is abrupt and precise. Len learns how to use syncon sequences we can call "words" to get Rob to do a few things Rob has been programmed to be able to do.

Now Len meets Sophie. Sophie learns from Len as Len learned from Rob. The syncons Sophie learns from Len are less precise than the values that Len learned from Rob, but because Sophie's synx trajectories as well as her syncon values can match Len's, her imprecise synx approximations to Len's syncons can become intelligible to Len. In the process, Sophie develops new synx trajectories to represent new concepts that Len learns from her. Len and Sophie both teach their child Dan, who learns not just how certain trajectories can be used but also how they differ between Len and Sophie. The result is that Dan hears similar sounds produced by the two of them in different contexts as meaning different things and different sounds as meaning the same things. These concepts are developed to higher perceptual levels, and the Chapter concludes with a discussion of phonetic symbolism and the stability of learned perceptual functions.

Chapter xII.13 returns to the ordinary human world, though we will return to the syncon world in Chapter xII.15, when we talk about larger communities. In Chapter xII.13, we discuss protocols and how protocols operate through a set of controlled perceptions of different beliefs and uncertainties each has about the other's beliefs and uncertainties. There are nine such perceptions in each partner in a fully developed protocol, the same nine in every protocol, but fewer if one of the partners is a computer that has not learned much about the world or about perceiving states in its user. Finally, we discuss the similarities and differences between protocols and ritual on the one hand and the use of protocols for deceit and camouflage on the other.

Chapter II.11. Private Language of Mother and Baby

In everyday language, a protocol is a way of communication between partners who cooperate in communicating, no matter how much they may disagree on the content of the communication. As the diplomat, George Ignatieff said: "I learned that protocol is really a language, a set of rules and conventions which enable people of different nationalities, social backgrounds and political persuasions to feel comfortable with each other" (Ignatieff, 1985).

A protocol is a negative feedback loop, to which an “initiator” introduces a disturbance that is influenced by the action of the “continuer” who may be controlling some unrelated perception. It is a two-node homeostatic loop[49] that may be a control loop. In PCT terms, a protocol is a structure encompassing at least two partners, an “initiator” (who we will call Irene) who controls a perception “iPerc” and a “continuer” (Charles) who controls a perception “cPerc” in such a way that her control of iPerc disturbs his cPerc. Charles’s action in controlling cPerc reduces the error in iPerc and decreases Irene’s disturbance to cPerc, reducing the error in cPerc[50].

For the protocol to work well for both partners, the continuer will often be controlling some perception such as of the initiator’s satisfaction, with a reference value for the initiator being satisfied. To ease the continuer’s ability to control, the initiator *displays* the nature of the perception she wants the continuer to influence, often but not always through the use of language. For example, if I say to you “Please close the door”, I display that I am controlling a perception of the door state, with a reference of it being closed, that I do not perceive it as being closed (though in fact it might be), and that I do not intend to close it myself.

I could *display* these same things with hand gestures or body language, if I thought you would perceive correctly what I am trying to display. By creating the display, I am also displaying that I perceive you to be controlling a self-perception of being a helpful person (at least to me), or perhaps to be controlling for perceiving me to be content or something similar, perceptions that I disturb by creating my initiating display[51].

Stated in this way, a protocol might be seen as just a way in which one person can help another to do something that would otherwise be difficult. But protocols can serve other purposes equally well. I might want to control my perception of something in your World Model, or perhaps to control for you to become more kindly disposed toward me, perceptions I cannot control without your showing me your state of knowledge, understanding, or friendship. Your cooperation in changing that state consists of using a protocol that both of us have reorganized to use, which requires a medium of communication.

Let us examine how a protocol might be produced, starting with a loop that includes side-effect interactions and partners who have no common language.

49. Though sometimes more than two people are involved, making it an N-node loop where N is the number of people concerned. We ignore such cases in this book.
50. In a notation described in Appendix 3 for compactness in describing interacting control loops, this would read *Irene*{*iPerc*[*Charles*{*cPerc*[*Irene*^), which might be transcribed as: Irene controls “{“ iPerc using as an atenfel “[“ Charles controlling “{“ cPerc using as an atenfel “[“ Irene, as at the start of the string “^” in a continuous loop.
51. Sometimes we refer to the initiating display as an “initiating protocol”, and the resulting action as a “continuing protocol”, because both the display itself and the action are likely to be implemented through a hierarchy of supporting protocols.

II.11.1 Baby's early interaction: getting someone to feed it

Not all influences between different control units are side-effects. They may be deliberate, meaning that the influence itself is the CEV of a controlled perception. An infant does not have a fully developed structure of elementary control units whose reference functions and perceptual input functions are simply sitting there waiting to be properly interconnected by reorganization. It may be able to move its body parts and to cry or gurgle, but any effects those actions have on anyone else are side-effects. Now we suggest how those side-effects may become intentional effects through reorganization.

The baby may well have a genetic growth plan for the types of perceptual functions it will later develop and control, as it does for the development of its physical features, but as a newborn, and indeed for at least a couple of decades thereafter, it does not have the perceptual functions it will have as a mature adult human. Much of its developmental course depends on the environment into which it is born, an environment both physical and social, in which the actions of more mature people are crucial for the baby's survival.

As is the case for any organism, the infant stays alive by controlling various perceptions in such a way that its intrinsic variables stay within genetically determined limits[52]. Like all living things, it can do this only by acting on its environment in ways that bring important controlled perceptions to, or maintain them near, their reference values. Chapter 7 and much of Part 5 of this book suggest how these ways may develop in a maturing organism.

Imagine a newborn infant left all alone for a long time. It might be able to cry, move its eyes and its various extremities, but, if it is alone, none of its actions upon its physical environment will provide it with food or water, and it will die. To control its perceptions effectively, it needs other people — or at least one other person. Almost the very first perceptual controlling it does must be social.

To control a perception such as the fullness of its stomach, the newborn infant, who we will now call Ivan, must act in ways that disturb some controlled perception(s) in another person so that the other person's control actions feed it. Ivan cries, not directing its output toward any particular end, and its mother (Cora) feeds him (Figure II.11.1). Ivan has not learned from experience how to cry, any more than a newborn foal has learned how to stand up and walk. Cora controls a perception of Ivan's happiness, and crying disturbs that controlled perception. If she can act to stop him crying, the error in her controlled perception is reduced. Cora controls her perception through Ivan, and Ivan controls through Cora, though initially only through the side-effects of his actions.

52. Remember that over the last four Chapters of Volume 1, we argued that the genetic reference value of an intrinsic variable may well vary over time, perhaps cyclically and perhaps as a consequence of other events in the body, because it probably is a variable in a homeostatic loop. The sensation of "hunger" may well be caused by changes in the value of an intrinsic variable over time.

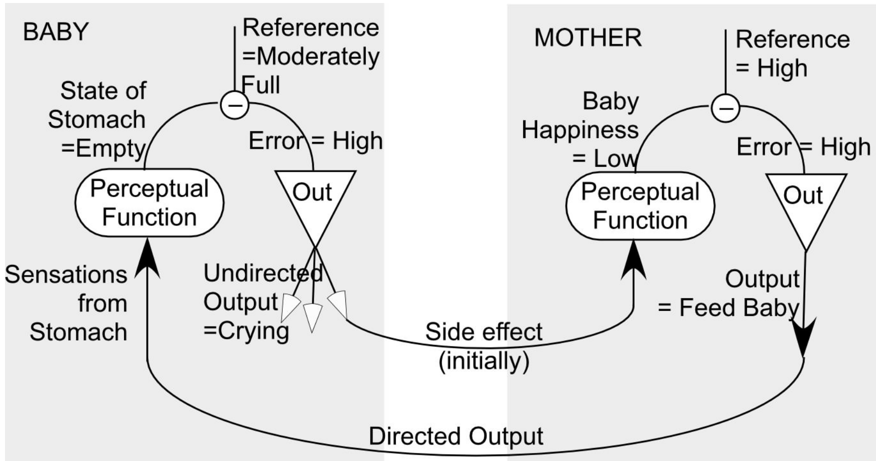

Figure II.11.1. Schematic of what might be the first externally observable control loop of a baby. It is a social control loop, incorporating a controlled perception in the mother, the environmental feedback path of which passes through the social partner (baby).

Since Ivan has not yet reorganized by learning from experience in the world, any connection between the sensation of an empty stomach and a reference for perceiving the stomach to be less empty must be genetic. He may not even perceive his discomfort as “hunger” or “empty stomach”, but what we, as Analysts, call ”empty stomach” leads the baby to cry, and the crying stops when the discomfort stops.

Crying is by no means Ivan’s only output to the environment when he is hungry or otherwise discontent, though it is the only one shown in Figure II.11.1. For example, he may output limb movements, facial expressions, and sounds of all kinds, and Cora’s “Baby Happiness Level” perceptual function might use any subset of these correlated features as possible contributors to her perception of Ivan’s state of happiness.

Since according to almost any possible version of PCT “all [intentional] behaviour is the control of perception” the initiating partner’s actions happen because some perception differs from its reference value. If the Cora’s actions do not reduce the disturbance caused by Ivan’s actions, in both partners a controlled perception remains in error.

In the Powers version of PCT, one consequence of failure to control a perception successfully is an increased rate of reorganization. The fact that the initiator’s (Ivan’s) environmental feedback pathway goes through another person (Cora) is irrelevant to the reorganization process in either of them. On the other hand, if the two reorganizations fail to mesh properly, neither will be able to counter consistently the disturbances that influence their respective controlled perceptions. Accordingly, the two reorganizations will continue until both can control their relevant perceptions or until Ivan dies of malnutrition. The reorganizations may eventually produce coordination, so looking from outside at both parties at once, we might consider this double process to be a single phenomenon we could label “co-reorganization”.

If Cora comes to perceive that Ivan waving his left leg and right arm while voicing a rising tone means he is hungry, and often feeds him after observing that pattern of behaviour, then Ivan will eventually reorganize so that his action outputs when he feels hungry will be to wave his left leg and right arm while voicing a rising tone. Ivan and Cora together have learned a “word”: Ivan-hungry (according to mother’s perception) = wave-left-leg-and-right-arm-and-make-rising-tone. But Cora will not use that “word” to let Ivan know she is hungry. Even in natural languages, including Japanese and Korean, the words used for the same thing may differ between protocol partners of, say, different gender or social standing.

More probably, if a baby cries and waves its arms about because genetically it has been programmed to do so, and its mother's midwife tells her that this means she should feed it, then those actions will continue to be used by the baby and will become its "hungry-word". Moreover, the mother will quite probably tell her friends who become new mothers that this is what babies do when they are hungry, so they will act in ways that induce their babies reorganize to use the same "language" of gesture and sound when hungry. That kind of "hungry-word" becomes part of the public language of babies rather than being part of the private language of one mother-baby pair.

Ivan did not actively choose the action components of the "hungry-word", which could have been any pattern of movements or sounds he was capable of producing deliberately. Rather, the structure of the "hungry-word" is the result of reorganizing the output components of the feedback loop involved in controlling the perception of general discomfort, which often includes empty stomach.

When Ivan emits this initially arbitrary display, Cora is likely to try feeding him. Feeding Ivan is her continuation protocol that complements his initiating protocol and brings controlled perceptions in both of them closer to their reference values. All language, whether in words or in gestures and facial expressions, can be seen as being, or at least as deriving from, the display element of a protocol. The concept of "language" is inseparable from that of a protocol.

It is crucially important to note the arbitrariness of the "hungry-word". All words of every language (except for those that mimic non-language sounds) are equally arbitrary, whether they are purely voiced, use voice and facial or other bodily gestures, or are purely transmitted visually. Language is not even limited to sound patterns, as is attested by the existence of various national sign languages communicated entirely through gesture.

Collective control, as we shall see, stabilizes over time the words and larger constructions used for different purposes, but though this stabilization allows, indeed requires, adults to correct a child's incorrect usage, maybe telling her that "farbling" doesn't mean what you do when you cook slices cut from a dead animal, but means what you do when you cook a whole dead animal. The word —any word — may have a correct usage within that community, and may had the same "correct" meaning for centuries, but its form is no less arbitrary for that. Indeed, if cooking whole animals goes out of fashion, "farbling" might come to mean simply "cooking meat". An incorrect usage would have become correct.

II.11.2 The Generic Protocol Form

Whereas a simple control loop with one ECU controls one perception, a protocol loop controls two perceptions of different kinds in two different people. Unlike the side-effect loops shown in the upper panel of Figure II.11.2, each partner in a protocol loop acts intentionally to influence a perception in the other, as shown in the lower panel of Figure II.11.2. The partner opposes this disturbance by some action that influences the other's perception.

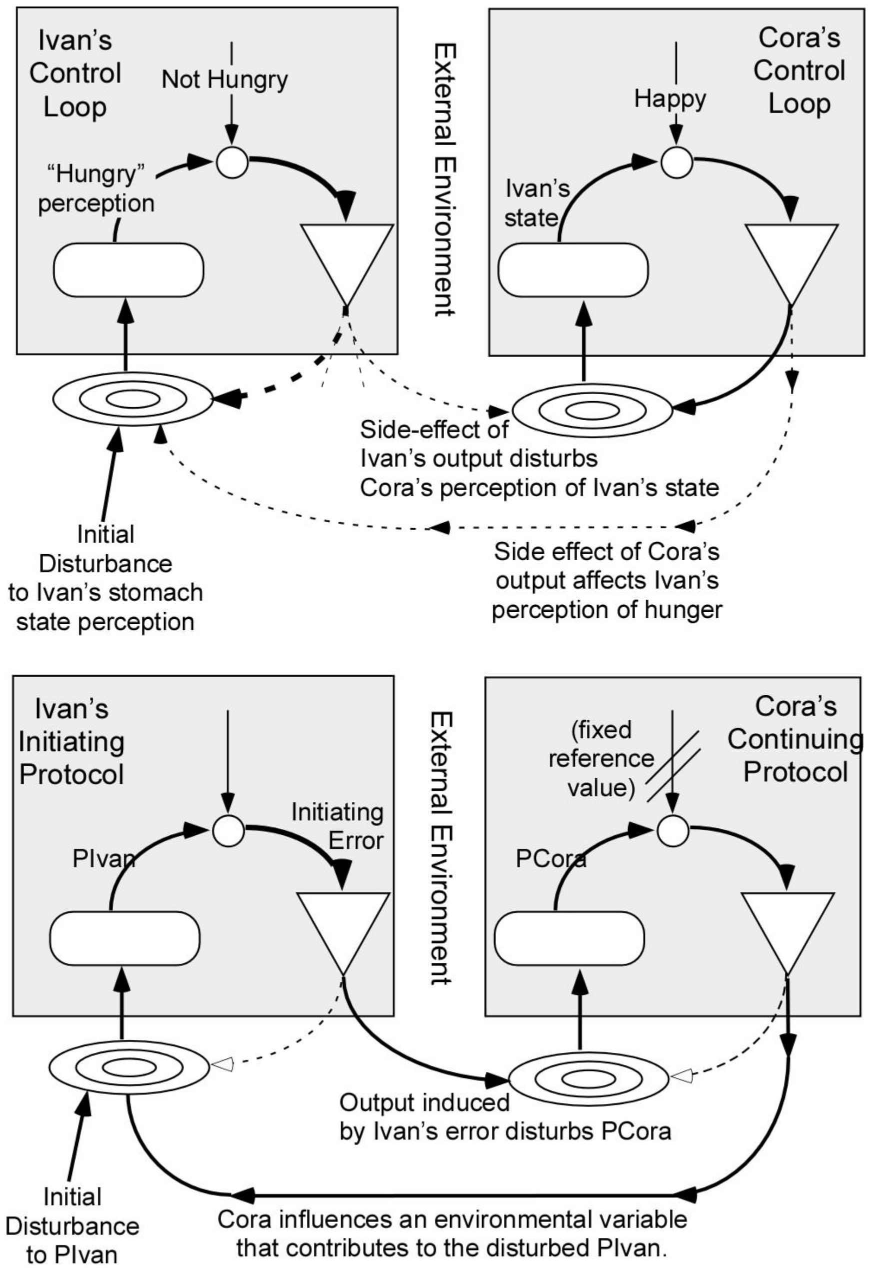

Figure II.11.2. The development and structure of a basic protocol pair. (a, top) baby Ivan and mother Cora each control some perception, oblivious to the presence of the other, though Ivan's control is ineffective. A side-effect of Ivan's ineffective control action disturbs a perception of Ivan's state that Cora is controlling, and a side-effect of Cora's control disturbs a perception Ivan is controlling. In (b, bottom) Ivan has a controlled perception that differs from its reference. He cannot influence it by his own actions, but Cora can. Ivan disturbs a controlled perception in Cora that Cora cannot influence directly. As Cora influences Ivan's perception in a way that reduces its error, Ivan's reduced output (crying) also reduces her error.

Since the influenced perceptions in the two partners are of different properties of the environment, no conflict is implied by the fact that two control loops are operating in concert. The mutual influences could, in principle, form a positive feedback loop, a negative feedback loop, or have a near-zero loop gain, but in the long run reorganization is likely to eliminate a positive feedback loop, and may improve the performance of a mutually beneficial negative feedback loop.

Tracing the protocol loop of the lower panel of Figure II.11.2, Ivan has a controlled perception we call "PIvan" that differs from its reference value, either because some external disturbance influenced it, or because its reference value changed. Baby Ivan might have been stuck by a pin, or maybe he can't reach a toy. Whatever the reason, Ivan either cannot or does not want to influence his own PIvan directly, but Cora can influence it (though she cannot perceive it) and probably will. Ivan's initially random action disturbs a perception controlled by Cora, which we label "PCora". Cora is unable to influence PCora directly, because it is a perception of a state in Ivan, but Cora can influence the CEV corresponding to PIvan and thus influence PIvan. She could give Ivan the toy, for example, or remove the pin. If Cora does something that reduces the error in PIvan, Ivan's actions alter the value of PCora. If the error in PCora is reduced, the loop is complete, and reorganization is unlikely to change it drastically.

Cora never perceives PIvan, but to control PCora she must act on some elements of its CEV that is influenced by Ivan's actions in controlling PIvan. Figure II.11.2 suggests that this action is initially done blindly, but that is unlikely to happen in a well-developed protocol. In a well-developed protocol, Ivan has displayed what he wanted. In other words, for different problems, Ivan acts differently in a way that Cora can perceive — uses different "Ivan-words" that Cora can recognize. His crying when he needs changing might be consistently different from his crying when he is hungry, for example. Since Ivan almost certainly uses outputs that differ in detail from one time of discomfort to the next, Cora's perceptual function for discerning the type of discomfort suffered by Ivan must be developed in Cora by reorganization that depends on the quality of control within each protocol loop.

Since Ivan (in this example) is an infant, he probably has a more labile control structure than does the adult Cora. After all, the reason he needs Cora to feed him, to remove the pin, or to move the toy, is that no control systems have yet been built that would enable him to feed himself, to remove the pin, or to go and get the toy. Accordingly, most of the reorganization in the co-reorganization process will probably occur in Ivan. Ivan has to learn to "get along" in the culture into which he has been born, because it will not only be Cora with whom he will have to interact as he grows a little older.

Most mothers (and fathers) use initiating protocols with babies that have not yet reorganized to create the corresponding continuation protocols. The parents coo and they cuddle, and they hope for a smile in return. The lack of a smile creates error in a perception they control, so they reorganize until their actions are accompanied by a baby smile. Maybe the smile indicates baby is happy; maybe it doesn't. Either way, it reduces an error in the parent, and the parent learns to act in ways that produce smiles more often than not.

Eventually, baby learns the appropriate continuing protocol, smiling and gurgling. But this does not necessarily mean the baby is happy. All it means is that the parent perceives the baby to be happy. This same uncertainty will remain for the rest of the baby's (or anyone else's) life.

Acting for another's benefit is called "cooperation", and is often claimed to be problematic. Ito and Yoshimura (2015) in the introductory sentence of their abstract, say: *"Why cooperation is well developed in human society is an unsolved question in biological and human sciences."* In contrast, as we saw in the last few Chapters of Volume 1, PCT is built on cooperation at very low levels, and the same argument suggests that cooperation is almost an inherent property of large enough groups of organisms, human or otherwise — and so, for the same reason, is conflict.

II.11.3 Cora and Ivan: Early Development of Meaning

Although unaware that living organisms act to control their perceptions, J. G. Taylor[53] (1962) theorized and demonstrated experimentally that perceptions could be built if (and probably only if) the person acted in a feedback loop that depended on or that directly influenced the perception. The feedback loop was, to Taylor, the essential element in creating the ability to perceive any aspect of the environment. Although his "Behavioural Basis of Perception" theory was based on Hullian reinforcement theory and Taylor was unfamiliar with control, it would remain essentially unchanged if the Hullian basis were to be replaced by Powers's reorganization process and the feedback loops had the asymmetry required for control. I therefore consider his work and that of Powers to be complementary and mutually supportive.

In a tenth-anniversary Festschrift for J. G. Taylor, I argued (M. M. Taylor, 1973) that new perceptions might be produced by the Hebbian-anti-Hebbian (HaH) process in a variety of ways: by modification of an existing perception, by statistical correlation among existing perceptions, and by splitting of an existing perception into distinct sub-types[54]. According to J. G. Taylor, and according to Powers's concept of reorganization, such perceptions would persist if used in active and effective feedback loops, but would degrade if unused and subsequent processes created interference. In the context of our discussion of interpersonal interaction, the active feedback loop is the protocol loop, and both parties must develop and stabilize the perceptions controlled in the loop if the protocol is to remain viable.

Also in the 1960s and early 1970s, Held, Hein and their collaborators (e.g. Cynader, Berman and Hein, 1973, 1976, Held and Bossom, 1961, Held and Hein, 1963, Held and Mikaelian, 1964, Hein, Held and Gower, 1970) were showing that new perceptions of various kinds (in humans and kittens) either could be developed only with much difficulty or were lost in the absence of active feedback, whereas they could be formed much more readily if they were part of an active feedback loop. Held and Freedman (1963) also demonstrated that when feedback became inconsistent, controlled perception also became less reliable, at least as manifest in the corrective actions. All of the results of this group were consistent with J.G. Taylor's theory, though they seem to have been unaware of his work. With this old theoretical and experimental work as background, we can return to the example of mother Cora and baby Ivan.

Taylor's theory suggests that if Cora is Ivan's primary caregiver and has reorganized to control her perception of Ivan's "discomfort" by feeding him, she would soon learn to perceive Ivan's differently correlated output patterns of sound and movement for hunger, pin-stick, and desire for toy as meaningfully distinct displays. Cora's "feeding" actions would have little effect on discomfort caused by a pin-prick, which would increase the likelihood of reorganization that would eventually enhance her ability to distinguish the "pin prick" display.

Unfortunately, however, Ivan's output displays are initially undifferentiated (or so we may assume for the purpose of argument). They might change randomly without relevance to the circumstances leading to the discomfort. The consequence sounds like an impasse, a metastable equilibrium situation in which all Ivan's displays could result in any of Cora's behaviours.

The equilibrium, however, if one exists at all, is unstable. In the process of reorganization Ivan changes his "discomfort" outputs for the various kinds of discomfort, perhaps randomly at first. If a certain display by Ivan on one occasion results in Cora acting to reduce the error in a perception that is the source of Ivan's discomfort, Powers's reorganization principle suggests that Ivan's pattern of outputs is less likely to be changed and more likely to be used the next time Ivan's control of that particular

53. No relation to the present author.

54. J. G. Taylor wrote a comment following my contribution in which he thanked me for having both found and fixed a gap in the logic of his theory.

perception has an error[55]. The symmetry that leads to the apparent equilibrium is broken. Ivan's random pattern of outputs on that occasion is likely to be different than the pattern he produces the next time a different "discomfort" occurs. If Cora finds a way to reduce Ivan's perceptual error for this other circumstance, the protocol loop is well on its way to splitting into different protocols, one for each different perception that Ivan is learning to control.

Important in this is the environmental context. If having observed one pattern of Ivan's outputs Cora finds a pin sticking in him and reduces the discomfort by removing the pin, and for another pattern of outputs Cora finds no obvious environmental problem but reduces Ivan's crying and limb-waving by feeding him, Cora will reorganize to create what we call "meaning". Pattern A of Ivan's output "means" to Cora that Ivan has a pin sticking in him, whereas pattern B "means" Ivan is hungry and pattern C means Ivan wants his toy. Patterns "A", "B", and "C" are quite arbitrary, but because of the co-reorganizations in Ivan and Cora, they mean something. They are "Ivan-words" with which Ivan and Cora are starting to build a simple, if private, language. When a visitor hears Ivan crying and says "Ivan is hungry", Cora can say "No, he is saying that he wants his toy".

To jump ahead in our story, when Ivan has after many years become a culturally and linguistically competent adult, he will have reorganized to produce on purpose sounds and gestures that Cora can distinguish easily, so that she can perceive Ivan's "meanings" in many different contexts. Ivan and Cora will have developed many, many protocols, the entirety of which constitute their particular personal language and culture.

The display languages of protocols help each participant perceive something about some state in the other party, provided both parties use the same complementary protocol pair structure. If only the two of them ever interacted, that would be the end of the story, but they are not the only ones. Other language-using people inhabit their two social worlds.

By the time Ivan has grown up to interact with many other people his own age or older, he will have discovered (probably not consciously) that his private "Ivan-words" used with Cora will not have the effect he wants when used with other people. They will not mean to other people what they mean to Cora. He will reorganize, so that his protocol display "language" comes to influence more people he meets in ways that reduce error in some of his controlled perceptions.

The Analyst (and perhaps the Analyst within Ivan) recognizes that this happens because Ivan's words begin to sound like words used by other people to convey particular "meanings". His early "Ivan words" become disused. He no longer waves his hands and feet and cries in a certain way when he wants his toy. He may say "Toy" instead, using voice with an appropriate intonation, and a year or two later he may say "Would you mind getting my toy out of the cupboard, because I can't reach it."

The structures that produced the early "Ivan words" may still be there, but they will be used less and less often as Ivan reduces the likelihood of getting frustrated and angry as he finds more precise ways of getting what he wants using language. Reorganization might eventually reconfigure them out of existence, but this is not guaranteed. It is said that one never forgets how to ride a bicycle, no matter how many years elapse since one used that talent. Crying in frustration might be a kind of bicycle-riding in that respect.

If you want another person to do something to reduce error in a perception you control but cannot influence, it might be useful to let that person perceive what you want, at least if you perceive the other as

55. We see here a primitive form of superstition, but it is a superstition that will in the end prove to be effective for Ivan; it is also the mechanism that we suggest below leads to the phenomenon called "phonetic symbolism".

being willing to help you. “To let that person perceive” something is to convey a meaning. But the concept of “meaning” is, like that of “information”, somewhat slippery, and we will return to it after a while, in Chapter xII.11 and more particularly in Chapter 1.

II.11.4 Protocols and Social Perception

Influencing one’s own controlled perceptions by disturbing perceptions controlled by others is a central feature of socialization, for humans and for any other sufficiently complex organism. Even bacteria and plants cannot survive and propagate in isolation from other living things. We all live surrounded by other life; indeed, we host far more bacteria than we have cells in our bodies, and it is becoming ever more evident that the actions of these individual bacteria influence our health in important ways. Indeed, if, as we suggested in Chapter II.3, we are a vehicle constructed on behalf of the welfare of our microbes, the relationship is reciprocal, a feedback loop that may be a kind of homeostatic loop.

Whether we know it or not, almost any action we take will influence some perception in some other living things, just as newborn Ivan probably does not know at first that his crying influences Cora’s controlled perception of his happiness. If one person’s control actions affect a controlled perception in another living thing, the influence is probably a disturbance[56] that induces error that the other may act to reduce, whether the other be a human, a pigeon, a bacterium, or a tree.

In most protocols that have co-reorganized to a reasonably stable state, an initiating protocol in one partner is matched by a continuation protocol in the other. If I meet you, a friend, on the street and stop to say “Good Morning”, I may expect you to say something like “Good Morning, how are you today”, to which I would be much more likely to say “Fine” than to list my aches and pains. The conversation might then continue.

If you simply return “Good Morning”, I might perceive you as being less interested in continuing the interchange. You might also ignore me, but I would not expect you to return “Beans are $3 today”, though it is possible to imagine circumstances in which you might say just that. If you did, it would probably disturb some other perceptions I was controlling and I might act to counter the disturbances. I might, for example, go and buy some beans. Or if we were two spies, the words might be a code for “Meet me at the park at three o’clock”, to which I might reply “But broccoli is also on sale” which might mean “Message understood”.

If I start with “Hey, Buddy”, you are not likely to return “Good Morning, sir”. “Hey Buddy” signifies a different implementation of a “Greeting” protocol, the continuation complement of which might be a simple “Hi” or “What’s up?”, both of which might lead me to perceive you to be presently of good will toward me, whereas “I’m not your buddy” might lead me to perceive the opposite, and a flat “Good Morning” might lead me to perceive that something is bothering you (meaning that you have some controlled perception that is in error and your actions are not currently effective in reducing that error).

The “Good Morning” protocol pair and its cognates may exist only within some cultures, but where they exist, they are fairly stable. On the surface, “Good Morning” seems to convey very little information, especially when said to someone seen often, since in such a culture, the probability that Charles will say “Good Morning” when first seeing Isabelle is near 1.0, and the probability that Isabelle will return “Good Morning” is also near unity. The greetings being so uninformative, why would either occur at all?

Suppose Isabelle failed to say “Good Morning”, or indeed anything at all. What might Charles

56. The influence may not disturb the perception directly; it might affect the other’s environmental feedback pathway so as to change the action the other uses to control the perception. For example, if one wants someone to choose a particular route, one might block a route the other usually takes.

perceive? Might something be wrong with Isabelle? Might Isabelle be angry at him for some reason? Might Isabelle be preoccupied and have failed to notice his presence? Might Isabelle have just arrived from a country with different protocols? Charles's "Good Morning" is primarily a disturbance to something Irene is controlling. How she acts to correct any error created by that disturbance might allow Charles to determine what that controlled perception might be.

Suppose Isabelle returned "Good Afternoon" to Charles's "Good Morning". Depending on the situational context, Charles might take her meaning as "Why did you sleep so long", "We had an appointment for an hour ago", "When did last night's party stop", all three of these at the same time, or a host of other possibilities. Just as "Ivan words" can have meanings that implicate a particular context to Cora, so anything Isabelle says to Charles has meaning to Charles that depends on both the situation as Charles perceives it and the situation as Charles perceives Isabelle to perceive it. The meaning is what Charles perceives it to be, not as Isabelle intends it to be, though Charles may wish to determine what Isabelle intended. How he can do this is the province of a "protocol grammar" (Taylor, Farrell and Hollands, 1999b), which we discuss in Chapter xII.13.

Well developed mature protocols need not use verbal language any more than Ivan does in his initial interactions with Cora. Imagine a prehistoric band of hunters stalking a prey animal. The leader might want to have the party surround it silently before moving in for the kill. He might use a protocol that if executed in verbal form might start "I want to talk to you, Albert" but when executed in gesture might be a pointed finger. When Albert is perceived to be paying attention, the leader might make a circling movement of the finger to show Albert that he should go around behind the prey. The gesture would be arbitrary, known to the other members of the party as representing a certain perception the leader wants to have, in exactly the way that sports teams may use private gestures to communicate their intentions in ways secret to the opposition. A wolf pack on the hunt may likewise use private gestures. Protocols are not limited to humans.

Protocols can, of course, be used deceptively. A greeting hug can disguise a stab in the back. But for the most part, and especially when used simply as part of controlling a perception such as the baby's hunger, they have been reorganized to form part of the regular perceptual control hierarchy, and are used unconsciously and honestly.

II.11.5 Cora's Conflict

Trading or bartering is a common kind of protocol in most cultures, but as a protocol it has some unique aspects. All of the following statements are taken for granted in most discussions of PCT:

- Perceptual control is all about reducing the difference between a perceptual value — the apparent state of the world — and its reference value. That difference is the perceptual error;
- interaction between the means by which two or more perceptions reduce their perceptual error can cause conflict, in which reduction of error in one perception involves increasing the error in the other or others;
- reorganization tends to reduce conflict.

In contrast, trading appears on the surface to *require* conflict, so one would naively expect reorganization to disrupt any trading protocol. Clearly, it does not. We examine why here, but we go into it in more detail in Chapter III.8 on Ownership, Trade, and Money.

First we look at why trading inherently seems to involve conflict, starting with mother Cora and baby Ivan. In Figure II.11.2 the protocol loop between them suggests how Cora and Ivan have found a common language that allows Ivan to control his perceptions of hunger, being stuck by a pin, or of having his toy by letting Cora perceive that he is unhappy and the reason for it. If we now think of Cora, we realize that she may well have a conflict.

If Ivan did not exist, Cora might at that moment be controlling for perceiving herself to be in a nightclub, reading a book, or simply sleeping. While she is controlling for perceiving Ivan to be happy, she can do none of those things, because they conflict. How often does a normal mother say to herself or aloud “I’ve only got two hands”?

Cora resolves this conflict according to which is more important at that moment, her perception of Ivan’s or her perception of something else she would like to have happening. We can use another word, and ask which is “worth” more to her, perceiving Ivan happy or perceiving herself to be reading a book. Supposing those to be the only alternatives in question, it is worth more to her to control one of those perceptions than the other, and either is worth more than controlling neither.

Back in Volume 1 (Section I.6.2) we took “worth” to be closely related to the ability to control one’s perceptions. We even gave it a metric, in the form of energy lost to the environment when a neuron fires. Powers analyzed control loops using the concept of “neural current”, the total firing rate of neurons in a “bundle” that tended to fire under similar conditions. If we use this concept again, we can say that smaller neural currents lose less energy, and thus deplete one’s total worth more slowly than do high neural currents.

The error value being one of those neural currents, as is the output from an Elementary Control Unit (ECU), it seems as though keeping a perception near its reference value has a “worth” greater than zero, as does controlling with a low gain. These two effects work in opposite directions, since a high loop gain implies rapid error reduction, so long as the high gain does not induce loop instability. Rapid error reduction, in turn, means rapid reduction in the output neural current (though not necessarily in the output to the environment, which must be sustained if the disturbance is sustained).

Using this metric, we can say that to Cora two things determine the worth of controlling one perception as opposed to the other: how much she cares about the perception that might be controlled (the gain or gain rate in her control loop), and how much error there is in each perception. By making Ivan seem happy, she might get to her book quickly, but if she reads instead of helping Ivan, she will continue to perceive error in her perception of his happiness. If not immediately, she will soon either attend to Ivan or go somewhere she does not hear or see him. That choice depends on a different conflict. She may believe that Ivan will soon go to sleep if he is left alone, for example.

The interaction between Cora and Ivan involves a conflict in Cora, but not in Ivan. Cora has to trade controlling one perception agains controlling the other, because she can’t control both at the same time. This introduces tensegrity considerations. The two perceptions are fixed relative to each other, forming a tensegrity “rod”, while the notional attraction of any perception toward its reference value, measured by the error and the loop gain, acts as a wire in tension. We could go back to considerations of energy and entropy to justify these claims, since we use energy dissipation as a measure inverse to worth, and control uses that energy to reduce entropy, but that kind of detail is unnecessary to make the point.

Ivan has no such conflict, and makes no choices. He appears in the analysis only insofar as his presence induces a conflict in Cora. To observe the tensegrity effects in a dyad, we must deal with two people who both have to make a choice as they interact. Where this choice issue is clearest is in an overt trading interaction, so in Chapter III.5.we introduce two new players, Tam and Bud.

Chapter II.12. A Geometric Interlude

Before we go further in our enquiry into the development of language, we must introduce a little preliminary material. In this short interlude, Section II.12.1 describes a simulation developed by Powers, which he called “The Little Man”. The Little Man views a target moving in a 3-D space and moves its fingertip to follow the movement of the target. In Section II.12.2 we will introduce an artificial environment, a three-feature space in which may be located identifiable points that we call “syncons” for no special reason other than to use a pronounceable word not encountered in normal English. Much of Chapter 1, “Growing a Language”, uses the syncon space, as does Chapter xII.15, “The Fractal Community”.

In the three-feature syncon space, an entity may move a point called a “synx”, analogous to the fingertip of the Little Man. Syncons typically serve as fixed target locations in the space, while the synx can represent a freely variable vector of feature values (a location in the space) that can be produced or compared with the syncon vectors. The synx is controlled by a talker and observed by a listener. At low levels of a perceptual control hierarchy, synx trajectories are analogous to syllables or words in natural human language, and at a higher perceptual level they are analogous to the panoply of possible phrase or sentence meanings.

II.12.1 The Little Man

We start our geometric interlude with Powers’s “Little Man” demonstration (Powers.1999)[57]. The Little Man (Figure II.12.1) will be a metaphor for much of the following development. Here is how Powers described the demonstration in the introduction to his 1999 paper:

> *Models of human arm position control (and other simple behaviors) have been of interest for many years. During the 1950s and early 1960s, it was assumed that such models would naturally take the form of negative feedback control systems.. Only a handful of behavioral scientists ever learned and simulated the properties of negative feedback control systems before that whole approach was overshadowed by the concept of the brain as a digital computer.*
>
> *The digital frame of mind led back to, or helped to preserve, a view of the brain that was put forth by the neurologist Sherrington at the beginning of the 20th century. In this view, higher “executive” centers in the brain issued general commands which were sent to lower centers where they were elaborated into more specific commands. Ultimately, all the proliferated detailed commands were funneled into the “final common pathway" where they commanded muscle contractions and subsequent physical events which we recognize as overt behavior. Even after the view of the brain as a Von Neuman digital computer was largely abandoned, this command-driven concept persisted.*
>
> *Neither Sherrington nor those who adopted his theory recognized the problems that this view would present to a modeler. There is no unique path from the general to the particular. An executive decision to “Buy 100 shares of AT&T” does not take into account that the stock market is closed today, that one's broker is sick at home, that one has forgotten his telephone number anyway, and that one has laryngitis. Unpacked, this decision is equivalent to saying “Generate any specific commands for action required in order that the consequences can be correctly described as my having bought 100 shares of AT&T”. There is a certain flavor of*

57. Available for the PC at http://www.livingcontrolsystems.com/demos/tutor_pct.html and on a CD included with LCS III (Powers, 2008)

"and then a miracle occurs" in this concept.

The specific commands required to produce a predefined result depend on whatever processes may intervene between the command and the final effect it is to have. So given the desired final effect, it is necessary to find the inverse of the intervening processes; that is what determines the command necessary to create a given result. This consequence of adopting a command-driven model has forced itself on modelers of behavior (and on roboticists) over the last several decades.

The literature of motor control has contained many examples of models in which commands for motor acts are formulated by computational processes that derive the inverse kinematics and inverse dynamics of limb movement. These models assume that goals are given in the form of desired trajectories of movement of a limb in space, with the end-points and the derivatives being specified in advance. The required end-points and derivatives are converted by inverse kinematic calculations into equivalent joint-angle derivatives and end-points; then the inverse dynamical calculations are applied to produce the torques required to cause the desired patterns of joint movement and final angles. The computing capabilities required of the brain in order to carry out this kind of feat, as well as the demands for knowledge about the physical properties of the muscles, body and external world, would be enough to give pause to anyone who did not have faith in the unlimited computing power of the brain.

This paper resurrects the analog approach to modeling behavior. The model described here uses few and simple computations. Inverse calculations are replaced by feedback processes in which no great accuracy is required. Behavior is produced not by specifying detailed trajectories of movement, entailing specifications of derivatives that must be consistent with each other as well as with the desired result, but by the generation of variable reference signals that specify positions. Instead of one large complex computation, the model uses a hierarchy of control in which no one control system has to do anything very complex.

The description of the "Little Man" program on the web site is as follows:

This program demonstrates a model of pointing. You can move a target around in three dimensions while a model person reaches out to touch it, following it as it moves. There are five lower-order control systems and three higher-order control systems that run the model; you can alter the basic parameters of all systems to see the effects. The model person uses binocular vision in three dimensions to detect depth information; all visual information is computed from reasonably accurate optical computations that calculate the finger and target angles that each eye will see.

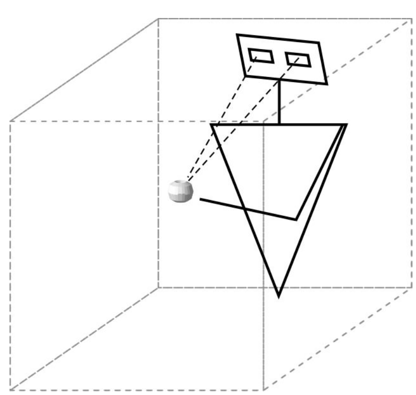

Figure II.12.1. Powers's "Little Man" or "Arm version 1". The Little Man "sees" with stereo vision and tracks the target by tilting and rotating his head. The "visual" data determine his perceived difference between the locations of the target and the "fingertip", which he controls with a reference value of zero.

As the target moves around, the Little Man's "finger tip" follows it closely but not exactly. To take an extreme case, if the target jumps instantaneously from one place to another, and the Little Man's control units have output functions that are leaky integrators (the configuration most often used for simulation in the PCT community) with appropriate rates for the integration and the leak, then the Little Man's fingertip will follow an exponential approach to an asymptotic position close to the target but remaining in error in the direction of the earlier position of the fingertip. If the target jumps to another new position during this exponential approach, the fingertip will approach this new position asymptotically in the same way.

The overall result is that as the target jumps around, the fingertip will trace a path inside a polygon that has the series of target positions as its vertices. Figure II.12.2 illustrates this relationship between target movement and fingertip movement as a time-space graph the trace in one dimension following a single step of the target, and in 2 dimensions as a trace of the target and fingertip movements over time with three steps of the target.

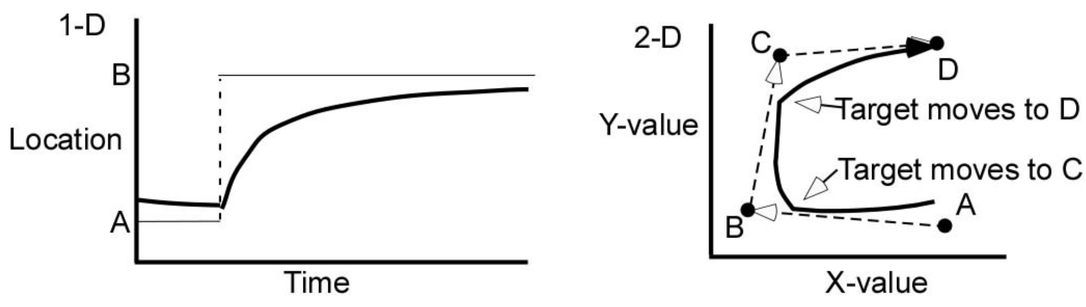

Figure II.12.2. Control is neither instantaneous or exact. It takes time for the "fingertip" to approximate the new position of the target. (a, left) The asymptotic approach to a new position as the target jumps from "A" to "B". (b, right) The trace followed by the fingertip in two dimensions as the target jumps from "A" to "B", and then jumps to "C" before quickly jumpimg on to "D". The concept illustrated by the 2-D trace is easily extended to three or more dimensions.

Control is never instantaneous or exact, in principle or in practice. Depending on the details of the control system, after a step change in the target the tracking "fingertip" may overshoot and oscillate around its final position, or it may follow some other trajectory. No control system can bring the error to exactly zero except momentarily by chance, and wherever the final position may be, the control system cannot bring the fingertip there instantaneously. There is always a trajectory of some kind as the "fingertip" moves. This pair of facts about control are at the heart of the proposals below about how languages and cultures evolve and split apart.

Since the Little Man contains only approach control units, one might ask whether the tensegrity construct is applicable to it. At the level of controlling the finger tip location relative to the target, it isn't. The Little Man controls three relationship perceptions in a three-dimensional space. What happens to one of those control units has no effect on any of the others, because apart from the target, the environment has no content. If there were an obstruction in the environment, then the fingertip and the arm as a whole would have to avoid it, and left-right movements would affect how up-down and in-out perceptions would have to be controlled. With an empty environment, the x, y, and z reference values for the fingertip location can change independently.

At the implementation levels, however, there are control units that bend the elbow, move the shoulder joint, change the head angle, and so forth. These involve tensegrity, as illustrated in Figure 8.7, because the lengths and achievable angles of the structure of the Little Man are aspects of the environments of those control units. Moreover, those control units do affect each other. For example, the orientation of the upper forearm interacts with the elbow flexion, and the movements of the Little Man as a whole reflect what we cited above:

> *The essential function of tensegrity structures is typically to define a shape, which can be quite complex, but by adjusting the tensions in the wires, motions can be created throughout the structure. Robots using controlled variations in wire tensions to create movement have a variety of useful characteristics (Piazza, 2015)*

II.12.2 Syncons and the synx

Imagine now a set of "syncons", locations in a 3-D space defined by three features that can take on values between zero and unity. A syncon then is a location in a unit cube. A syncon could be a location in physical space, as in the case of the Little Man's targets in Figure II.12.2, but the space might be more abstract, defined only by its three specified features. For example, in the "*Story of Rob and Len*", that occupies most of Chapter II.9 and Chapter xII.15, a syncon will arbitrarily be defined by the fundamental frequency, pitch height, and pulsation rate of a buzzy sound.

The more features define a syncon, the easier it is to discriminate among large numbers of them. Appendix 6 shows how rapidly an increase in the number of available features can increase the discriminability of syncons. With 27 syncons, their average pairwise separation in a 3-feature space can be as much as 13 times their average separation on any one feature.

Now imagine an active entity we can call a "syncontroller" analogous to the Little Man. The syncontroller can perceive the location of the active syncon, and in the same feature space can move a "synx" that corresponds to the Little Man's fingertip location. A synx differs from a syncon in that a syncon is a fixed and identifiable location in its space, whereas a synx is mobile and can move anywhere in the space. For the syncontroller to move the synx from one syncon location to another will take time. The space may contain many syncon locations, but no more than one syncon is "active" at any moment, and any one controller controls only one synx.

So that we can talk about them, individual syncons may be given arbitrary single-letter names, in the same way as we labelled the target points in Figure II.12.2 (right panel). Using their letter names, their activity time sequence can be represented by a string of letters such as ABBCDDDD, which is the sequence shown in Figure II.12.2, or ABACABDAB.. which has a kind of syntax (A is followed by B except when it has immediately followed B).

The syncontroller perceives the difference between the locations of the synx and the active syncon, and controls its perception of the difference between the active syncon and the synx with a reference value of zero. If control were perfect, the synx would always be at the location of the active syncon, but control is neither instantaneous nor perfectly accurate; so the synx chases smoothly around the feature space, inside the polygonal trace defined by the sequence of active syncons, in the same way that the Little Man's fingertip follows the jumping target inside the polygonal trace defined by the target locations (Figure II.12.2).

That is what we could do if we were to engineer a control system that would incidentally identify which syncon was active by monitoring the movement of the synx, in the same way that Oliver measures the weight in the rock pan by monitoring the weight in the scale pan. But we are not going to do that. We will be asking how a reorganizing control system might develop a multi-level structure in the process of learning to track the movement of activity around a set of syncons, if the activity pattern is structured.

Powers's Arm-2 demonstration (Powers, 2008) shows that multiple control systems can reorganize to coordinate well in a space of as many as 14 dimensions, but this is done by varying the parameters of a predesigned structure of ECUs. We want our systems to create new types of perceptions that they then control, in the way that the "chair-level controllers" of Figure 4.15 are new types of perceptions built from perceptions of the different parts of the chair that always move together.

What new perceptions should we expect? To anticipate the conclusion, we might expect that perceptual functions should be developed that produce categorical outputs corresponding to the different syncon locations used by the source of the inputs to the perceptual function, and above them, recognizers for types of trajectories that might be categorized into a form of syntax. We follow this question particularly in the extended stories of Rob and Len, and of the rest of Len's family and descendants starting in Chapter II.9.

Chapter II.13. Growing a Language

In Chapter II.9, we considered the first interactions between a baby and its mother. Now, we look at a later stage, at which a child begins to learn the protocols of its culture. To divorce this process from the specifics of any particular culture, we start with an artificially formal situation in a syncon world and investigate how it should be expected to evolve into an informal culture for which there are no precise rules, only approximate ones that may be violated on occasion. We will return often to this artificial environment.

II.13.1 The story of Rob and Len[58]

We imagine a preprogrammed robot called “Rob” whose programming contains many protocols, and a newborn organism called “Len” who will learn those formal protocols sufficiently well to use them successfully in living with Rob. Len will learn Rob’s protocols and a “language” consisting of sequences of syncons (which were described in Section II.12.2 of the “Geometric Interlude”). Later, we will introduce new, younger members of Len’s family, and in Chapter xII.15 we meet descendants in later generations of the family, who have learned the language from their older living relatives. After a few generations the artificial “Len-Rob” language loses its strict formality and evolves into a kind of “natural” language with many conventions but no absolute rules.

When Rob, the robot, “speaks”, it produces only discrete syncons. Each syncon is defined by a specific set of values of three features in some space of description. To make the idea concrete, let us say the features are auditory, and that a syncon is presented as a pulsating buzzy sound with a fixed long-term average intensity when Rob is “talking”. A buzzy sound has a fundamental frequency with many harmonics. The three features are the fundamental frequency, the pitch height (the frequency region where the harmonics are strongest), and the pulsation rate. The features could equally well be gestural, defined by such features as the height of a hand, the upturning or downturning of a hand, and the distance between two hands. Ut really doesn’t matter what the features are, and we will continue with buzzy sounds that partially emulate the vowels of human speech.

Rob can perceive each of the three features and set them to arbitrary reference values, which is to say that it can change its synx, analogous to the Little Man moving its fingertip. Rob could in principle produce sounds anywhere in the feature space, but has been programmed to shift its synx instantaneously from one syncon to another with no intervening trajectory. When Rob’s synx is at a particular syncon, which we call the active syncon, it emits either no sound or the sound that is defined by the three features of that syncon. Rob does not *control* its synx, it *commands* the synx values it wants, and the synx instantly changes to those values.

Rob’s synx movement produces a language of discrete syncons, rather than of arbitrary feature value sets, as shown for two feature dimensions in Figure 6.9 (repeated here as Figure II.13.1). The syncons are shown as distributions around the target values to indicate that although Rob produces the sounds exactly, Len has some uncertainty in the perceptual value of each feature, at least from the Analyst’s viewpoint, since he is unable to discriminate small differences in any of the feature values.

58. This section is loosely based on part of a talk given at the Control Systems Group conference in 1993 (As of 2019/01/11, a video is available for download at http://www.mmtaylor.net/PCT/Movie/TaylorCSG1993.mp4; the relevant portion starts 19 minutes 15 seconds into the talk).

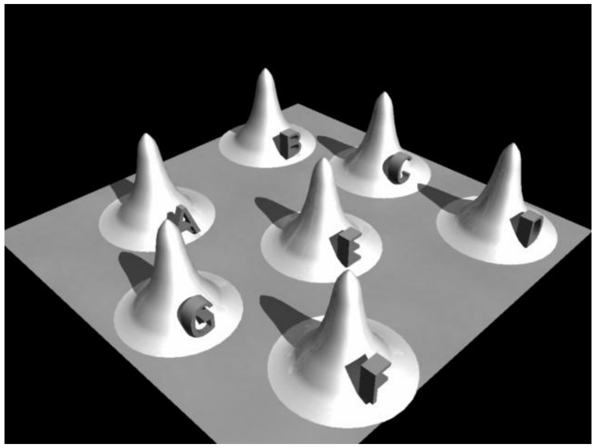

Figure II.13.1 Two dimensions of Rob's syncons in a 3-D display. The spreads around the peaks represent the listener's inherent perceptual uncertainty around the actual feature values.

The newborn Len can perceive and can produce acoustic waveforms that vary over the three feature dimensions. In addition to the fundamental frequency, pulsation rate, and pitch height, Len can vary the loudness of his output using a fourth function, "L". For Rob, "L" is irrelevant, because the loudness Rob will perceive will also depend on how far away Len is when he speaks.

The range of values of perceptual features that Len will learn to accept as a particular syncon, such as "C" in the Figure, will define a macrostate for his linguistic perception. Each of these macrostates extends over all possible values of the "L" feature from inaudible to destructively loud. The Figure shows seven of many such macrostates. If Len hears a microstate (set of feature values) not in one of those macrostates, he will learn to perceive it as irrelevant noise. Likewise, when Len produces values for a microstate not in one of the macrostates as Rob perceives them, Rob will not perceive Len as having "said" anything[59].

Initially, like a newborn baby Len has no control over the individual feature values he perceives from his own outputs. Each perceptual input other than overall loudness is linked with some random weight to all the feature output functions, as suggested in Figure II.13.2. Initially the deviation of Len's perception of pitch height from its randomly varying reference value will affect his output of all three acoustic parameters, as will both the other error values. Len simply makes noises that vary all over the feature space, and there is no obvious reason for his (dis-)organization to change.

59. . Just as we cannot perceive any aspect of the environment for which we have not created a perceptual input function.

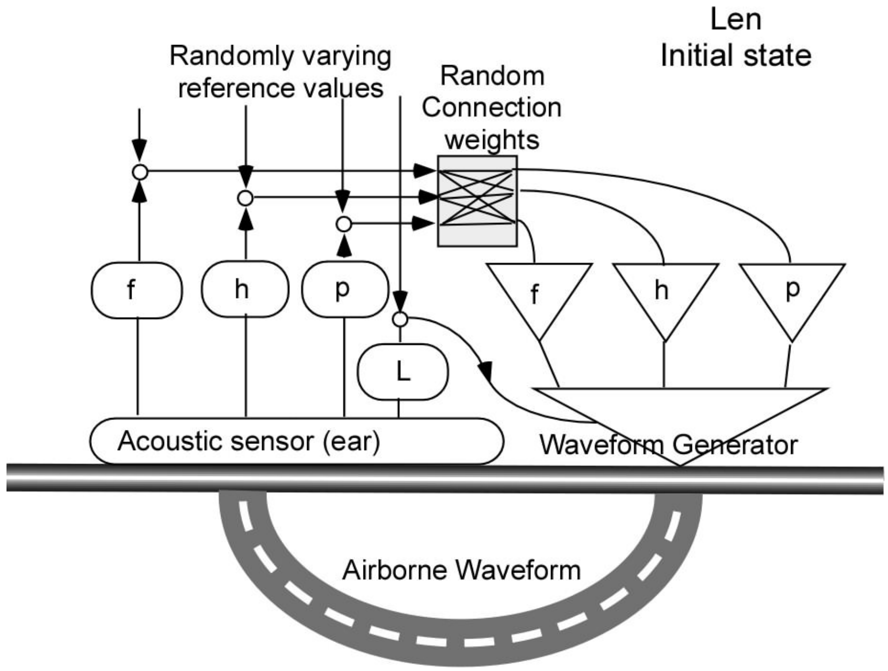

Figure II.13.2. Len starts of with the ability to output waveforms of variable frequency, pitch height, and pulsation rate, and to be able to perceive the frequency, pitch height, and pulse rate of input waveforms, but since each of the input perceptions affects all of the outputs by an initially random amount, Len cannot control his feature values. With reorganization, Len acquires the ability to control by pruning the connections that connect unrelated perceptions and outputs.

Good control is neither better nor worse than no control unless some important variable is influenced by the quality of control (QoC). In an evolutionary sense, however, QoC must act as an intrinsic variable, because if control is poor, the side-effects of control will be inconsistent, even if the mechanism of control is invariant. Since it is the side-effects of control that influence the intrinsic variables, reorganization will proceed more rapidly when control is poor and intrinsic variables are poorly maintained than when control is good. The effect on reorganization is exactly as though QoC was itself an intrinsic variable.

Although Powers apparently did not make this argument, he did use QoC as an intrinsic variable, which allowed him to demonstrate the effectiveness of reorganization for a 14-degree of freedom control system originally organized randomly, as is the 3-D control system in Figure II.13.2 (Powers, 2008). It is therefore not unreasonable to expect Len quickly to reorganize so that the fundamental frequency perception is linked to the fundamental frequency output, and so forth, diminishing and eventually pruning away the cross-links that generate cross interference in controlling. The result is that Len's organization soon comes to look like Figure II.13.3. Len's reference values still vary randomly, but now Len can make whatever noises he chooses ("chooses" means, of course, "setting the reference values for the three features").

In real life, it seems that auditory feedback is needed if babies are to babble and develop the kind of control suggested in Figure II.13.2 and Figure II.13.3 (e.g., Koopmans-van Beinum et al. 2001: *"Detailed analyses of early vocalizations of deaf and hearing infants revealed that auditory feedback is needed to lead to coordination of movements of the phonatory and the articulatory system, and that this coordination capacity is a prerequisite for the development of normal speech production."*).

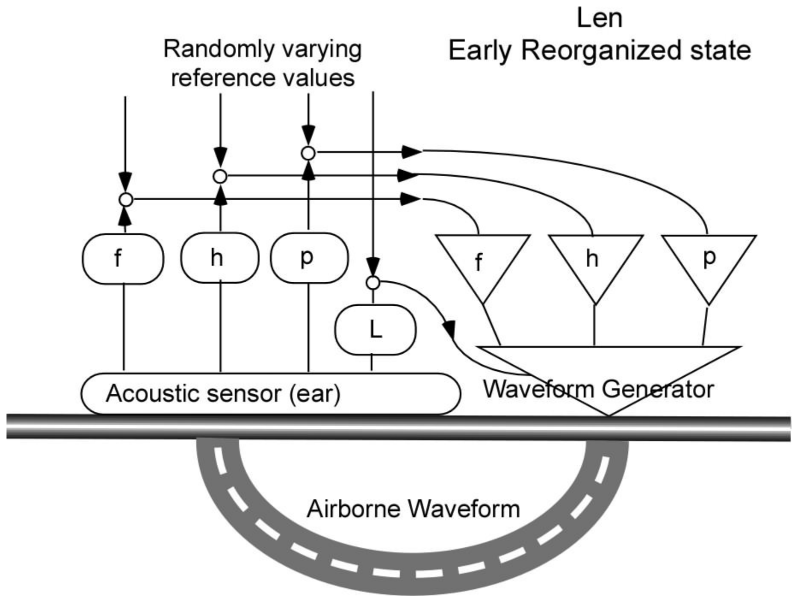

Figure II.13.3. As Len reorganizes under the influence of quality of control as an intrinsic variable, the connections between perception and output are pruned so that error in the perception of fundamental frequency affects only the output fundamental frequency, and so forth.

Len lives with Rob for some appreciable time while his perceptual control apparatus matures. Len makes his noises, but when Rob is "speaking", the robot generates only those feature patterns that specified by syncon locations in the feature space. Accordingly, when Rob is "speaking", Len hears much more consistency in the patterns of features than when he makes his own noises with their random reference values. When they are both "speaking" at once, Len's ear senses sounds with two fundamental frequencies, two pitch heights, and two pulse rates, which together allow for eight different combinations of three features. Despite this, the combinations he hears still concentrate at the locations of Rob's syncons (as in the right panel of Figure II.13.4, because those are the only patterns Rob produces, whereas Len produces patterns widely distributed over the feature space, as in the left panel of Figure II.13.4.

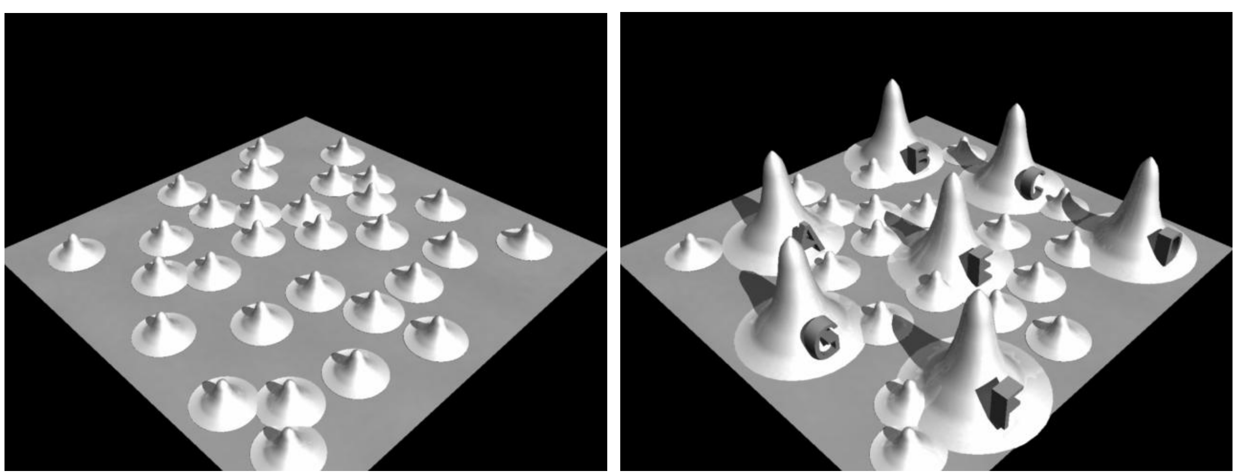

Figure II.13.4. A 2D illustration of how the consistency of Rob's output feature patterns overwhelms the statistics of the feature patterns when Len and Rob both produce sounds heard by Len. The height of a cone represents the recent likelihood of encountering a particular grouping of feature values represented by its x-y location on the base plane. In three or more dimensions, the disparity is considerably greater. (Left panel) A few of Len's outputs. (Right panel) Len's and Rob's outputs as heard by Len. The figure shows only seven syncons for Rob, whereas we will later allow Rob to use about 27 in the 3D feature space.

If HaH processes play their expected part, Len will develop a new kind of perceptual function that takes as input the features labelled f, h, and p in Figure II.13.2 and Figure II.13.3, but not L (because L varies widely and independently of the other patterns, depending on, among other things, how far Rob is away from Len). Len, or Len's descendants, may, however, begin to use systematic variations of L in their protocols[60]. According to the HaH hypothesis, new mutually inhibiting perceptual functions that have inputs from all three features should be produced in Len, one corresponding to each of Rob's syncons. We can call the array a syncon category recognizer.

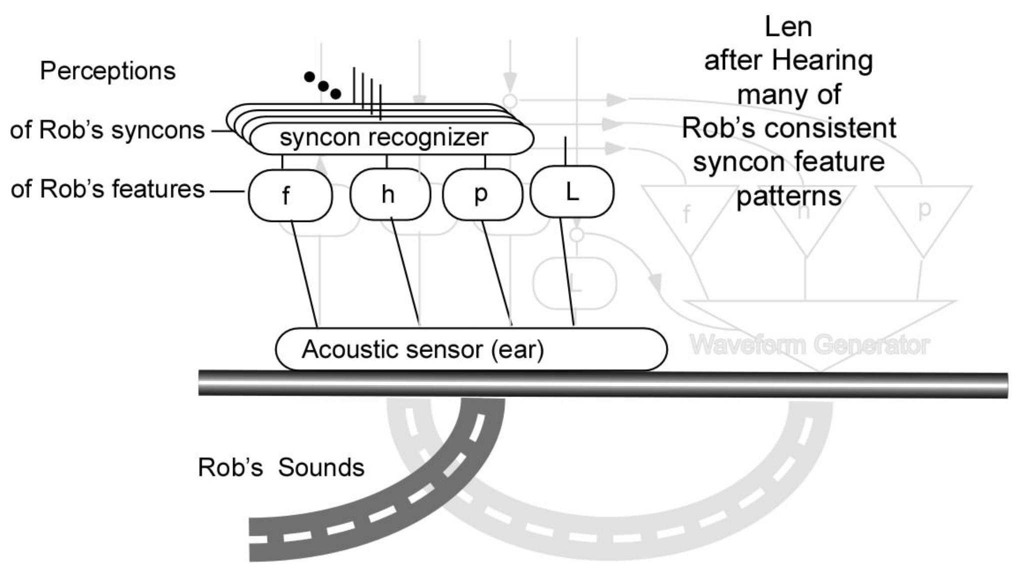

Figure II.13.5. The HaH process creates labile recognizers for the consistent patterns of features heard from Rob, whether or not Len emits consistent feature patterns. At this point the outputs of Len's syncon recognizers are not controlled perceptions, but they are perceptions.

The individual syncon recognizer elements will be tuned by the co-occurrence of triplets of feature values that are the syncon microstates of Rob's output. The situation is analogous to the development of "chair-level" perceptions out of "chair-part-level" perceptions that we discussed in Volume 1. Such tuning, however, does not mean that the exact triplets are the only values that will result in an output from the syncon recognizers. Depending on the structure and base activity level of the recognizers, the nearest syncon recognizer might produce an output for any input that has a harmonic structure, a defined pitch height, and pulsates in intensity, or it might be necessary for a fairly exact triplet to occur before any of

60. Since Rob is not designed to vary L, the only variation Len will hear in Rob's output will be the effect of changing their locations in physical space. Such variations will not easily contribute to "meaning" in the development of protocols.

the syncon recognizers produce output. Such considerations determine the sizes of Len's syncon macrostates, both for perception and for production.

One way of describing a syncon is by tabulating the relative importance for discrimination of the three features and the central values that define them when they are output by Rob, the authoritative source. A "parallel coordinates" feature profile representation is suggested in Figure II.13.6 for an arbitrary three possible syncons labelled A, B, and C. In the figure, A and C have much the same fundamental frequency, but B has a very different fundamental. Fundamental frequency would therefore be important for Len to distinguish between B and either A or C, whereas pulsation rate would be important in distinguishing A from either of the others. If Rob were to vary some of these feature values, as Len and his successors are likely to do, Len would be able to tolerate considerable variation in the fundamental frequency of B or the pulsation rate of A without misidentifying the syncon. Such profiles as these may be instrumental in the drifting and splitting of natural languages, but for the moment we are dealing with Ron, who controls his synx precisely, and Len, who is only just learning to control his synx.

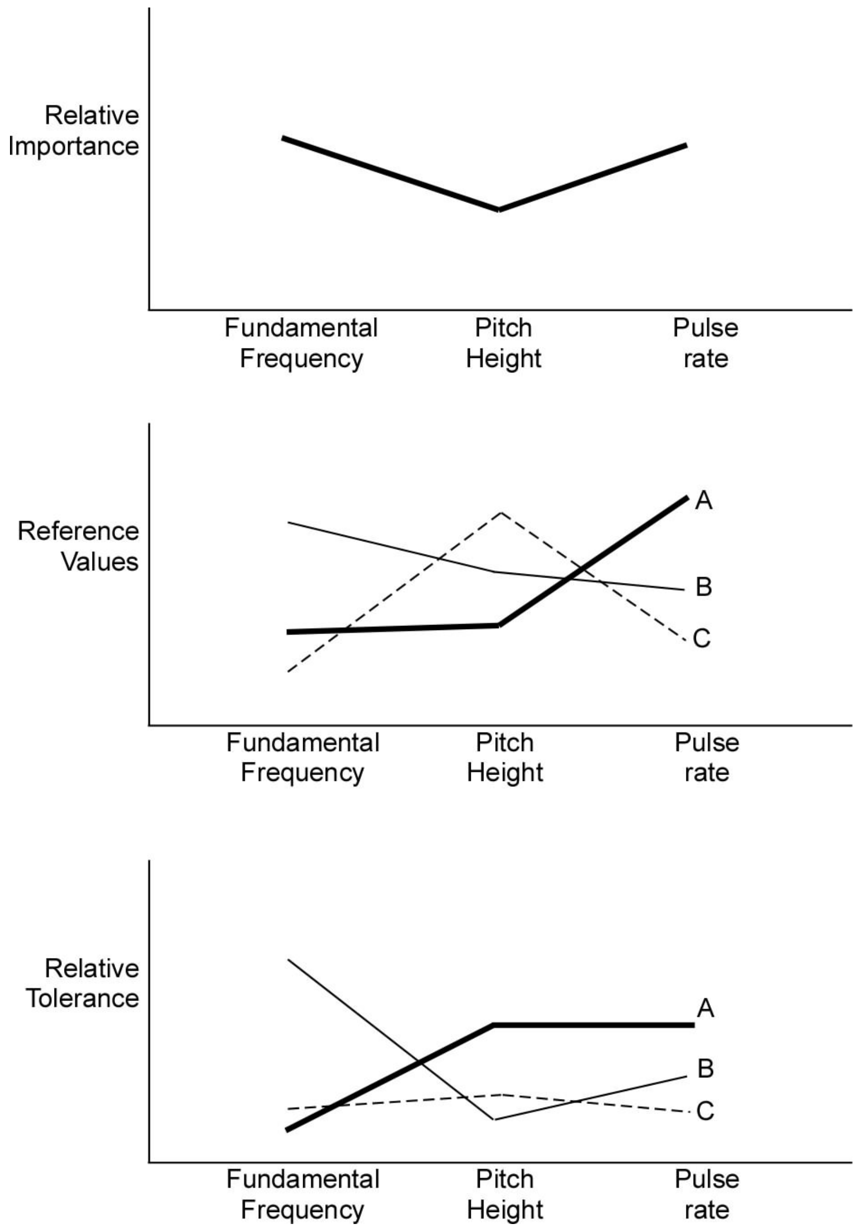

Figure II.13.6 Feature profiles for the perception of syncons. The upper panel suggests the influence of each feature, while the middle panel shows the reference values appropriate to some possible syncons, and the lower panel suggests relative tolerance limits for the different features before the pattern ceases to be perceived as, for example, "a good B". The "relative importance" profile shown is an average, as the actual importance must be different for each pairwise discrimination. For example, "Fundamental Frequency" (f in Figure II.13.2 and successors) is of little importance in discriminating A from C, but is important in discriminating B from either A or C.

II.13.2 Baby Len babbling syncons

When little Len produces his noises, occasionally they will form a pattern that is sufficiently close to one of the syncons to produce an output from a syncon category recognizer element in one or both of Len and Rob. As yet, Len is in no position to produce a synx (a sound pattern) that approaches a syncon he chooses, since the new syncon recognizer perceptual function has no connection to an output process. They are new perceptions, but they are not controlled, and if Rob goes away and Len hears only his own outputs, the syncon category recognizer functions will soon dissipate as the HaH processes continue to operate without the statistical consistency of input provided by Rob's sound patterns.

However, just as the random connections reorganized to allow Len to control the three features independently, so the uncontrolled perceptions of syncons should begin to become preferentially connected to particular patterns of feature outputs — analogous to chairs rather than chair-parts — allowing Len to begin to approximate control of syncons as well as of features. Even if Len's synx originally drifted randomly throughout the 3-D feature space, it will soon begin to be preferentially located near enough to the locations of Rob's syncons so as to produce outputs from his (Len's) syncon recognizers, and perhaps from Rob's. Len's resulting output will have many of the qualities of the babbling of a baby, which soon begins to cluster around the intonation patterns of the language of its social environment (e.g., de Boissons-Bardies et al., 1984) in preference to the those of other languages.

The all-seeing Analyst realizes that Len needs to implement the associative memory implied by Powers's Figure II.11.1 in B:CP, so that when Len wants to produce a synx corresponding to syncon "X", the correct pattern of feature values will be provided to the reference inputs of his feature control loops. Would the necessary connections occur naturally?

In Section 8.4, we said:

> *With HaH, we might expect to see the development of neurons that show "association", meaning that if a particular pattern of inputs recurs, the synapses that were strengthened the first time will be further strengthened, and those that were weakened before will be further weakened. […] If this happens, the postsynaptic neuron will be increasingly likely to fire when only a subset of the "associated" inputs fire, providing a perception of the entire pattern even though some of it might be missing or deviant.*

The HaH process tends to create associations on the perceptual input side of the control hierarchy, and the syncon category recognizers could be seen as content-addressable memories, addressed by the f-h-p triplets of values to produce a syncon identity recognition. Could the process, without forcing, be expected to produce the reverse effect, taking the syncon identity and producing an appropriate f-h-p triplet? The answer seems likely to be "Yes". The polyflop structure suggested in Figure 8.14 implements it if the figure is inverted so that the upgoing data paths in the figure are seen as downgoing reference output paths that address associative memories that output the required feature set. The positive feedback that creates "labels" on the perceptual side of the hierarchy in Figure 8.14 now becomes the means of producing coordinated action of lower level control units on the output side without conflict.

Len's synx will initially match Rob's syncon locations only occasionally, by accident, but at some point Len, whether by the HaH process or some other mechanism, becomes able to control his synx to approximate Rob's syncon locations in feature space fairly well. The polyflop structure may have developed enough to begin the fold structure of Figure 8.13. Len can now produce sounds having feature patterns that produce much more output from the desired syncon recognizer than from any of the others. If his synx is close enough to a syncon location that only one of his or Rob's syncon recognizers produces appreciable output, his control of syncon identity is perfect, even though Len's "voice" or "accent" will be easily distinguishable from Rob's.

A consequence is that Len has become less able to control the individual features independently. To do so would conflict with control of perceiving his synx as a producer of syncons rather than just of individual features. If Len meets someone else who uses a different set of syncons, Len will initially find it difficult to mimic them, just as an adult who moves to a different country may become perfectly fluent in the words and phrases of the new language, but always retains "a foreign accent".

II.13.3 A language of syncon sequences

While all this is going on, Len is also reorganizing to control perceptions of inanimate objects in his external world, for example by picking up or moving various items. We will not worry about how this reorganization happens, but the progress of reorganization must involve some unspecified intrinsic variable or variables, one of which is Quality of Control, that are affected by Len's control actions in his environment.

Side-effects matter, as we saw when we analyzed the first protocol between Cora and her new baby Ivan, though they are usually ignored in discussions of PCT. Unless you are controlling for perceiving someone else to be observing your actions, anything anyone else sees you do is only a side-effect of your control. Language is an obvious case in which you do control for perceiving someone else to be influenced by your action, but until you have established a language of some kind, most of your interactions with other people happen through side-effects.

In order for Len to develop a language, language must do something useful for him, which means that it must enable him to control some perceptions that affect his intrinsic variables better than he could do without using language. In Chapter II.9, we examined how a baby might develop different kinds of cry that its mother could use to determine whether feeding or pin-removal would reduce the crying that disturbs her perception of the baby's happiness. These kinds of cry had meaning to mother Cora. They were the first words of a private language between mother and baby. But now we consider how Len might develop a language full of meanings by interacting with Rob in the syncon-synx environment.

Because Rob is Len's only possible language partner, we must give Rob some appropriate abilities beyond his ability to speak in precise syncons. If that was all Rob could do, Len could control nothing but his speaking voice by interacting with Rob. We already assumed that Rob has been programmed with many different protocols, and now we will assume that Rob can act, as Len can, on objects in the environment. For example, Rob might be able to give Len a ball Len cannot reach. We will further assume that Rob represents a ball by moving his synx abruptly from syncon "B" to syncon "A" to make a word "B^A"(using the caret "^" to indicate an abrupt transition with no hiatus or co-articulation smoothing).

Rob is programmed to give Len the ball only if Len lets Rob know he wants the ball. We assume that Rob controls for Len to reduce error in his controlled perceptions (becoming more contented), as this is the essence of using protocols. Displaying to Rob what he wants is something Len could in principal do by using the language of syncons. If Rob is programmed to say "B^A", when Len has a ball in hand, then by the same associative process described above, Len should be expected to perceive {"ball in hand" + syncon sequence B, A} as an associated entity in the same way that the feature patterns associate together to create syncon perceptions. Eliding a lot of steps in the reorganization process, Len should learn to say the syncon sequence "BA" if he has a reference to perceive a ball in hand.

Of course, Len does not produce "B" exactly as Rob does, and nor does his synx move stepwise between his "B" and his "A". Instead, it follows a relatively smooth trajectory between the two syncon locations, as is shown in Figure II.12.2 for the Little Man following an abrupt shift in the target location. Reorganization will improve the fidelity of Len's "B" and "A", and speed the trajectory between them, but only until the result is accurate enough for Rob to identify the “BA” trajectory as having the same meaning as his “B^A” sequence, and give Len a ball to hold. When that happens consistently, Len’s control of “ball in hand” is successful, and reorganization slows dramatcally, at least in respect of the perceptions involved in that protocol.

Using sequences of two syncons, a sufficient level of accuracy is not difficult to achieve. Figure II.13.7 suggests that whereas “B” and “V” might be hard to distinguish, if Rob never uses and Len never learns “VA”, “B^A” will not need to be distinguished from “V^A”. If Len uttered something that was acoustically closer to “VA” than to “BA”, Rob would still hear “BA” because “VA” is not a sequence known to him. Rob might, perhaps, use the sequence “V^Z” for something, but this will not affect the distinctiveness of “B^A”, at least if “Z” is quite distinct from “A”. In linguistic terms, at this stage of Len’s development of the syncon language “B” and “V” will be allophones of the same phoneme. Later, we will see that these allophones may become separate phonemes, but that has not happened yet for Len.

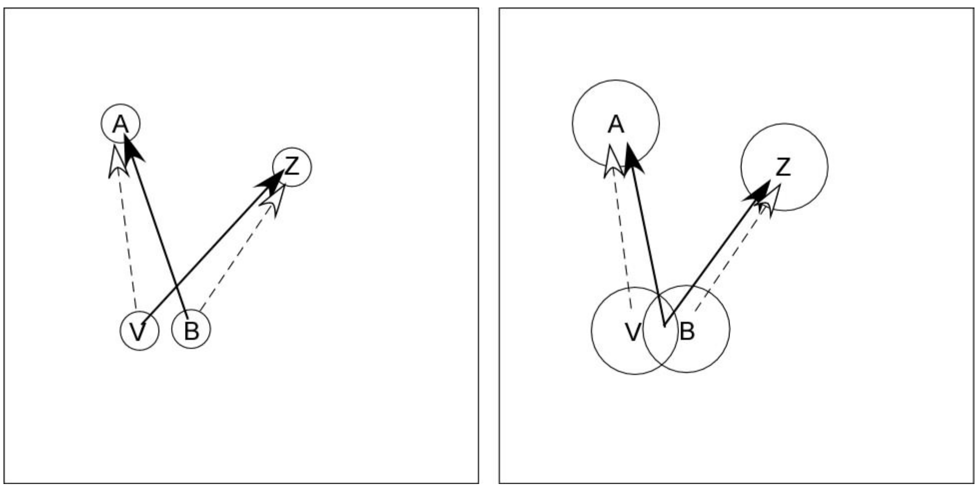

Figure II.13.7. Even though Rob’s B and V might be hard to distinguish, if Rob uses BA but never VA (left), it become obvious that the first syncon was B and not V. So even if Len’s “near enough” B is closer to Rob’s V than to Rob’s B (right), Rob would still be likely to hear “BA” and not “VA”. Rob does use “VZ”, so even if Len’s”B” of BA is physically identical to his “V” of “VZ”, Rob will still hear Len as having said “BA” or “VZ”, rather than “BZ” or “VA”.

Len is, as pointed out above, physically unable to “speak” his syncon sequences the way Rob does. Whereas Rob must switch his synx instantly from one preprogrammed syncon location to the next, Len has no such switch, but instead uses perceptual control of what he hears or imagines himself “saying”. Len’s control systems require time to bring his perceptions to their new reference values up and down and across his control hierarchy, as suggested in Figure II.13.8 for one of the three features of his and Rob’s synxes. Rob’s synx, on the other hand, always coincides with one of his syncons.

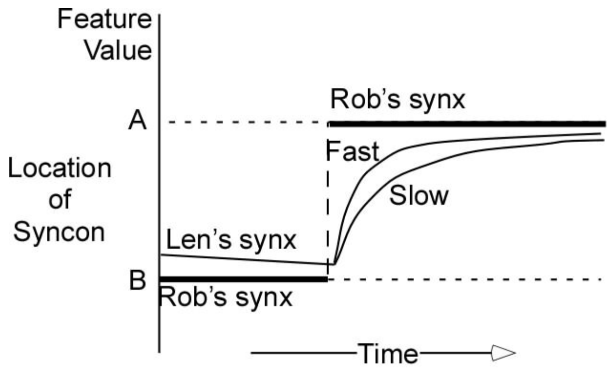

Figure II.13.8. Len's synx changes smoothly, whereas Rob's changes instantly between the locations of their B syncons and the locations of their A syncons. Len's transitions may be fast or slow, and this rate could come to act as another feature if Len talks to a listener who could discriminate the transition speeds.

This "disability" turns out not at all to be a disability on Len's part, because it gives him a flexibility unavailable to Rob, to use his control trajectories as features in higher-level perceptions, modulating, or perhaps even entirely changing, the meanings of particular syncon sequences, as we shall see later.

Another person perceiving Len's synx movements might be able to use Len's smooth transition trajectories *in addition to* the relatively stable syncon target locations, in identifying which syncon sequence Len is trying to produce. The trajectory provides information might otherwise be lost because of Len's inability to match and hold the target locations precisely. In natural language, such features come under the general heading of "prosody".

After a while, "BA" would only be one of many syncon sequences (perhaps including "VZ") that Len might have learned from Rob. Between them, Rob and Len would also have developed some protocols, including a protocol initiated by Len we could call "Give me X", in the way baby Ivan and Cora developed a protocol which Ivan could initiate by making the appropriate cry for "Feed me" or "Get this pin out of me".

The difference between Rob's and Len's speech is effectively that between de Saussure's[61] Langue (precise speech, Rob) and Parole (variable speech, Len)[62]. Len's trajectories of synx control moves following reference values that may change stepwise makes Len's movements from one "phoneme" to the

61. it is not possible offer a precise dated reference for this. (From Wikipedia "Langue and Parole": *Saussure did not publish his notes in relation to linguistics and langue and parole. Unfortunately Cours de linguistique générale was published after his death in 1916 (later translated into English in 1959 as Course in General Linguistics) and was made up of remaining lecture notes by Saussure, course notes provided to students and notes taken by former students of his lectures he performed between 1907-1911 in Geneva*.) For those who want to read an English translation, the reference in English given in Wikipedia is "de Saussure, F. (1986). Course in general linguistics (3rd ed.). (R. Harris, Trans.). Chicago: Open Court Publishing Company. (Original work published 1972)."

62. A note on Langue versus Parole. Rob "speaks" in syncons, discrete and invariant, in a well-defined "Langue", whereas Len uses physical components that cannot move instantaneously from one place to another. Len controls for hearing himself make the intended noises, and takes time for the control process. Len's output varies from one utterance to the next—"Parole".

next dependent both on the prior and the succeeding contexts. His trajectories through the space of the three distinct features are continuous, and depend both on where they came from and where they are to go next, whereas Rob's are discrete and independent of acoustic context. Every time Rob says "BA" it will the same "B^A", whereas when Len intends to say BA in different contexts, it will be different in detail.

In Chapter 9 we described two separated hierarchies, one the analogue hierarchy as described by Powers, the other a "logical" or "categorical" hierarchy built on precise identification of perceptions as members of particular categories. Categorically, this perceived colour IS red, it is not a kind of reddish orange. Or it is NOT RED. This observed political behaviour IS or IS NOT democratic. The category interface makes decisions and admits no half decisions.

II.13.4 The story of Len and Sophie

Len now meets Sophie, who is as language-naive he was before he met Rob. Len teaches Sophie the syncon sequences Rob taught him, which Sophie learns from him just as he did from Rob. Rob had a limited repertoire of actions that Len could use to control different perceptions, and Len has created protocols for most of them, which Sophie is likely eventually to learn from him.

Len can now control his synx so that it produces good approximations to Rob's 27 syncons. When speaking sequences such as "BA", however, Len uses smooth transitions to move his synx from one target syncon location to another, and the trace of these transitions is reasonably consistent when moving from one syncon to the next in a given context. If Len uses primarily only a small subset of the 729 possible transitions (allowing for repeats) among the 27 syncons he has learned from Rob, Sophie might be expected to develop perceptual functions that identify not only Len's sequences of syncon target, but his trajectories as well, a yet higher level of controllable perception that Len did not achieve from his interactions with Rob, but that he might develop in his interactions with Sophie.

In natural language, a "word" uses particular temporal pattern, a syntax, of phonemes, just as a phrase uses a syntax of words. For example, many English words start with "bl" and "fl", as in "black" or "flick", but none start with the equally easily pronounced "vl" that occurs in the Russian Vladimir or Vladivostok. If Len uses such a syntax in his syncon language, Sophie should eventually reorganize to perceive it and to control her perception of it. For example, Len is likely to use the "BA" sequence that references a ball in hand, or "VZ" meaning something else, but Len would never use "VA" or "BZ", which Rob never used. Sophie might develop a recognizer that tells her whether something Len seems to say is a "properly constructed" word such as "BA", and if it isn't, such as "VA", perhaps to recognize what he said as an acoustically similar "real" word "BA". But before Sophie graduates to an implicit perception of (and control of) Len's syntax, she must acquire perceptual functions as Len did, for syncons and "words".

Sophie, like Len, must use perceptual control for synx variation, with the result that both create smooth transitions from one syncon to the next through feature space. Using these trajectory variations, she and Len have an opportunity to develop new protocols unknown to Rob, even without changing the repertoire of syncons or syncon sequences ("words") used by Rob and Len. They can add trajectory variation in itself as a new perceptual feature, and that variation might include variation in the "L" feature that Len could not use with Rob because Rob did not produce trajectories. For example, they might distinguish between a smooth "BA" and on in which most of the transition trajectory was silent, creating something like a glottal stop "B'A", or between a quick "BA" and a slow "B..A".

Sophie develops syncon and word recognizers by listening to Len, just as Len developed them by listening to Rob. We can take those for granted, even though Len's synx when producing syncons is much less precise than is Rob's. And just as Len's syncon targets differ slightly from Rob's, so in the same way and for the same reason, Sophie's will be different from Len's. Since, ignoring trajectory variation, syncon targets are defined in a 3-D feature space, or 4D if they come to use variation in "L" as a significant feature, only a small proportion of her syncon targets will be closer to Rob's than they are to Len's[63].

Len has learned that "BA" is associated with the ball in hand. Not only that, but with Rob he developed a "Give me X" protocol, in which his initiating part was to say "BA" when he did not have the ball. So if Sophie by chance happens to output "BA" or something close to it, Len is likely to perceive that Sophie wants the ball in hand, even if she was simply babbling syncon sequences. If Len gives her the ball often enough when she says "BA", or if he says "BA" with the ball in hand, Sophie should, possibly by the HaH process, develop an association between the sound trajectory and the event of getting a ball in hand, or perhaps just with the ball as an object.

For Len and Sophie, "BA" has begun to have properties that in a natural language, we would call the "meaning" of a "word", just as Cora and Ivan developed "meanings" for his different "Ivan-word" cries and gestures. Other transitions would acquire other meanings, private to the two of them, though those meanings would in many cases be similar to the meanings Len had acquired from his association with Rob. Just as with the syncon target locations and the synx transition rates, Len's "meanings" might have different boundaries than had been programmed initially into Rob, and Sophie's "meanings" might differ slightly from Len's. Both Len and Sophie presumably would associate "BA" with a ball, but they might not perceive the same set of objects as being "BA"s. Len might allow a smooth oval object as a ball, whereas Sophie might accept only objects very close to spherical.

Just as Len's target syncons are likely to differ from Rob's, and his variable trajectories will be inside Rob's stark polygon corners of syncon locations, so Sophie's targets will differ from Len's, though perhaps not to the same extent; Sophie is physically capable of reproducing Len's synx trajectories exactly, whereas Len could not make his synx perform Rob's abrupt shifts between active syncons. Sophie's ability to control her perception of variation in her trajectories enhances the likelihood that Len's trajectory variations will come to have meaning for her, and reciprocally for Len. Trajectories can gain meaning in just the same ways as did Ivan's varieties of crying.

Because Sophie's target locations are determined by the syncon sequences she hears from Len, her "V" location cannot be created by hearing "VA", since Len never says it. But he may say "VN", "VQ", and so forth. If Len also says "BN" and "BQ", Sophie will perceive B to be phonemically different from V. If, however, Len never uses "B" and "V" in the context of the same following syncons, her "B" and her "V" recognizers might fail to develop as distinct categories. If the syncon language were a natural language, they would be described as not being different phonemes, just as /r/ and /l/ are not phonemically distinct in Japanese or Korean.

---

63. In any Euclidean space of N dimensions, the distance between two points is the square root of the sum of squared differences over the dimensions. Let's say Rob's target for some syncon is $\{0, 0, 0\}$ and Len's approximation is $\{1, 0, 0\}$. Sophie's approximation to Len's approximation is somewhere in the unit sphere centred on $\{1, 0, 0\}$. Sophie's approximation is closer to Rob's only in the region where Sophie's sphere intersects the unit sphere around $\{0, 0, 0\}$. The volume of the unit sphere is $4\pi/3$, while the volume inside both spheres is $5\pi/12$, so the probability Sophie's syncon is closer than Len's to Rob's original is 15/48, or approximately 1/3. With more features, the probability becomes much lower that Sophie's syncon is closer to Rob's than is Len's. Sophie will acquire Len's "accent" rather than Rob's "formally correct speech".

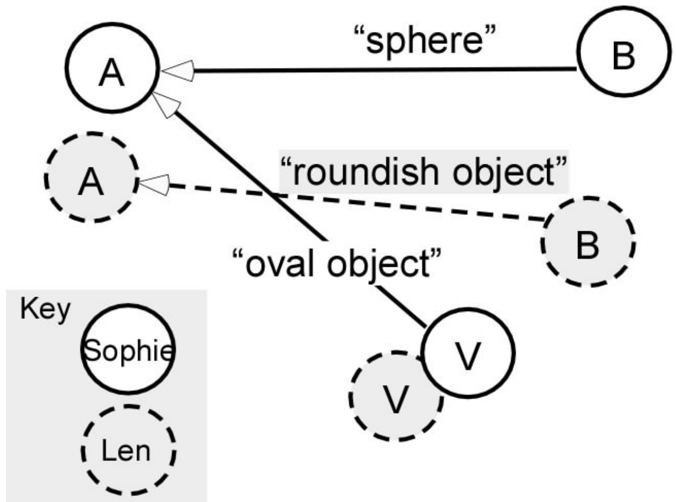

Figure II.13.9. Len never intends to say "VA" but Sophie hears him say "VA" as often as she hears him say "BA".

Figure II.13.9 assumes that "B" and "V" are different phonemes in Len's syncon language that Sophie has largely learned. If Len's B and V syncons are close to each other, it is quite possible that Len's "B" will be equally close to both Sophie's "V" and her "B". If Sophie sometimes hears Len's "B" as a "V" and sometimes as a "B", she will hear him saying both "BA" and "VA", whereas Len never heard "VA" from Rob, and doesn't perceive himself as saying it. If Sophie is controlling for "getting X in hand" and by chance happens to hear "VA" more often when she sees Len with an oval object than she does when Len has a sphere, she may well create a new "word" "VA" to distinguish the two situations.

Because Sophie is controlling different perceptions in a protocol with Len when she uses "VA" rather than "BA", "VA" will come to have a meaning for Len. He will hear "VA" when Sophie says "VA", even though he has never said it himself. The change in categorical language perception comes about and is stabilized because it corresponds to controlling different perceptions of the non-language environment, as J.G.Taylor proposed (J.G.Taylor 1962). Len will be likely to reorganize to perceive Sophie as intending "VA" and "BA" to be distinct words, and might develop a similar (though probably not identical) oval-spherical perceptual distinction among balls. On such small differences do languages drift.

Sophie's experience of hearing smooth trajectories differ's from Len's experience with Rob. Whereas Len encountered only sequences of Rob's precise syncon locations in developing his syncon and syncon pair (word) recognizer functions, Sophie will also encounter Len's trajectories from one target to the next, introducing a timing and a loudness variation perception into the mix, if Len is consistent in the loudness differences between syncons in specific trajectories. For example, in "BA" he might consistently produce the "B" louder than the "A", but in "BR" the "B" might consistently be softer than the "R". Sophie's syntax perceptions would in that case come to include the L property and a new "T" property as syncon features along with the other three, thereby greatly increasing the size of the space available for distinguishing syncon "words". Here we have another instance of interactions creating productivity, much as we saw in the autocatalytic and homeostatic creativity of Chapters 14-17.

Trajectory variation might come to have a more important part to play than simply easing discrimination among syncon sequences. Since Rob used no trajectory variation, Len's variations do not distinguish among the different protocols he used with Rob (Figure II.13.8). But that does not mean his variations are independent of context. If Len uses trajectory variations consistently in different contexts, Sophie should develop recognizers for them, even if only as a perceptual component of the words she learns to perceive. A slow “BA” is more different from a fast “BL” than are Rob’s precise “BA” and “BL”, even if Len’s syncon feature values are only approximated by the end points of his synx trajectory rather than being exactly maintained at the start and end of the trajectory (Figure II.13.10).

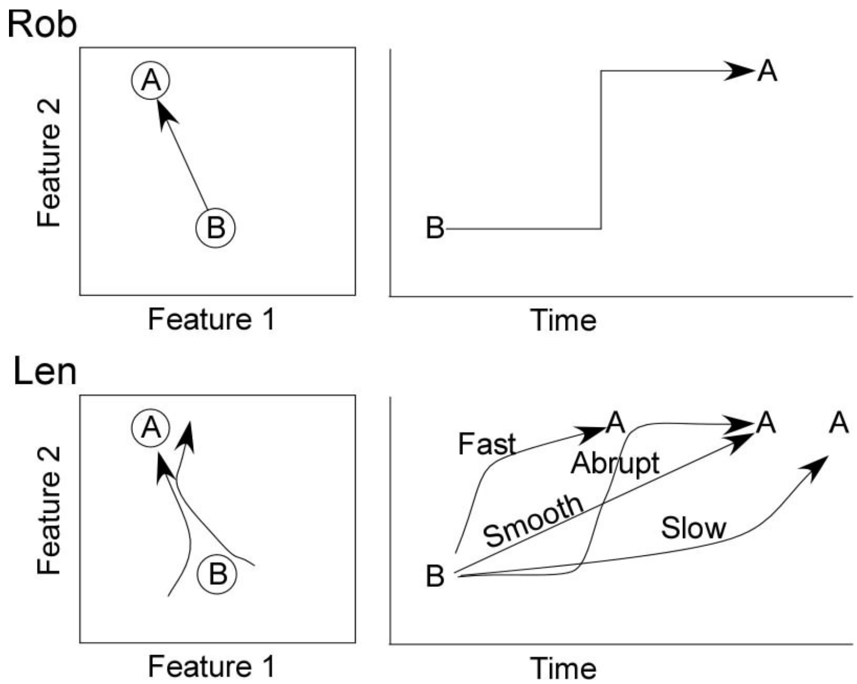

Figure II.13.10 Possible “BA” trajectories for Len as opposed to Rob. Len’s trajectories can be much more variable, both in the path through feature space and in the timing of the move. Whereas Rob’s trajectory changes all the features simultaneously from one syncon to the next, Len’s could be different for each feature. The trajectories in each timing diagram represent only a few possibilities, not timings for a set of different features in a single “BA” transition.

Even if Len’s trajectory variations are initially as random as were baby Ivan’s cries and limb-waving, his interactions with Sophie are quite likely to give them meaning by differential association with the perceptual contexts in which they (randomly) happen to be used, as did Ivan’s cries in his interactions with Cora. For example, Len might say “BA” with one trajectory, such as a fundamental frequency rise and fall, when he holds the ball and shows it to her, and with another trajectory, say a falling-rising fundamental frequency when he gives it to her.

But just as Ivan might have distinct cries for “Feed Me”, all of which result in reduction of hunger because of Cora’s action, so Len might develop distinct ways of saying “BA” depending on some context he (but perhaps not Rob) perceives. Len’s trajectories with even only three features suggest that he might have a vast number of distinguishable ways of saying “BA”. For example, if feature 2 happened to be fundamental frequency, the left panel of Figure II.13.10 might show Len saying “BA” with two different tones, rise-fall and fall-rise, while the right panel might show whether the tone transition was abrupt or gradual, early or late in the transition. All of them would be the same to Rob, but not necessarily to Sophie.

We can make an analogy here to the 1515 dot array we discussed in Chapters 11 and 12. Rob can control perceptions of whether an individual grid square contains or does not contain a dot, whereas Sophie and Len would perceive more possibilities within the array, as suggested in Figure II.13.11. In

Panel (b) they perceive the internal structure of each dot, whereas in (c) they can perceive both where within the array square the dot is positioned, and a "flow" of smooth curves through the array squares containing dots, analogous to the feature trajectories we have been discussing. In our natural world, we might say that the three panels showed the same letters (gross macrostates) in different fonts (refined macrostates within the gross macrostates).

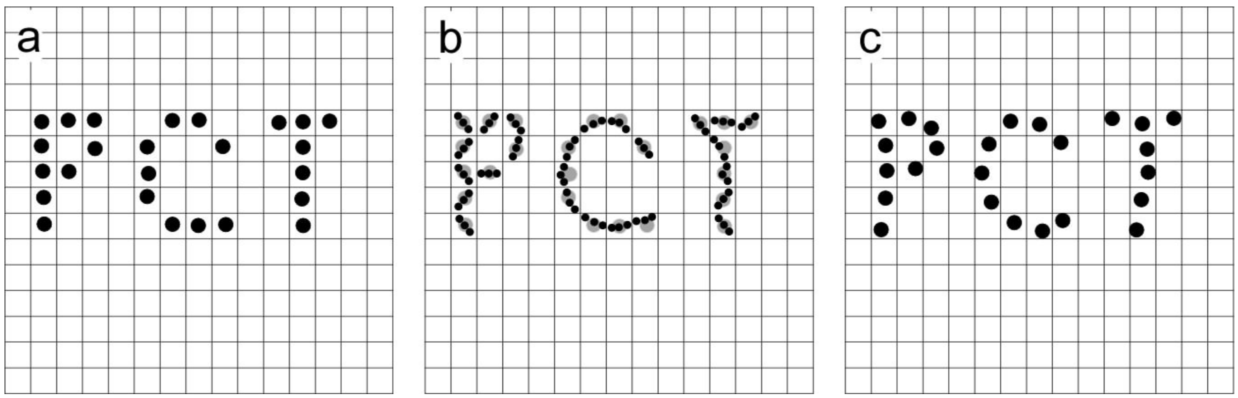

Figure II.13.11 A dot pattern as might be perceived (a) by a "Rob" who controls only the gross locations of the dots in the matrix and (b, c) by a "Len" or a "Sophie" who can see finer connections and structure within and among the dots inside the squares. In (b) they can see that the dots are constructed from triples of smaller dots, and in (c) they can see where the dots are within the grid squares. Panels (b) and (c) are two of a very large number of possible microstates that belong to the macrostate defined by the dot locations of panel (a).

Suppose we ignore the internal structure of the dots illustrated in Figure II.13.11b, and restrict to nine locations Len and Sophie's ability to perceive and control where the dot lies within its square (top-left, centre, middle right, and so forth in Figure II.13.11c). How does their refined ability affect the information they could transmit using the 1515 dot array with 25 dots as a medium? In Chapter 10 we found that the uncertainty of a universe of 25 dots in a 1515 array was 105 bits. Now we add $\log_2(9)$ = 3.17 bits for each of the 25 dots, or 79.2 bits in all, giving Len and Sophie 184.2 bits from which to construct their variety of structures, among which might be "font varieties" within the "three letter" sub-Universe. That sub-Universe allowed them 13 bits, to which we can add the same 79.2 bits for the dot locations in their squares, giving them about 82 bits available for communication, which is probably more than they would have in a "three object" sub-universe of dot occupancy for the cells.

Identifying the three-letter sequence still is done within the 225-bit universe of dot occupancy, because if that is all they wish to communicate, the more refined macrostates of Panels (b) and (c) become irrelevant, but they could communicate much more, if they want to take advantage of the "font variation" universe of dot locations within their squares, and enormously more if they want to use the dot microstructure as well.

Transfer this approach to the three-feature syncon universe. At Rob's level of perception, there are 27 syncon locations and 729 two-syncon trajectories, which provides Rob and Len with the possibility of transmitting about 9.5 bits per word. If Len and Sophie could distinguish just two levels of trajectory speed and five variations of each feature (level, rising, falling, rise and fall, fall and rise)[64], they would have at their disposal 729×2×5×5×5 = 182250 two-syncon "words" instead of 729, or about 17.5 bits per "word". These extra eight bits would be available for redundancy that would enhance the correctness of

64. . In a natural language, these would be "tones" if the feature is fundamental frequency.

interpretation of imprecisely "spoken" words and for conveying contextual information such as the importance to the speaker of what was said — or for both.

The point of this little exercise is to highlight the enormous gain in communicative capacity Len and Sophie have available to them because of their ability to vary the feature trajectories when moving the synx from one syncon to another, as compared with the capacity that Len had in his initial interactions with Rob, who could only make abrupt transitions of his synx between fixed syncon locations. The question, which almost answers itself, is whether Len and Sophie are likely to use their expanded universe of possibility. We argue, of course, that they will.

II.13.5 Len and Sophie have a child

A child, call it "Dan", growing up hearing only Len and Sophie as sources of "voice" other than his own output will naturally come to perceive the syncon patterns they use, in the same way that Len learned from Rob and Sophie from Len. But since Len's and Sophie's syncons differ somewhat, the statistics for Dan are less precise than were those from which Len or Sophie developed their syncon recognizers. Dan's Perceptual Input Functions (PIFs) for protocol-relevant perceptions are likely to be less sharply tuned than are Len's, who learned only from Rob, or even than Sophie's, who learned only from the less precise Len.

Dan's syncon and trajectory macrostates occupy a larger proportion of the global space of possibility than Len's or Sophie's, though still a very small portion of it once his "speech" begins to be intelligible to Len and Sophie. But now their space of possibility is very much larger than was Len's just after learning from Rob. The redundancy gives Dan a larger leeway for deviation from the region that encompasses both their syncons than either of them had when they were first learning. Nevertheless, he still must get them to recognize the he is approximating a "B" and an "A" if he wants a ball, rather than showing by his prosody that he is failing to control something unidentifiable. Referring back to Figure II.13.9, Dan's "V" and "A" recognizers are likely to be tuned to accept both Sophie's and Len's versions, but what of the rather different "Len's B" and "Sophie's B" (Figure II.13.12)?

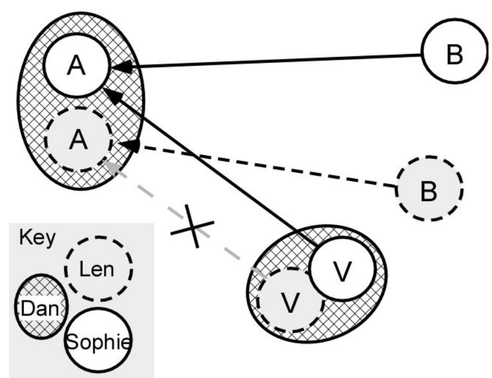

Figure II.13.12. Dan is likely to have more widely accepting recognizers for A and V than does either parent, but what of his recognizer(s) for B. Does he have a single very wide acceptance region or a separate one for each parent?

The answer is that either of two things may happen. Thinking outside the Len-Sophie-Dan "syncon-speaking" community for a moment, in the real world someone exposed to a variety of different accents does learn to accommodate a wide range of forms for many phonemes, as one learns to accommodate voices of different pitch ranges. However, one recognizes not only that this word was "dog", but also that it was spoken by a male or a female, by someone from a different linguistic background, or by a particular person.

In Figure II.13.9 and Figure II.13.12, Sophie's syncon locations are characteristically shifted from Len's, all having a higher value of the feature represented by the "y" axis of the diagrams. Such a consistent shift might easily be implemented in Dan's syncon recognizers, which would associate not only the feature patterns but also the perception of who was "speaking". High "y-feature value" means "Sophie speaking" if Dan recognizes "VA", or "VA" if Dan perceives that the speaker is Sophie. This could be seen as a separate message ("It's Sophie speaking") multiplexed onto the syncon trajectories in addition to "Give me a ball" or whatever protocol was being used.

Recognizers take form initially from the consistency of feature associations in the form of patterns or profiles. If one version of the feature pattern for B is always accompanied by the presence of Len but only sometimes by the presence of Sophie, we should expect Dan to develop a Len-B recognizer, and separately a Sophie-B recognizer. But if the two feature patterns for a syncon are close enough to make them hard for Dan to discriminate, as are the V feature patterns in Figure II.13.12, he should develop just a "V" recognizer. The consistent shift in one of the feature values (the y-axis in the figures) should also allow him to develop a separate "Sophie-versus-Len" recognizer, which recovers the individual messages from the multiplexing mentioned above.

Len would not have been able to perceive repeats in Rob's output unless Rob gave every syncon the same duration, because Rob's output consisted of abrupt changes from one active syncon to the next. Sophie, however, could perceive them in Len's softer output trajectories, which would usually settle for a while near the feature pattern of a repeated syncon but not on a syncon that is followed immediately by a different one. Dan will be able to perceive repeated syncons in the outputs of both parents.

Pairs of syncons, as we saw above, not only act as "words" in the language of syncons, but also improve the ability of the listener to determine which syncons were intended. However, when perceptual ambiguity persists, one transition can split into two, as with Sophie's novel distinction between BA, meaning "Sphere object" and VA, meaning "oval object". Dan presumably will develop a perception of "Oval-object-Sophie" but not "Oval-object-Len" or necessarily even of "Oval-object" — unless he acts to control one of these latter percepts as distinct from a "roundish object" percept.

Let us imagine what might happen if Dan attempts to use the "Give me X" protocol of gesture and voice but uses "VA" rather than "BA" when interacting with Len before Len has learned "VA" from Sophie. Len might not recognize what Dan wants, as an English speaker might not if someone asked him if he had a "vall"; would the questioner want to know about a wall, a ball, or would she be asking if he had a fall?

If Dan's "V" is close enough to Len's "B", Len might recognize Dan's "VA" as "BA" and give him a ball — a spherical object. But that is not what Dan wants. Dan wants an oval object, which is what he gets when he uses "Give me VA" with Sophie. So Dan uses a rejection protocol that Len can perceive as such (for example, Dan might output "NO" while putting the ball down). Figure II.13.13 suggests the protocol state at this point.

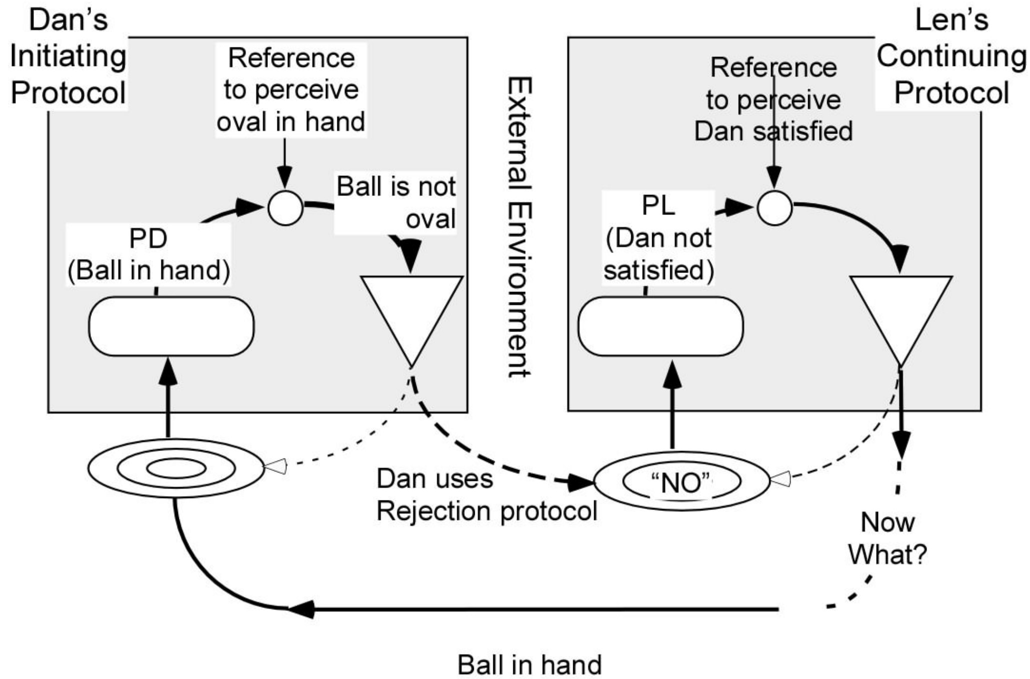

Figure II.13.13. The "Give me X" protocol that was originally initiated by Dan, at a stage after Len has misunderstood Dan's "VA" request for an "oval object" as a "BA" request for a "Ball", and Dan has used a lower-level "Reject" protocol to tell Len that a ball was not what he wanted.

Let's assume Len does not recognize Dan's "VA" as having started with "V", since in Len's language, "VA" is a non-word. It simply does not exist. The best Len could do to continue the protocol is to try giving Dan a variety of objects, finally arriving at the oval object Dan wants, at which point Dan may well output "VA" rather than using the rejection protocol "NO".

Len does have a syncon recognizer for "V", even though he has no word (trajectory) recognizer for "VA", but if this same pattern recurs, Len might easily build one, and learn from Dan the VA "word" trajectory that Dan learned from Sophie[65]. Another likely result of reorganization is a shift in Dan's pronunciation of "V" to make it more distinct from his "B", or a shift in Len's range of acceptance for "B" so that Len later will find it easier to hear that Dan actually wants something different when he says "VA" from what he wants when he says "BA". Indeed, parents often can understand their children when others cannot. The "B"-"V" phonetic distinction will have become phonemic.

Sparseness of representation represents redundancy in Len's patterns, and therefore is likely to ease the development of specific recognizers for the patterns Len does use, especially if Len usually makes a fast transition between B and A, but a slow one between F and Q. The same multiplier would apply if Len used consistent loudness variation. Eventually "BA" might come to mean something different from "BA".

We will not pursue this further now, other than to point out once again the considerable increase in discriminability loudness and timing variation permit among the relatively few trajectories that were used

65. . As usual, "learn" is a short form for "reorganizes to be able to perceive and possibly to control a perception of …".

to distinguish among the protocols in the language Len learned from Rob. This is, however, a double-edged sword, in that it allows new learners to produce identifiable syncon pairs despite having low precision in their feature values for individual syncons and trajectories. If new learner Abe deviates in one direction and new learner Zeb, who never meets Abe, deviates consistently in the other, Abe and Zeb might not be able to understand one another, though both could understand and be understood by Len and Sophie.

II.13.6 A Note about Language Divergence

In the Introduction to Volume 1, we discussed some drifts in the route from Proto-Indo-European to English, including the shift from **p** to **f** in words like *pater→father*. What happened to the PIE **p**? When it becomes **f**, pretty well all that changes perceptually is the duration of the aspiration, the puff of air that accompanies initial p in many English words. However, the easy way to produce the two sounds differs in the places where the lower lip closes off the air flow, against the upper lip or against the upper teeth.

Changing the mode of production alters the feature trajectories from and to neighbouring phonemes in the sequence. The drifts are all interconnected. And just as Sophie's accidental introduction of "VA" for a conceptual distinction previously unused by Len left Sophie and Len with a "BA" syncon "word" unchanged, so English still has a **p** sound to go with the **f**. Watkins (2000) lists no Proto-Indo-European words starting with **f**, so one may assume that English has both **p** and **f** because of the kind of split that gave Sophie "VA" while allowing her to keep "BA".

Such drifts and divergences do not occur as a result of decisions by some authority. They happen in the course of using various protocols in ways that may differ subtly from one person to another. Ohala (1992) provides evidence from many languages belonging to different language families that the source of the drifts is to be found in perceptual confusion, not in production ease, as was the case for our hypothetical Sophie when she confused Len's "BA" with the "VA" that he never intended to produce.

Just as Sophie created a new syncon pair "VA" and gave it, for herself, a new meaning "oval object" distinct from the newly restricted meaning "spherical object" she gave to "BA", so in real language do sounds and meanings split, drift, reconverge, and generally change over time, until the protocols used by groups that had common ancestry but little contact fail to work when members of different groups meet.

Sophie's conceptual structure was different from Len's because for some reason she treated oval objects as providing her with atenfels for controlling perceptions that she could not control as easily using spherical objects, or vice-versa. Len cannot have been controlling perceptions that required such a split, which would mean for him that he would not even have perceived the difference had Sophie not persisted in using what were clearly different words.

Len's reorganization would then have continued as we saw in Section 5.3 when chair parts were associated into a single complex "chair" because they moved synchronously, or when Len first perceived and controlled his synx to match Rob's syncons. The complex of Sophie, "VA" and an oval object would have become a perceptual complex that could not be controlled if Len produced what to Sophie would be a "BA". The "BA"-"VA" pair would have become two different words referring to two different concepts.

Now we can jump ahead a little, and imagine a new Sophie and Len descended a few generations from these originals in each of two families that now live apart and seldom if ever meet. We assume that the "BA"-"VA" split occurred independently in each family. In one, that younger Sophie does as we described above, but in the other, the younger Sophie perceives that a "BA" with a flat surface on one side will not roll, and calls it a "VA", meaning to her a defective "BA". If the two younger Sophies were to meet, they both would have the word "VA", but each would give it a different meaning. In translating between them, “VA” would have become what is often called a “false friend”, such as is “gift” in German-English translation (poison in German versus something donated in English.

II.13.7 Autocatalytic Creation and Words for Abstract Concepts

As Len, Sophie, and Dan reorganize their control hierarchies for perceptions of passive objects in their environment, they encounter in the world more objects and produce more concepts than were in Rob’s preprogrammed vocabulary. We have already seen that issue in Sophie’s invention of “VA” for the concept of an oval object. If they want to talk about new concepts they need new “words”.

Among these new concepts might be geometric relationships, such as one thing resting on top of another, a pattern that occurs with objects that already have individual names in the “Give me X” and other protocols Len learned from Rob. Another protocol Len might well have learned from Rob, or by interaction with his family members, is “Put X there” or “Look, X is there”, or something similar. But they do not yet have a protocol in which “there” is “on top of Y”.

If the only time one thing was on top of another was when a “PZ” was on a “HR”, the physical “PZ”-on-“HR” complex might be treated as a singular object, a “JQ” in its own right, in the way we treat a chair as an object in its own right (Section I.5.4), even though we could equally well call it a “seat with a back, resting on four legs”. One could imagine one of the family wanting to use a form such as “Look, there is a JQ” or “Make a JQ”, as we might say “Here’s a chair”.

In the new family’s experience, however, it turns out that a “PZ” is not the only thing that can be on top of an “HR”. Many things can be on top of many other things. A “BA” might be placed on an “HR” or even on a “KV”. An “AC” could be on top of a “GN” or vice-versa, an “FP” could be on “WS”, and so forth. Would they create a new word for every such possibility?

The number of complex objects (one on top of the other) to be labelled grows roughly as the square of the number of distinct object names. If a new two-syncon “word” had to be created for each of them, the new words would quickly fill the available feature space, even allowing for timing and loudness variation, making them hard to discriminate. Moreover, one family member would find it difficult to ask another to create a novel “on top of” configuration, something that would be easy for us real humans. If we said to a friend “Could you put the chair on top of the piano, please”, the friend would know what we meant, neither of us having ever previously seen a chair on top of a piano.

On the other hand, as with the oval-versus-spherical distinction, for a few frequently observed or needed “on top of” configurations in which the entire configuration usually moves as a whole, such as the chair seat on top of the legs, the invention of a distinct new word meaning “chair” would not unduly fill the available space of words, and might occur naturally like Sophie’s new word for an oval object.

There is, however, another way Len’s family might naturally solve this communication problem. Following the principle we have used whenever we have discussed the construction of new perceptual functions, we would expect Dan (and the others) to develop a perception type that Powers called a “relationship”, that would be signalled by the use of a separate syncon “word”, such as “TP”, equivalent to the English “on” or “on top of”. Rather than a “PZ” on an “HR” being a “JQ” it might be a “PZTPHR” sequence of syncons, and a “CX” on a “GT” might be a “CXTPGT”. “TP” would have acquired the meaning of the relationship “on top of”.

In the syncon language, “TP” would be unlike “BA”, in that Len could never have asked Rob for a “TP” by itself, without specifying what was on top of what. On the other hand, when talking with the quasi-human Sophie, Len could conceivably have used "TN" alone, in circumstances such as we describe in Section II.15.3, when Len believes that Sophie knows the objects and needs to know only the relationship he is controlling to see.

It is not possible to touch or to move a relationship, though a relationship can be made and unmade by changing the values of some property of its members. Nor could it be used alone in a "Make an X" protocol. "Give me an on-top-of" cannot be done, although "Give me an example of one thing on top of another" is very easily done. If there are too many objects to permit discovering a new word for each possible relationship with every possible set of objects that fit the relationship, then to use them in a protocol, the protocol interface would have to use a compound word or a sequence of words.

Synx trajectories for these longer sequences are more complex than those for simple two-syncon "words". Quite apart from the kind of idiosyncratic variation that could lead to using trajectory features to convey meaning independent of the words individually, even the simplest form of control ensures that the trajectories will not be straight. Instead, as shown in Figure II.12.2, they are likely to curve away toward the next syncon target before they reach the original target (Figure II.13.14a) — except, of course, if a syncon is immediately repeated (as is "B" in Figure II.13.14b) or terminates a sequence, in which case the synx is likely to dwell on the syncon momentarily, as at "D" in panel a or "C" in panel b.

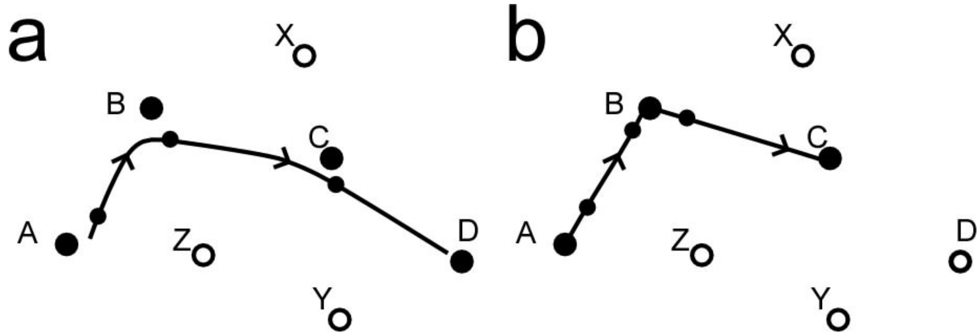

Figure II.13.14. The trajectory of the synx in feature space for the syncon sequence A-B-C-D (a, left) or A-B-B-C (b, right) The arrowheads mark the approximate location of the synx at the moment when the target shifts to the next one in the sequence, and the small dots on the trajectory suggest the moment when the next target might be identified by a listener. At no moment during either trajectory is it plausible that X, Y, or Z might be the next syncon in the sequence. Notice the relative stasis of the trajectory at the repeated B syncon in the right panel, as suggested by the proximity of the small dots on the trace to the syncon B.

In a higher-dimensional space such as the three-feature syncon space, the effect is more dramatic and the early effect of deviating toward the next target becomes clearer. Syncon repetitions, which would be impossible to distinguish from single occurrences in Rob's "speech", become easily distinguished by the abrupt change of trajectory direction before and after the doubled syncon in the speech of the family who depend on perceptual control of their synx features in both production and reception.

If the trajectory curve becomes part of the perception, as it will if the feature space is sparsely filled with trajectories, then particular curves can become recognizable components of a trajectory, just as the pairwise trajectory between syncons became a "word" with a completely arbitrary meaning. The curves help to identify word pairs, just as the simpler trajectories help to identify the syncons in a two-syncon "word". The 3-D feature space begins to contain syncon targets, "words" (target pairs) and permissible "word sequences" (synx trajectories)— the start of a basic syntax of the syncon language.

Figure II.13.14a illustrates that if the recognizer depended only on the synx being identifiably at a syncon location, as is necessarily the case when Len listened to Rob's abrupt transitions, it is quite possible that the A-B-C-D sequence would be perceived as A-C-D, omitting B, because the trajectory bypasses "B" by some margin. One hears this effect in the speech of a TV newscaster who in running speech pronounces the name of the capital of the USA almost indistinguishably from the "wash" in "wash my hands". However, the trajectory for "ABCD" is quite different from the trajectory for "ACD". The same newscaster would probably make the "ing" distinct when saying "washing liquid" rather than "Wash~~*ington*~~".

If perception of the trajectory allowed perception of the "AB" word, when Len is talking that word could have been perceived as distinct from "AC" by the time the trajectory reached the small dot near "A". Perception of "B" as a syncon target would then be a *consequence* of the perception of "AB", rather than the perception of "AB" being a consequence of the sequential perception of syncon "A" and then of syncon "B" (Figure II.13.15).

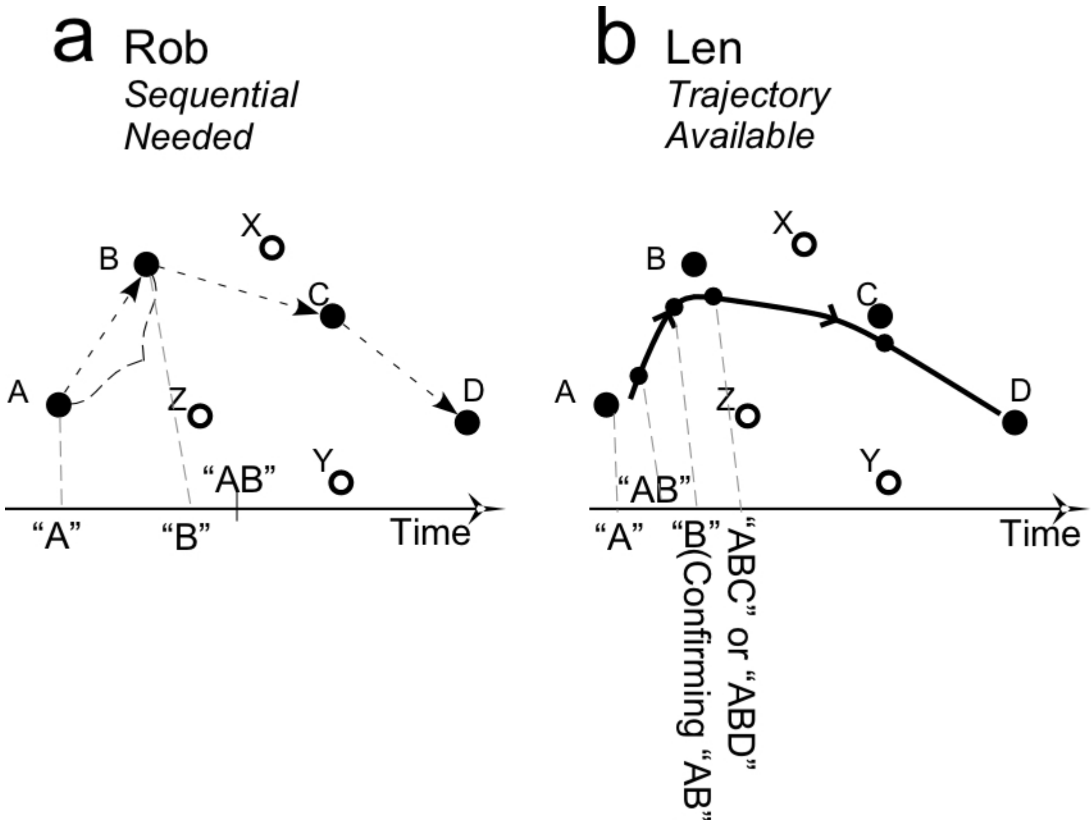

Figure II.13.15 (a, left) If the syncon sequence is perceived by perceiving "A" and then "B" in a sequence perceptual function, as it must be if Rob is "speaking", the perception of B must occur before "AB" is perceived, whereas (b, right) if the trajectory of Len's "speech" is perceived as an item of its own, "AB" can be perceived almost immediately after the trajectory leaves "A", with the direct perception of "B" coming later. By the time the Rob-listener had perceived "AB", the Len-listener might have substantially perceived the third syncon in the trajectory.

A related effect occurs in natural language: the perception of consonants depends to a large extent on the trace of the formants leading into and out of a vowel (***Ref***), rather than on static properties of the acoustic waveform during the time of the consonant, though the stable portion of a long vowel (like the B-B sequence in Figure II.13.14b) helps the recognition of the vowel.

Relationships are not the only way in which the need for new word types arises. It may be necessary to discriminate within classes of objects for which words have already been established. For example, Dan may ask for “BA”, while Len and Dan have in their common environment a baseball, a soccer ball, a balloon, and a marble. Dan wants the soccer ball, but how can he let Len know within the “Give me X” protocol frame, when all of them are “BA”, spherical objects?

One answer is that Dan might make an arbitrary new noise, a new syncon sequence — a new “word” to be used when he wanted a particular object from the set. Why would this happen? No mystery; it is what would be expected in the normal course of reorganization. If a control system doesn’t control very well — as would be the case for a protocol loop that provides the right result only half the time — its connections are likely to change until it does. The only possible output for a protocol loop Dan could use at any instant is a pattern of the three acoustic features, and such a pattern may define a syncon. All the permissible syncon locations are used, so the only effective change is in their sequencing, which means a new “word”.

To see how Dan might introduce a new word in a way that the others could learn it, refer back to Figure II.13.13. At the point where Dan uses the rejection protocol by using the trajectory “NO”, he might instead use a “correction protocol” and say a word he made up in the process of reorganization, “SM”. When Len produces a small object, the protocol is complete and Dan ceases to disturb whatever of Len’s controlled perceptions is driving the continuation half of the protocol loop.

Notice what happens here. If Len reorganizes to produce only small balls when Dan says “BA”, Dan will again use a rejection or a correction protocol when he wants a big ball. He may even use a different made-up syncon sequence, “BG”. Len should eventually reorganize so that when Dan says “BA” and then “SM”, or “BA^SM”, a small ball completes the protocol, but when Dan says “BA” and then “BG” or “BA^BG”, he wants a big ball.

Dan, for his part, might well reorganize further, and say the sequence fluently rather than as separated “words”, since after “BA” only half the time will he get the ball he wants, whereas if “SM” follow “BA” he will always get it. On the other hand, if he just says “SM”, he might be offered any small thing when all he wants is a small ball. The sequence “BASM” gets what he wants, whether or not Len tries offering different balls after Dan’s separated “BA”. So it is reasonable to assume that “BASM” might be what Dan produces when he wants a small ball, and “BABG” when he wants a big ball. Len should be expect then to develop trajectory recognizers for both sequences, as suggested in Figure II.13.16.

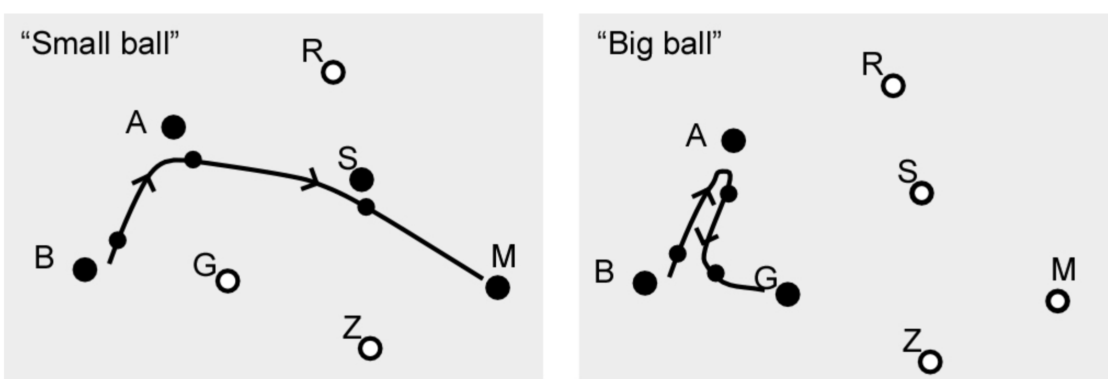

Figure II.13.16. Feature-space synx trajectories in two dimensions that Dan might use to refer to a small and a large ball. In the case of the small ball, the trajectory does not pass very close to the "A" syncon target, and in the case of the big ball, the trajectory does not get close to the "B" target for the second "B" in "BABG".

The two-dimensional synx trajectories in Figure II.13.16 assume that Dan says "BASM" and "BABG" fluently, rather than as "BA^SM" and "BA^BG" separately. The listener might not independently identify the use of all the syncon targets as the synx passes them, though the trajectory level recognizer would be clearly distinct from most other trajectories that used any three of the four targets. With more than three features, the trajectory-level distinction would be even clearer.

As a consequence, the sequences "BASM" and "BABG" might come to be interpreted not as "SM" and "BG" classes of balls, but as two words referring to distinctive classes of objects, "BSM" objects and "BAG" objects. When used in protocols, those classes of object might well come to be conceptually distinct, just as their trajectories are perceptually distinct, much as Whorf (e.g. 1944, 1950) claimed that conceptual and linguistic structures are linked in feedback loops, and as Taylor and Taylor (1965) did when arguing that feedback loops are the generators of phonetic symbolism.

The evolutionary process of elision is suggested in Figure II.13.17. The evolution could happen in the outputs of one protocol user, such as Dan, or over the course of generations, each one learning a slightly more elided version from its predecessor.

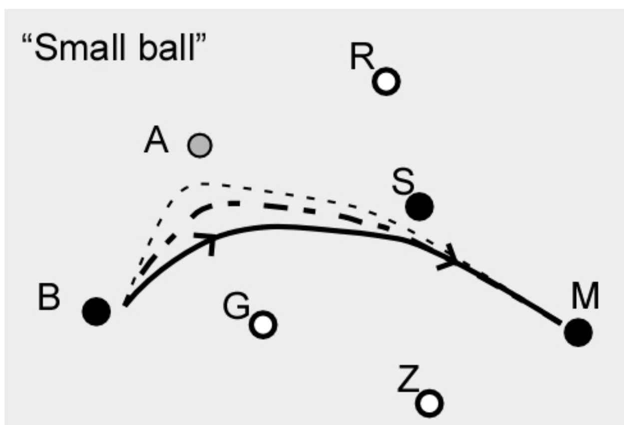

Figure II.13.17. Evolution of a trajectory "B(A)SM" that elides one of the syncons in the original trajectory, creating a new "word" that will be used in protocols, and might thereby induce a new perceptual category of object.

This sort of drift can be observed in natural language (e.g., Ohala, 1992, who gives many examples of different kinds, all of which can be understood as instances of the foregoing). It is evident in some of the so-called peculiarities of English spelling. As just one example, the "gh" in "light", "bright", and "night" once signified a "kh" sound, a sound lost in most English dialects, but retained in the Scots "*It's a braw bricht moonlicht nicht the nicht*" (It's a beautiful bright moonlight night tonight, in which the elision of the "kh" sound is accompanied by the short "i" sound turning into an "ei" diphthong).

Of course, small and big are not the only subclasses of ball that Dan might want. Colour is another possibility, and like small and big, colour is applicable to other objects as well as to balls. So if Dan created a word “RG” for “Red” when he wants a red ball and “GZ” when he wants a blue one, the words “BARG” and “BAGZ” might be evolved through the process described above, the latter subsequently being elided to “BAZ”. “BARG” might not be elided if there was already a “BAG” elision of “BABG”, since perceptual control would be lost if sometimes “BAG” resulted in Dan receiving a red ball and sometimes a big ball.

If, however, the class of “red ball” is the kind of class that would be used sufficiently often to lead to the elision, then its elision would be likely to be distinguished by trajectory variants from that of “BABG” such as those illustrated in Figure II.13.17. Maybe the “BA(B)G” elision would additionally evolve to a fast trajectory past the “A” while the “BA(R)G” elision would evolve to a slow one (Figure II.13.18).

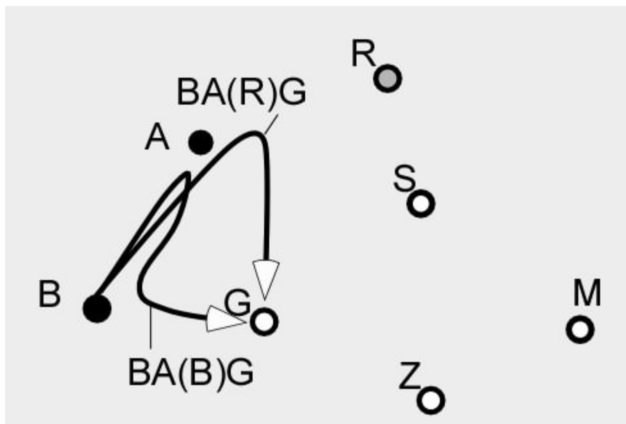

Figure II.13.18 Although both BABG and BARG may by elision result in BAG, the “A” may be pronounced quite differently in the evolved versions. If the syncon language was ever written, the two forms would be spelled the same but pronounced differently.

Distinction of this kind is akin to the distinction in English between “***de***sert” (a dry area) and “des***ert***” (a soldier leaving his or her unit without permission). Homophones do, however, exist in natural languages, and they can be distinguished only by trajectories at a higher perceptual level, or, to say the same thing in other words, by the fact that just one of the words fits the context.

A small red ball, by the process described above, might initially be a “BASMRG” or a “BARGSM”, interchangeably. Dan might not, however, want a ball, but a small red toy. If “toy” was signified in the “Give me X” protocol by the syncon sequence “VF”, then Dan might use different classifiers for small and red than he used for balls, but equally he might use the same ones, having developed a set of perceptual functions that used the similarities among small objects and among red objects as distinct from big ones and blue ones.

If he used the same classifiers, then he probably would not elide the “small red ball” trajectory, since part of the trajectory would be the same as that for a “small red toy”: “VFSMRG” or “VFRGSM”. But there is another possibility, that the adjectival classifiers might be ordered differently for balls and toys, which would enhance their ease of discrimination. The lateral inhibition of the HaH process (Chapter I.9) plays no favourites as to the direction in which the cores of perceptual functions separate. SMVFRG is more distinct from BASMRG than is VFSMRG.

Verbs are another kind of descriptor of relationships, but the relationships implied by most verbs are dynamic, evolving over time, whereas the geometric relationships are static. The “Give me X” protocol was described as a frame in which X could be a ball, a toy, or any other object, the frame being constructed from gestures. Only the “X” involved the use of syncons and syncon “words” and “phrases” described by the synx trajectory.

But why should not the “Give” and “me” parts of the protocol also be constructed from words? We are talking about verbs. Here we simply assert the obvious fact that protocols can and do include them, rather than tracing the very similar path by which they join the vocabulary of the syncon-speaking family.

Beyond verbs, real-world protocols include such “courtesy words” as “Please” and “Thank you”, which do not obviously signal error that needs to be corrected through the actions of a partner. When considered together with features such as intonation, they allow the partner to perceive something of how the speaker is disposed toward them, as well as perceive how the speaker speaks (see the chapter by Nevin in LCS IV). Does the speaker, perhaps, welcome further communication? Even though the words themselves may not signal error that needs to be corrected, their trajectories in the expanded feature space of timing and variation of contour may help the partner to choose the most effective action in initiating or completing subsequent protocols.

If Len happened to do the same trajectory shift in particular contexts with other words and the corresponding objects, the shifts themselves might come to have meaning. A rise-fall relative to the straight-line change of fundamental frequency between syncons might come to mean “This is called a…” while a simple fall could come to mean “I give…” or “Passing … between us” where “…” would refer to whatever object was labelled by the spoken word. Alternatively, depending on how the relationships of trajectory and context happen to align, a simple fall might mean “I’m sad” and a fall rise “This is a question”. The meanings that come to be associated with trajectory variations are in principle quite arbitrary, and could be very subtle.

Trajectory characteristics over syncon sequences have a status similar to the feature characteristics of the individual syncons. Trajectory-level features, for the longer syncon sequences we shall be considering, become “prosody”. At every new perceptual level, the dynamics of the variations at the simpler old level become features that could become perceptual variables at the new level. For ease of discussion, though, we will continue to concentrate on the three features that define the individual syncons, treating the role of trajectories in easing discrimination while ignoring (temporarily) their role in the meanings of protocols.

Recognition ease depends on the separation of the individual macrostates within the appropriate universe of description. Len’s interactions with Rob used “BA” but not “VA”, so if Len made a sound that Rob could have categorized as either “B” or “V” Rob could be sure that Len intended “BA” and not “VA”. When dealing with Sophie, who creates a word “VA” with a new meaning, Len might always say “BA” rapidly and “VA” slowly, or give them different intonations, maintaining or even increasing their separations in trajectory space. But if Len is unable to control some aspect of the trajectory precisely, then that dimension of the universe of possibilities is reduced, just as would be the case if, in the 25-dot universe, Len and Sophie could identify only four locations of the dot within its square, rather than nine. Instead of gaining 82 bits of communication possibility, they would gain only 50.

Humans and other social animals incorporate gestural features in their communicative strategies, which further augment the feature space. Our pseudo-human users of the syncon language are unable to do this. Nevertheless, their autocatalytic inventive productiveness of novel features has the same kinds of effect that we here associate with trajectory shifts. What for a human might be a difference between, say, uttering something with a hand across the mouth and uttering the same thing with the hand by the side might, for Len, Sophie, and Dan, be accomplished by using different speeds of transition of the synx between syncons.

A generic pattern of trajectory types based on the perceptual differentiation of features was found by

Everett, Blasi and Roberts (2015). Perceptually, an irregularly produced feature is harder to discriminate than is one produced with precision, so they predicted that since vocal fold vibration was more jittery and irregular in dry or cold air, people in dry or cold regions would have more trouble producing unambiguous rapid pitch fluctuations, and languages in those regions would be less likely to use trajectories that include pitch variation (called "tones") than would those used in warm humid climates. Their survey of over 3700 languages found that almost all tonal languages, especially those with a complex tonal pattern, are found in warm humid regions. Even within China, the northern Beijing dialect has four tones, whereas the southern Cantonese dialect is variously reported to have between six and nine (Taylor and Taylor, 2014).

II.13.8 Higher-level trajectories

The synx trajectory itself is labelled by the sequence of labels for a syncon sequence that would easily lead to that trajectory. The labels "spell" the syncon sequence, but they do not represent the trajectory any more than spelling a spoken word indicates the speaker's mood and intent.

"*Could you spell that name for me, please*". When one listens to foreigners attempting to speak one's language, or when one tries to speak another's language, having access to the written versions of the words, the resulting speech is often recognizable as derived from the spelling, but does not sound much like the version spoken by a native speaker. The more recently the language was given its current written form, the more likely it is that the feature trajectories are reasonably close to those that would be derived from the spelled versions of the words. English spellings began to be stabilized with the advent of popular printed material and the development of dictionaries three to five centuries ago. Few native speakers of most varieties of English now pronounce the "k" or the "gh" in words such as "knight", "night", or "know", though there are dialects in which those letters can still be detected in the speech of natives. Other languages have been given written form much more recently, and most still remain unwritten.

The converse of "*Could you spell that name for me, please*" can also occur. Language in Indo-European languages and many others is written by a sequence of identified letters, but even there, the concept of "trajectory" exists. Letters do not exist in isolation, and neighbouring letters are not perceived necessarily as individuals. Words have what Taylor and Taylor (1983) called "Bouma shape" after Bouma (1971), who categorized lowercase letters into seven groups. Letters within a group were more easily confused with each other than with letters in other groups (Table II.13.1). Many words are uniquely identified and can be read if their letters are individually represented by a nondescript shape akin to a generic member of their group.

Table II.13.1 Bouma's seven letter groups

Short	Tall	Projecting
II.1. a s z x	II.8. d h k b	II.13. g j p q y
II.2. e o c	II.12. t i l f	
II.4. r v w		
II.5. n m u		

Bouma went further, however. In a now closed science discovery museum called the Evoluon in Eindhoven (The Netherlands), he arranged a display of a small written paragraph that was initially shown very blurred, but slowly came into a focus when the viewer pushed a button. At an intermediate stage, the paragraph was easily read, but as it came into better focus the reader could see that it contained no letters. Instead each word consisted of just a wiggly line with appropriate projections above and below the central band. Handwriting can in any case be seen as a trajectory, a trajectory of pen movements, one that elides the precise trajectories of the letter forms, and as in this example may elide entire letters provided the "Bouma shape" of the word is visible.

Figure II.13.19, which is spread over the next six pages, is an attempt to represent on the printed page something of this effect. The successive pages show the same image with successively less blur. You should be able to read it by the fourth example. Here is the first and most blurred. Probably all you can see is that there are some vaguely defined blotches arranged in more or less horizontal lines.

Figure II.13.19a The most blurred image of the text.

This is the same text, slightly less blurred. Perhaps you can make out that the blurs seem to be handwritten words of a scrap of text.

Figure II.13.19b The second most blurred version

The third level of blur. By now you may be able to read some words. It might be easier to read if you simply glance at it rather than working to make out the words.

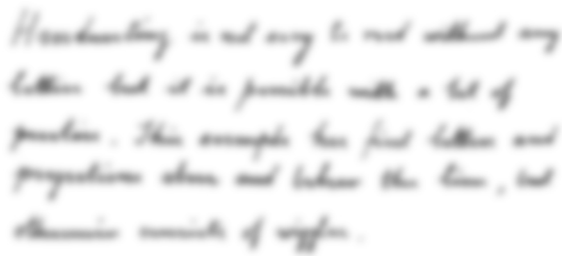

Figure II.13.19c The third level of blur.

At the fourth level of decreasing blur, you should be able to read most of the text quite easily if you don’t try too hard, and just look at it briefly.

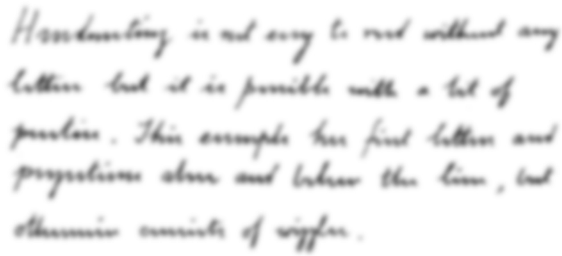

Figure II.13.19d Level four of the deblurring.

This version or the previous one should be the easiest to read.

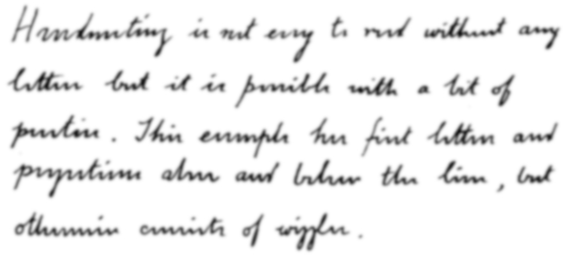

Figure II.13.19e The text only slightly blurred.

And this is the original, any blurring being due only to the graphic processing and printing. You may find it harder to read than the preceding two blurry versions, because now you can see that what you previously perceived as letters (because you had perceived the whole words that contained them) are not letters at all apart from the first letters of the words, but simple wiggles all alike other than the projections that create the Bouma shapes of the words.

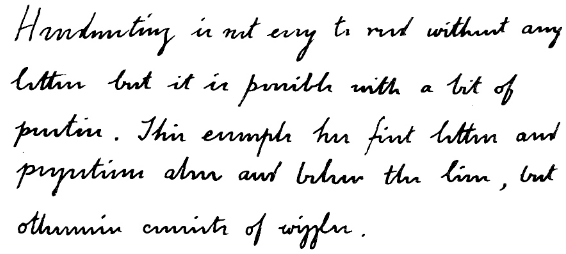

Figure II.13.19f The original "text", which you can now see contains no letters other than the first letters of the "words" together with dots on "i" and crossbars on "t" letters.

Here are the six levels of blurring again, shown together in a single display.

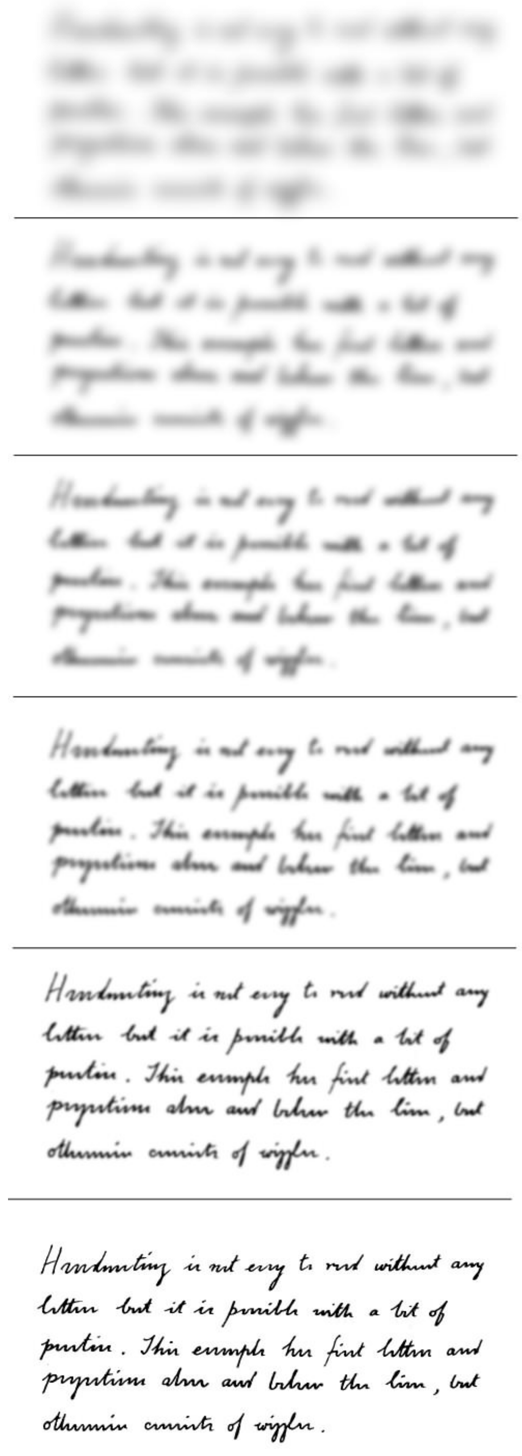

Figure II.13.19g Handwriting is not easy to read without any letters but it is possible with a bit of practice. This example has first letters and projections above and below the line, but otherwise consists of wiggles.

Returning to the world of synx and syncom, the “meaning” of the synx trajectory resides in what perceptions it might help control, and a meaningless trajectory will be easily lost from the recognition repertoire. If an elided trajectory, such as “BSM” for “BASM” is not recognized by the protocol partner, it is likely to disappear through reorganization, leaving as a survivor the version that is understood. However, as with “knight”, the elided version might well be the survivor if in normal usage there is little likelihood of a confusion with a homophone (“night”) or a near homophone (“might” or “mite”).

“Normal usage” is a key phrase. Just as we can evoke trajectory recognition as perception in a higher-level feature space than the space in which individual syncons are perceived, so we can argue that perception of sequences of feature trajectories (“words”) is perception in a still higher-level space we might call “phrases”.

In Section I.12.7 we discussed the automatic filling-in of missing and the correction of wrong items in contextual perception. Now we see that this is exactly what is involved in the perception of trajectories. If necessary, the component syncons could be filled in because of the perception of an identifiable trajectory, or the “words” filled in because of an identifiable semantic, syntactic, or pragmatic context. Of course, “if necessary” means neither “automatically” nor “always”.

Perception of trajectories of trajectories is just another way of talking about the perceptual side of the perceptual control hierarchy, more complex perceptions being based on ever more inputs at the most basic level. At each level, the missing items can usually be asserted (Section I.12.7), and so it is quite possible to replace sequentially from coarser to finer the elements of a high-level trajectory from which much has been elided. A wife’s sigh might in the right context be filled in with the letter sequence: “Aren’t you ever going to ….” (fill in the blank from the pragmatic situational and historic, and possibly the conversational, context).

At an even higher level, did either Holmes or Moriarty omit anything from their messages to each other in Conan Doyle’s expressive interaction between them?

> Professor Moriarty, "The Napoleon of Crime," visits Sherlock Holmes, the two having never previously met..
>
> Moriarty: *You evidently don't know me.*
>
> Holmes: *On the contrary, I think it is fairly evident that I do. Pray take a chair. I can spare you five minutes if you have anything to say.*
>
> M: *All that I have to say has already crossed your mind.*
>
> H: *Then possibly my answer has crossed yours.*
>
> M: *You stand fast?*
>
> H: *Absolutely!*
>
> (Arthur Conan Doyle*, The Final Problem*)

Moriarty found out exactly what he wanted to know about Holmes’s understanding and intentions, as did Holmes about Moriarty’s understanding of the pragmatic situation. Nothing unnecessary was said, and nothing was omitted. The trajectory of the conversation at a very high level was perfectly discriminable from other possibilities to them both, though it would not have been to Dr. Watson, as a third party observer who did not know all that the protagonists knew of each other. Had the interchange been conducted with either party intending that Dr. Watson should understand what was going on between them, very much more would have been said. In that case, the interaction as quoted here would have been found by Watson to be full of omissions, despite being the same as the original complete version according to any conceivable physical measure.

II.13.9 Phonetic Symbolism and the Regularization of Verbs

We next consider a phenomenon known as "phonetic symbolism", which, to judge from the references in Google Scholar, was the subject of many studies between about 1930 and 1965, after which it received academic attention only sporadically until fairly recently, when it became of interest to those wanting to create brand names that sounded appropriate to their products.

Phonetic symbolism is introduced here as a possible model for a more general process whereby word parts may become associated with different domains of experience. It also offers an example of the same process that we illustrated in the handwriting example, of filling in from context based on perceiving trajectories. More generally, we will later suggest that similar processes occur in the development of social attitudes toward cultural phenomena, such as food or music preferences, clothing choices, and even attitudes toward immigrants, among others. Here, however, we restrict our inquiry to the connotations of the sounds of words.

In its purest form "phonetic symbolism" refers to the effect particular sounds have on the feel of words. To what degree does the (nonsense) word "vam" feel, say, warm, or "ket" feel pleasant, and if people consistently agree, is this agreement reflected in the words used in their language? If speakers of a given language do agree consistently, do speakers of unrelated languages feel the same way about a range of such distinctions?

Taylor and Taylor (1962) investigated these questions directly, by asking monolingual speakers of four unrelated languages (English, Korean, Japanese, and Tamil) in their home countries to judge a series of CVC nonsense syllables on four apparently unrelated dimensions of meaning (size, movement, warmth, and pleasantness). The syllables used 12 consonants and six vowels selected to be as closely as possible available in all four languages. With very few exceptions, they found strong agreement among speakers of each language on each of the semantic dimensions for the independent effects of the sounds of each of the initial consonant, the vowel, and the final constant. However, they found no correlation across speakers of different languages on any semantic dimension for any of the phoneme positions (except for a small correlation with English in most cases), and no correlation across semantic dimensions.

Writers on phonetic symbolism often start by pointing out that in English, "i" as in "bit" is associated with words with a connotation of small, whereas words with low sounds such as "oo" in "boom" are associated with large, loud, or unpleasant things. They then go on to suggest this must be universally true because small things usually make high noises and big things low noises. Taylor and Taylor refuted this by showing that in three of the four languages, nonsense words with the vowel /i/ were given the largest or nearly the largest rating among the vowels.

Overall, the results reported by Taylor and Taylor (1962) show that the effects of phonetic symbolism are real and strong, and as arbitrary as is Len and Rob's use of the syncon pair "BA" to represent a smooth roundish object. If there is any underlying physical or evolutionary underpinning for it in the different dimensions, that influence is very subtle as compared to random language-specific influences. Measured by the degree of agreement among speakers of a language, phonetic symbolism for vowels is strongest in Japanese, followed by Tamil, Korean, and finally English, whereas for consonants the order is Japanese, Korean, English and Tamil. English, the language most studied in the literature, has the weakest phonetic symbolism overall among these four unrelated languages.

A search in Google Scholar reveals no subsequent studies that have compared phonetic symbolism in unrelated languages other than English. Nor do there seem to have been later studies that compared several semantic dimensions across even two languages. Since Taylor and Taylor found low but statistically significant correlations between English and all three other languages, studies that include only comparisons with English (and probably other Indo-European languages) are irrelevant to the question of whether phonetic symbolism is universal, even perhaps having a physical or physiological underpinning, as has often been suggested.

Taylor and Taylor attributed the frequent but low correlations of the other languages with English to the pervasive nature of English as a global *lingua franca*, to which even apparent monolinguals are likely to have had appreciable exposure. Furthermore, many languages have a body of loan words from English that might also contribute to the small observed correlations with English. This especially true of Tamil, spoken in lands that were English colonies, and of Japanese and Korean, spoken in countries that had massive infusions of English-speaking military personnel after the Second World War and the Korean War.

I. Taylor (1963) and Taylor and Taylor (1965) proposed a symmetry-breaking feedback theory of the development of phonetic symbolism very like the theory suggested in Chapter II.11 for the interaction of mother Cora and baby Ivan as they developed protocols and created meanings of Ivan's initially random sounds and gestures. To quote a section from Taylor and Taylor (1965):

> *Suppose a language is produced without reference to existing languages, and is used as the sole language of some community. ... In our hypothetical protolanguage, sounds occur at random in association with concepts. Sampling fluctuations ensure that some sounds will actually occur more often in connection with particular concepts than in the total vocabulary of the language. For example, let us suppose that the language has 20 words with the connotation of hot, and that /R/ is the initial letter in five of them,. If /R/ is the initial letter of less than 25% of the words of the language, then it is overrepresented as the initial letter of "hot" words.*
>
> ...
>
> *If, in our protolanguage, a person uses a word with the connotation of heat, he is likely to use a word starting with /R/. Whenever this happens, /R/ and "heat" tend to become associated with one another. If /R/ is also overrepresented in "cold" words, then /R/ and cold also become associated. There would then be no bias toward /R/ feeling cold or warm, but it would be associated with temperature in general. However, if /R/ happened to be underrepresented as the initial letter of cold words, then its temperature-related associations would be to heat. The speaker of the language would feel /R/ to be proper in a "hot" word, and might feel that it was in a way wrong in a cold word. He has developed subjective phonetic symbolism for /R/ on the hot-cold scale of meaning.*
>
> *If the speaker has developed such subjective phonetic symbolism, it will influence him in the choice of words when he speaks. To give the impression of heat, he can use /R/ words, even if they do not directly denote temperature. When he wants to use a word directly related to heat, /R/ words will seem appropriate. Hence the effective vocabulary becomes biased so that the proportion of /R/ words used for heat increases, and the proportion used for cold decreases. This increase in bias aids the development of subjective phonetic symbolism in other speakers, who are brought up using the changed language. (Taylor and Taylor 1965, pp. 424-425).*

The long-term stability of sound patterns with specific semantic associations, such as the proto-Indo-European "dwo", which appears in many languages in words representing "two-ness" (Watkins, 2000), argues for the idea that phonetic symbolism is not a curiosity, but a fundamental aspect of language, stabilizing spoken language through a feedback process that occurs with almost every word spoken. The relatively low level of phonetic symbolism in English may be attributed to the mixture of sources for the English language, which is basically Germanic, but which has a substantial Romance overlay and includes many words such as "algebra" or "kinetic" that have been borrowed or derived from other languages, some in different language families.

The present work on Perceptual Control in language and culture is based around the argument that the same process as was hypothesised for the interactive development of both subjective and "objective" phonetic symbolism (the overrepresentation of particular sounds in words with particular connotations) also plays out in the development of protocols more generally, and is central to the idea of language as an artifact.

II.13.10 Maintenance of Language Functions

Much of the preceding material has been devoted to the question of what perceptual functions the "syncon-speakers" should be expected to develop. Equally important is the question of what perceptual functions for language patterns should be expected to persist. Here, we can refer back to both Powers and J. G. Taylor. Powers describes reorganization not as confirming or reinforcing structures that exist (as a "reinforcement" process would do), but as modifying those structures that do not control their perceptions in ways that help maintain intrinsic variables near their reference values. Perceptions that are not part of effective control loops would tend to be lost over time as reorganization proceeds. J. G. Taylor demonstrated that for a perception to be developed and retained, it must be a part of an active behavioural feedback loop. Together, they make a compelling case for the proposition "use it or lose it". Both deny that what is lost is lost through decay, and agree the loss happens because the behavioural feedback influences other perceptions that might use some of the same resources.

In English, most verbs form the past tense by adding "-ed" to a stem, as in "fail"-"failed", but many common verbs are irregular, as in "go"-"went" (which derive from two different Proto-Indo-European roots, meaning respectively "to release" and "to turn" or "to wind"). According to Watkins (2000), the participle "-ed" derives from Proto-Indo-European "-to-", which, when attached to a substantive, indicated accomplishment of the meaning of the base word (as "-ed" does in "beard-ed"), so one might have expected "*go-ed" (spoken as a homophone of "goad") to have been a word meaning "to have released" and "wended" to have meant the past of "to wind". Indeed, we still use "wended his way" in the sense of his having taken a winding route. However, by elision as discussed above, "went" could be an elided way of saying "wended", perhaps to connote a less circuitous way of going. The irregular "go"-"went" can thus be seen as the use of one verb for the present tense and a different verb for the past, supplanting a confusable regular form of the past tense[66].

Using the rationale suggested above for phonetic symbolism, we can argue that a conflict is likely between the perceptions of "*past action in general → -ed*" and "*past action of going → went*", when the past of an irregular verb such as "go" is wanted. To say "go-ed" may sound right to a child because "-ed" just does sound right for past actions, but if a child uses "go-ed", collective control by adults is likely to correct it. "Correction" means the creation of error in a controlled perception, and according to Powers, this would lead to an increased rate of reorganization. Powers could not have known of rattling, but we might now argue that unexpected corrections would increase the local level of rattling and similarly would usually lead to an increased rate of reorganization. Since the past of "go" is frequently wanted, correction would happen fairly often, whereas for an infrequent irregular verb it would seldom happen. In such a way, by collective control, may the regularized form of an infrequent verb, but not of a frequent verb, become prevalent and not disturb the perception of a potential corrector.

In the early days of the evolution of English from Anglo-Saxon after it was overlain by Norman French, a higher proportion of verbs were irregular than is the case now. Lieberman et al. (2007) found

66. Whether a philologist would agree with this folk-derivation of "go – went" is questionable; the example is intended only to illustrate one kind of process that might lead to irregularity.

that the rate of regularization depended on how frequently the verbs were used, the half-life of the largest group (by frequency of use) being about 2000 years. They found that the probability that a word would pass from irregularity to regularity in any given decade or century is proportional to the square root of its frequency of occurrence (at least in written text, there being no evidence of word frequencies in speech before the era of sound recordings). Seldom does a regular verb become irregular, though it does occur and is occurring now (e.g., Michel et al. 2011).

Cuskley et al. (2014) claim to have refuted the finding of Lieberman et al., based on an analysis of US writing over the last 150 years, but the data they report actually conforms very well to the findings of Lieberman et al.. What they report as showing no decline in the number of irregular verbs actually conforms better to an average half-life of about 2000 years than it does to zero decline, in spite of the fact that they counted all the verbs found in their corpus rather than tracing the history of individual words as did Lieberman et al. (2007).

Back-extrapolating the shift toward regularity into prehistory, Lieberman et al. suggest that in the distant past, almost all verbs would probably have been irregular. Reading very crudely from their graph, this state could have been as recent as 5000 years ago, but of course it is uncertain as to whether it ever was the case. If the irregularity of early English happened to be caused by the mixture of Anglo-Saxon and Normal French, the trend should not be extrapolated beyond the time when the two languages were merging, some 800 years ago. This point could be tested, however, by doing a similar analysis for the four major descendants of Latin, namely Italian, Spanish, Romanian, and French.

However that may be, the point is that there are two conflicting effects, one tending toward regularization because the regular formation is used often, the other tending toward irregularity because the regular formation leads to more perceptual confusion than does the use of a form based on a different root. Thinking of a language as a network of relationships, it is interesting that these opposing tendencies apply to different links in that network, suggesting that its structure may have the form of a tensegrity structure. Since as far as I know nobody has investigated that possibility, I leave it as an unfounded speculation worth research.

Chapter II.14. Protocols

I have used the word “protocol” casually a few times, particularly in the introductory Chapters of Volume I, but now we treat the concept of a “protocol” in a more formal way, as it is used in “Layered Protocol Theory” (LPT, e.g. Farrell et al., 1999). In that context, a protocol is a means by which two people, each controlling their perceptions of the other’s actions, can pass information from one to the other.

Mother Cora and newborn baby Ivan (Chapter II.11) did not start out by using a protocol for communication, but Cora learned to distinguish between Ivan’s “hunger’ and ”pinprick” cries and limb movements, and Ivan learned non-consciously (perhaps by Hebbian enhancement and anti-Hebbian reduction of synapses) to cry and move in different ways so as to get the pin removed or to get milk. When each has learned consciously or non-consciously to perceive these different behaviours in the other, they have started to use an elementary protocol in the LPT sense of protocol, and they have begun to create a private language into the bargain.

Just as “perception” has a casual everyday meaning and a formal meaning within Perceptual Control Theory, so does “protocol”. In Chapter I.8 we introduced the concept of “Motifs of Control”. A protocol is such a motif, a rather elaborate one that in its full form involves 19 different control loops, while being based around a core of only two, one in each partner. Furthermore, protocols exist in levels, simpler ones supporting longer-lasting ones in the same way that control of simple perceptions supports control of more complex ones that have a longer “now” (Section I.7.2). It is not easy to imagine a scenario in which all 19 possible loops would be used at a single level, though most non-trivial protocols are likely to use around seven.

Protocols may form a hierarchy akin to the perceptual control hierarchy, though different in that the perceptual control hierarchy is built from the sensory-motor interfaces with the external environment, whereas a protocol hierarchy is built from a complex level down through simpler protocols to perceptual controls through the environment, each of which controls some perception of the communicating partner by trying to disturb the partner in a simplified version of the Test for the Controlled Variable (TCV, Section 1.2.5).

Everything we will say about protocols between human and human applies equally well if one of the participants is not a human. All that differs is the set of belief states that the human participant may be able to control about what the partner may know or understand, whether the other is a living control system like a pet dog or a machine such as a self-driving car or even a TV set. The difference is only in quantity, either of protocol levels (only one in a TV set interface) or of variety within a protocol level. One might add a qualitative difference, in some cases, of whether and to what degree the non-human partner is capable of learning.

Many animals that interact much with humans can learn a lot, whether they be family pets or wild birds, fish, reptiles, or mammals. Most machines cannot, but the protocols applicable to human-machine interaction are exactly the same, even for a machine as simple as a hand-pushed lawn mower. All that differs is which of nine or ten perceptions about states of the machine can be controlled by the human partner, the operator of the machine. So let us explore the structure of the two sides of every protocol, whether it be one that might last years, such as a protocol for teaching a student the mathematical intricacies of General Relativity Theory, or parts of a second, such as a simple protocol that is effectively just control of a single perception such as the on-off status of a light switch.

Exactly the same protocols can be used for honest or for deceptive communication. We discuss deceitful use of protocols to the next Chapter, but mention it here in preparation. For example, since presumably long before Sun Tzu (ca. 600 B.C.E) set it down in writing, military commanders have tried to lead their opponents to perceive a state that does not exist so that their actions will help the commander to win the battle. They can do this only by communicating with the opponent, whether or not the opponent knows that the communication is happening. Contemporary salespersons and advertisers, as well as politicians, are in the same position, if they are trying to get the potential buyer or supporter to believe something about the product that the advertiser or politician believes to be false. All of them must use the same protocol structure if they are to discover whether and how well their ruses are working.

II.14.1 Protocols and belief control

Nineteen control loops may seem to promise a very complicated motif, but like most motifs that appear frequently in normal behaviour, its structure is not very complex at all. Eighteen of the control loops form six groups of three, three groups performed by the initiator of the protocol, three by the "Continuer". These eighteen loops all are about beliefs concerning the current state of the communication of a single Primal Message, and they are all at their reference value when the transmission of the Primal message is either complete or rendered no longer necessary by events perceived by at least one of the partners.

Each group of three control loops in each partner is based on their levels of belief about one of three different propositions.

- P1: *That the Continuer has understood the Primal Message.*
- P2: *The quality of the communication process is sufficient for an adequate interpretation*.
- P3: *That it is not worth continuing to transmit this message.*

P1 mentions a "Primal Message". Each instance of a protocol has a Primal Message, which is the communication analogue of the error in the perception whose control required the Continuer to do something. If the Continuer has understood the Primal Message, she has perceived what the Initiator was controlling for her to perceive.

The Primal Message should not be confused with another term we will use: the "Primary Message", which consists of the overt actions, including language, that the Initiator first uses to describe the Primal Message to the Continuer.

If we call a generic proposition "P" then each group of three propositions in each partner has control loops for belief perceptions with these three recursive reference values: (1) Self believes P, (2) Self believes Other believes P, and (3) Self believes Other believes Self believes P. When all three beliefs match their reference values (believed at least strongly enough to be within their tolerance levels) for any particular proposition P in both partners, no further action is taken with respect to that proposition unless some external disturbance affects the relevant belief perception.

Many of the eighteen propositional recursions in the two partners start at their reference values to begin the particular protocol instance, and so require no action to bring them there. For example, two friends in a quiet environment probably both believe all the recursions of P2 before they start. They believe, for example, that the partner uses and understands the same language as they do, that the partner's hearing in adequate (or that the partner can read, if the communication is in writing), that the environment is quiet enough, and so forth. If the perceived state of belief for all the recursions of P2 in each partner is within its tolerance level, what is left is only to determine whether one has understood the Primal Message and is believed to have understood what the other wanted them to understand (P1), or whether more needs to be done to bring about this understanding (P3).

The 3x3x2 pattern (three propositions at three levels of recursion in two partners) defines 18 of the 19 control loops. What about the last one? The Initiator, say a grown Ivan, started the protocol because there was some perception that for some reason he thought was more easily controlled if the Continuer, say Ivan's mother Cora, helped either by a physically effective action or by displaying a change in their understanding of the Primal Message. The protocol is the means for Ivan to control that perception. When the Ivan uses the protocol, the Continuer, Cora, takes for him the place of an atenfel, a component of the environmental feedback pathway in the control loop of his to-be-controlled perception.

In normal everyday conversation, both partners already believe all three recursions ((1) A believes, (2) A believes B believes, and (3) A believes that B believes A believes) about at least P2, and perhaps the Continuer even has a good idea about what Primal Message the Initiator will hope to get across, which would imply that the beliefs even about P3 are not far from their reference values at the start of the protocol. Since "believes" is the reference level for all of recursions of P2, no action is required on those controlled perceptions. But if a loud noise interrupts the conversation and continues, at least one of the partners is likely to perceive their own inability to understand the other (lowering their level of belief in P2 at some recursion level), or at least to question whether the other can understand them. In this case, at least one will act to correct the error, perhaps by signalling that they might move to a quieter place. That signal initiates a new (lower level) "supporting" protocol as being the appropriate action to reduce the error in one or more of the P2 control loops.

A supporting protocol is to a protocol hierarchy just what a lower-level perceptual control loop is to the perceptual control hierarchy. In fact, since a protocol is simply one way of controlling a perception when it is different from its reference value and the hierarchy inside the individual has no atenfels for control of that perception, a protocol hierarchy can be seen as a perceptual control hierarchy spread across two individuals. The parallel is not exact, because although the perceptions controlled in the 19 control loops of a protocol all require parts of their feedback loops to pass through the shared environments of the two partners, the recursive feedback loops of the protocol itself are within the partners, with no reference to how they are physically implemented.

So long as both partners control for all three recursions of P1 to come to their reference values, they will maintain error in P3 and probably its recursions, continuing the conversation. But always there is a possibility that one of the partners comes to believe P3 and commits an action that has been called an "Abort". One possible reason for an Abort action could be that some controlled variable that contributed to the protocol being initiated has taken a value near enough to its reference value that the protocol action (sending the Primal Message) is no longer necessary.

Perhaps a higher-level protocol has had all its P1 controlled beliefs come to their reference values and has ended, no longer requiring this supporting protocol to complete its work in getting its own Primal Message understood. In effect, the Abort in this protocol is performed by the Continuer saying "I get it now", though of course at different levels if would not be done that way. Simply continuing the conversation that is ongoing at a higher level is often enough. A very low-level example of this occurs in reading, when the reader fails to notice a missing letter in a word, or even the duplication of a word, because the meaning of the passage is already clear by the time the misprint could have been identified.

Protocols can fail, meaning that they are aborted with P1 still in error, which will not be the case if the higher-level protocol does the equivalent of saying "I get it now". In that case, the protocol, though aborted, did not fail. It failed if it was aborted by one of the partners because they perceived that it was not reducing the error in whatever perception they were controlling by using it. The same effect might occur if a person quit hammering to control a perception such as driving a nail into a rock when the hammering had no apparent effect on the nail or the rock.

A protocol is a way of communicating with reasonably high confidence a meaning in the form of a message from one entity to another. Either entity could be a human, a non-human living control system such as a dog, or might even be electronically embodied in silicon.

The form of the protocol is the same no matter whether one or both of the communicating partners are humans, machines, or other living creatures. The protocol structure was initially developed for human-computer interface design, and was later found to be applicable with minor extensions to the symmetric human-human situation. We will develop it in a context of inter-human communication. It is here, rather than in the directly observable physical actions, that error reduction is the important point. The one with the information to communicate wants to perceive that the other has understood the information. The other's physical actions either succeed or fail in bringing the perception-reference difference within tolerance bounds.

The communicator of that message may be the Initiator who asked for the information, or who provided it because they wanted to perceive the Continuer to understand it (a reference value) and perceived that the continuer did not know it. Either way, the message has been passed (in a cooperative protocol as opposed to a deceptive one) when both partners perceive that it has been understood.

Often, especially at lower protocol levels, only two of the 19 possible controls result in any observable action, while all the others are already at their reference values. The performer of the first of these two loops we call the Initiator, who I often personify by assigning a name starting with "I", such as Isaac or Irene. Using a "Continuer" (who we might call Carla or Charles, for example) the Initiator controls some perception that is in error. The Initiator therefore acts in a way intended to reduce the error, by disturbing some perception controlled by the Continuer so that the continuer's control action reduces the error in the Initiator's perception. This pattern, between mother Cora and baby Ivan, is shown in Figure II.11.2 (lower part).

Here is a classic example of a completed basic protocol, in which only one party uses language, but each party performs some action, either speaking or some other kind of action:

Irene (in a wheelchair): *Please would you close the window*.

Carlo closes the window.

Irene: *Thanks*.

When any mechanism works smoothly and well, it may be hard to see anything of how it works. In this case, we have a perfectly ordinary exchange, in which Irene asks Carlo to do something for her and he does it easily. How many times a day do we engage in such interactions? "Please pass the salt", "Hold the door for me a moment", or a thousand other things that pass without a second (or even a first) thought, and are done without discussion.

We will use this simple interaction to probe a little deeper from a PCT viewpoint into what might be happening in such everyday interactions. We ask what each participant might perceive of the other's behaviour that disturbs a controlled perception in ways that result in acts that can be perceived by the first, creating feedback loops that correct different kinds of problems in the interaction. Our answers are certainly not definitive, but they may be suggestive.

Depending on the circumstances, the same "close the window" protocol might have been performed without either Irene or Carlo using language. Irene might, for example, have been able to pass her Primal Message by simply pointing at the window she wanted closed. The form of the protocol is independent of the way in which its functions are instantiated. As we shall see, all of its functions can be described as the control of one or more of the 19 controlled perceptions in the motif, 18 of them being perceived levels of belief, while only one concerns the content of the message itself.

Before Irene asked Carlo to close the window, she had a reference value to perceive it to be closed, whereas she perceived it then to be open. That error was her “Primal Message”. When he closed the window, the error vanished, and the protocol was complete. Always, protocols are completed when the initiator’s Primal Message (the error) is corrected and both participants perceive that to be the case[67].

When Irene acted, Carlo perceived something that disturbed some perception he controlled, perhaps a perception of Irene’s state of happiness. Irene relied on his control of some such perception, in the same way that baby Ivan’s use of their private language allowed mother Cora to know whether to feed Ivan or to find the pin that was sticking into him.

Of course, Carlo might have been ill-disposed toward Irene at that moment, in which case Irene displaying increased discomfort would reduce any error in his perception of her happiness. Both these hypothetical situations (Carlo well or ill-disposed toward Irene) illustrate a fundamental requirement of any protocol, the ability of each party to display to the other in a way that the partner can perceive accurately, perhaps categorically, what the other intends to display, which might possibly be deceitful. These requirements lead to the different complexities of the full 19-control loop form of the protocol. We shall deal with them gradually.

II.14.2 The Protocol Motif

Chapter I.8 described several of the simpler basic perceptual control motifs, recurring patterns of relationships among control loops, that perform some function as a group. Each motif provided an emergent property that was not expressed by any simpler structure of control loops, for example the “conflict” motif, which leads to “stiffness”, which in turn leads to “tensegrity” in a two-level hierarchic structure. Now we explore more deeply a more complex PCT motif, the Protocol, the only motif we have seen so far that extends across independent individual control hierarchies, as opposed to independent individual control loops. Later, we will see the Protocol Motif appear as a component in other social structures, in particular, the Trade Motif (Chapter III.9).

We repeat this example of an entire protocol:

Irene (in a wheelchair): *Please would you close the window*.

Carlo closes the window.

Irene: *Thanks*.

Irene’s objective in initiating the protocol is to reduce the error in some perception she is controlling (the Primal message, in this case the open-closed status of the window). To do so she uses Carlo as an atenfel, in the same way that the human operator is an atenfel in the control of the balance-scale weights in the Rube Goldberg system of Figure I.3.5.

One thing we can say nearly for sure is that Irene *displayed* to Carlo how she wanted him to act. I say “nearly” because there are many possible reasons she might have acted the same way. For example, she might have been testing whether Carlo would control for pleasing her by doing as she asked. But we can ignore these other remote possibilities and assume that she was controlling for perceiving the window to be closed.

Restating what we mentioned at the beginning of the last Section, how she wanted him to act (or as we

67. Notice that as stated here, the participants need to perceive errors directly, which Powers does not allow, but the Seth-Friston equivalent circuit connection for the hierarchy (Figure 7.2 in Volume 1) does allow. Later. however, we will find that direct perception of one’s own error is not required. It is the partner who perceives those errors.

shall see later, what she wanted him to learn) was the "Primal Message" that she wanted Carlo to perceive. The Primal Message is simply the perceptual error she wants Carlo to reduce or eliminate by his actions. What she displays to Carlo, however, is her "Primary Message", which in this simple interaction should be sufficient, if Carlo is of goodwill toward her, for him to do the deed that corrects the error that is her Primal message. In a more complex protocol, this is seldom the case. The Primal Message could be as complex as Professor Irene wanting to perceive Student Carlo as understanding the mathematics underlying quantum chromodynamics, a Primal Message that might take many years to transmit through a hierarchy consisting of many levels of supporting protocols.

Other than by closing the window herself, Irene can correct the error that is the Primal Message in two ways, only one of which uses a protocol. The other is to disturb some perception in Carlo in such a way that his control action has a side effect of reducing the error in her perception, without Carlo needing to perceive that this is what she is doing.

Military commanders try to do this with their enemies. They might act so that the enemy will control some perception by acting in a self-defeating way. In preparing for the invasion of Normandy in June 1944, the allies displayed to Hitler that they were preparing an attack through the Pas de Calais (Straights of Dover), by installing models of tanks and other invasion necessities in Kent, and creating suitable radio traffic for the Germans to intercept. The US General Patton commanded this formidable non-existent invasion army, to add verisimilitude. It was all part of a "standard" deceitful protocol used time and again over the millennia.

We are interested here, however, in *collaborative* control using a protocol, which implies that Carlo is a knowing and willing participant, acting with the intention of changing his perception that an error exists in some perception Irene controls. In the example, he correctly perceives that Irene controls for perceiving the window closed, which at present it is not.

In order for Carlo to act to reduce the error in the perception Irene controls, he must discover what that action might be. In other words, he has to discover what Irene wants or what she wants him to do. If Irene is to get what she wants from him, she has allow him correctly to perceive what she wants him to understand or to do. She must successfully communicate her "Primal Message", in this case that she should perceive the window to be closed.

II.14.3 The Three Basic Propositions

In analyzing the progress of a protocol, we consider the three propositions we presented above, all of which both parties must believe to be true if the protocol is to complete. Once again, they are:

- P1: *That the continuer has understood the Primal Message.*
- P2: *The quality of the communication process is sufficient for an adequate interpretation*.
- P3: *That it is not worth continuing to transmit this message.*

Each of these propositions apply to protocols whose Primal Message is at any level of complexity, from a simple Yes or No to the generally accepted Theory of Quantum Chromodynamics. Let's consider what they mean, in inverse order. P2 and P3 describe conditions under which one or other partner does not believe sufficiently strongly that the Primal Message can be or needs to be properly understood by the continuer (Carlo in our example) and that something must be done to improve the likelihood that it will be.

If P3 is believed, one of two states must be true. Either the Initiator (Irene) has by some other means eliminated the error whose description is the Primal Message, or one or other partner believes that understanding cannot be achieved, at least not using this protocol, and the conversation should stop. Even if one partner believes P3, the protocol is not complete if the other partner continues to disbelieve P3 and continues the conversation.

Carlo: "*I can't make any sense of what you are saying*".

Irene: "*You seem to be getting pretty close. I think it's worthwhile trying a little longer*"

Although Irene is the Initiator of the main protocol, Carlo initiates the protocol tells her that he wishes to Abort the protocol. Before Carlo speaks, Irene does not believe he wants to stop, and Carlo believes that Irene does not believe he wants to stop. Carlo cannot control her belief, though he may be able to influence it by what he says. He can control only his belief about her belief, which he can do by saying something that disturbs Irene's perception of his belief state,

Continuing this interchange, Irene lets Carlo perceive that her belief state about P3 has moved from disbelief to uncertainty or weak belief. She starts a new supporting protocol which we say is at a lower protocol level, in the same way that lower level perceptual functions support perceptual functions at a higher, more complex, level.

Irene: "*Are you telling me you want to stop*?"

Carlo: "*It might be best*"

Irene: "*OK*"

At this point both of them believes that the other believes that Carlo believes the dialogue should stop. Irene's final "OK" allows Carlo to believe Irene also believes it is not worth continuing. Both now believe P3, that the other believes P3 and that they believe P3. The Abort is complete, through the performance of these three levels of supporting protocol.

The same kind of supporting protocols may be appropriate when the abort occurs because Carlo understood the Primal Message.

Carlo: "*I think I've got it now*"

Irene: "*Convince me*"

Carlo paraphrases or offers a consequence of what Irene has been trying to explain. This may take seconds or it might take hours or even months.

Irene: "*OK*"

Just as with the "impossible to complete" reason for an abort, the "Primal Message understood" requires Irene and Carlo both to believe P3, believe that the other believes P3, and believe that the other believes that they believe the other to believe P3.

We move on to P2.

For Carlo to understand the message, the conditions must allow him to perceive what he needs to perceive about Irene's states. He may need to speak the same language, have a sufficient knowledge background, be attending to her, and so forth. If Carlo is unable to interpret Irene's different displays, for whatever reason, he will not be able to control for doing what she wants. The second proposition is thus:

P2: *The quality of the communication process is sufficient for an adequate interpretation*.

If P2 is not believed by either partner, they may initiate a loop in which the partners reorganize (or renegotiate) their means of communication. The problem might be that one party cannot see or hear (or feel) the effects of the other's actions because of ambient background noise, insufficient lighting, or that they just had a numbing injection at the dentist. Or perhaps at the level of this protocol, Carlo does not understand what is being communicated. Maybe Irene is unilingual in, say, English, while Carlo doesn't speak English.

At a higher level, Carlo may not have the mathematical background required to make sense of the equations involved in what the Irene is trying to explain. Perhaps Irene is trying to describe General Relativity, but Carlo never learned tensor calculus. Whatever the reason, the protocol cannot complete successfully unless something about the communication process changes. The solution here is usually some form of reorganization, perhaps as simple as moving to a quieter location.

There is not much more to be said about P2, but the same three levels of recursion in the control of belief apply in the same way they do for P3. When each believes (1) Other believes P2, (2) Self believes Other believes P2, and (3) Self believes Other believes Self believes P2, all the controlled belief perceptions are at their reference values, and no further action is necessary. In most everyday situations, all these beliefs are at their reference values before the protocol begins, and no action is required. This is almost never true of P1, though hypothetically it is possible.

P1: *That the continuer has understood the Primal Message.*

When Carlo believes he has fully understood the Primal Message and Irene believes he has, the protocol is complete. Whether Carlo has correctly understood Irene's Primal message, whatever its simplicity or complexity, is a matter of *fact* to which neither party has access. All either of them has is a level of belief in whether or not he has. That level could be anywhere between strong disbelief and strong belief, so the condition for ending the protocol could be restated as "Both Irene and Carlo have a strong belief in P1", believe the other also does and that the other believes they do, too. Again we have the same three levels of recursion.

In principle, no matter which of the three propositions is concerned, Taylor, Farrell, and Hollands (1999) reported finding occasions in natural, "Wizard of Oz", or synthetic dialogues in which all three levels of "I believe that you believe that I believe" occurred, but never a fourth level. One might surmise that each level of recursion introduces uncertainty into the situation and the uncertainty would be excessive at a fourth level, or that three levels in each partner are in principle enough to bring the parties to the same belief about each other's belief in the proposition. Whether these surmises have any relation to the truth is something that requires further research.

If either Irene or Carlo does not strongly believe that the Carlo has understood the Primal Message, that party will continue to control his or her level of belief in P1 until it reaches a reference value we might call "strong". The whole of the operation of the protocol consists of achieving this state at which both party's perception of their own state of belief in P1 is within its tolerance bounds of the reference value "exact knowledge", to which "strong belief" is the closest possible approximation.

Carlo's action in closing the window might be a test of his weak belief that Irene wanted it closed, just as in the Test for the Controlled Variable (TCV) the Experimenter disturbs what might be the Subject's controlled variable to discover whether it is, and what is the environmental equivalent of the reference value if so. Irene's "Thanks" confirms that his weak belief was correct. At that point, Irene believes P1, Carlo believes P1, Irene believes Carlo believes P1, etc, and if all these recursions in both partners are true of P1, they necessarily are true also of P2 and P3, completing the protocol.

Controlling beliefs in PI is the only place in the protocol in which to control may require any reference to the content of the Primal Message. For example, if Carlo isn't sure whether Irene wants the window to be closed or to be opened wider, he might signal his uncertainty with a display of his own, a Primary Message in a supporting protocol such as waving a hand upwards and then downwards near the window with a questioning expression. Irene would then perceive "Carlo does not believe P1" and "Carlo now believes Irene believes Carlo does not believe P1" both of which she might be able to correct by making a downward hand gesture. If Carlo then closes the window, Irene has corrected the error that led to her issuing her Primal Message, and the protocol is successfully completed once she has allowed Carlo to perceive "Irene believes P1".

Any way of ending the protocol, other than a successful completion, is considered an "Abort" of the protocol. An "Abort" overrides P1, terminating the protocol even if one or other party disbelieves P1. In our "close the window" example, Irene might send an Abort if Joe walks by and closes the window to control one of his own perceptions before Carlo even finds out what Irene wanted. If that were to happen, the entire protocol might go like this:

Irene (in a wheelchair): *Please would you close the window*.

(Joe walks by and closes the window).

Irene (to Carlo): *Never mind*.

II.14.4 Control of belief: the Syntax of a Protocol

'Twas brillig, and the slithy toves
Did gyre and gimble in the wabe:
All mimsy were the borogoves,
And the mome raths outgrabe.

(Lewis Carroll*, "Through the Looking Glass",* 1871)

This first verse of Carroll's "Jabberwocky" has good syntax, but does it convey much meaning? To some extent I can believe that it does, though all we know is that "*slithy toves*" and probably "*mome raths*" are animate beings, as perhaps are "*borogroves*". We can determine that some words mean properties of things designated by other words, such as that something about the environment was "*brillig*" and that "*borogroves*" can be "*mimsy*". The verse is not wholly unintelligible, even though one cannot say in any detail what it all means, other than that it seems to describe a rather pleasant environment apart from those *slithy* toves.

But does this amended version, which changes nothing except that it substitutes six nonsense words that appear to represent a property or an activity? I think that this version seems to be about a rather unpleasant, even forbidding, place with nasty inhabitants, even though the dictionary meanings of the six words are equally empty in both versions.

'Twas trobug, and the sloothy toves
Did fode and wumble in the wabe:
All doozly were the borogoves,
And the mome raths offgroze.

If you agree with my perceptions of the two verses, which differ only in that the revised version substitutes six nonsense words of the original with six rhythmically matched different nonsense words, then we can see that it isn't simply the syntax of the clauses in the verse that offers the pleasant or unpleasant impression of the environment described. Why? I suggest the reason is phonetic symbolism (Section II.13.9). But this is irrelevant to the point I want to make, which is that we can loosely or fuzzily divide the words in the verse into "content words" and "function words".

In making my substitutions, I made a further mental subdivision, between "entity words" and “property words”, and changed only what I considered Carroll intended as “property words”. This is a distinction based on PCT. As discussed at length in Section II.1.8 in Volume 1, we never control entities as such. We control or use for control only their properties, such as their location, their colour, and so forth. When we were dealing with Black Boxes and White Boxes in Volume 1 (Chapter 10) we discovered that the only aspect of Real Reality that we can discover is how what we perceive as entities can link to each other, performing different functions. These functions are the properties described by the words I changed.

In Volume 1, a property is a potential atenfel, while an entity or object is a frequently encountered bundle or syndrome of potential atenfels, an atenex. The distinction is somewhat fuzzy, as a property such as “colour” could itself be treated as an entity, with many possible uses in controlling different perceptions, but these usually turn out to be available only when the property is an attribute of different entities — property bundles or syndromes. Here, the entities (toves, wabe, borogroves, and mome-raths) have no properties already in the mind of the reader, so their characteristics are entirely defined by the property words. In most written or oral communication, this is not the case. The reader or listener already has some idea about some properties of most of the entities mentioned.

Here we begin to deal with the syntax of a protocol. Consider first the trivial protocol with which we started this Chapter:

Irene (in a wheelchair): *Please would you close the window*.

Carlo closes the window.

Irene: *Thanks*.

Irene is controlling for the open-closed state of the window. Her reference value for that perception is “closed”, and to convey that fact is her Primal Message, which she instantiated in a verbal “Primary Message”. As we noted earlier, her Primary Message could equally well have been a gesture, or could have been elaborated to include a reason she wanted the window closed. But Carlo understood some properties of a window, among which was that it had a potential state of being open or being closed.

Irene believed that Carlo also knew whatever property of the window he could use to close it, lowering the sash, swinging the entire frame, pressing a button, or whatever. Irene did not need to know how the window would be closed, so long as she believed that Carlo did, and could understand that window-closure was what she wanted to happen. Irene was controlling for perceiving the window to be closed, but there were many ways she could have influenced that perception without involving Carlo. The protocol provided just one available atenfel for reducing or eliminating the error in her “window-state” perception.

Within the protocol, at a control level below her control of the open-closed state of the window, what she was controlling was her belief that Carlo understood the Primal (not Primary) message, with a reference value “strong belief” (which we may symbolize with a “B”, to distinguish it from her current belief state, which might be “b” for weak belief, “u” for uncertain, “d” for weak disbelief, and “D” for strong disbelief.)

If Carlo failed to close the window when she asked, there are two different kinds of possibility, each of which might be amenable to control by Irene. The first kind is based on the possibly that Carlo understands her, but simply does not want to do as she asks. We won’t pursue this possibility yet, because it deals with non-cooperative dialogue, but will concentrate on the second kind, based on Irene’s belief that Carlo did not understand what she wanted.

In other words, Irene has a weak disbelief in both P1 and P2, for both of which her reference value is “B”, strong belief. She needs to act, and her actions should disturb a perception Carlo is controlling, so that his action helps her to control her belief perceptions, not the perception that she wanted to control when she initiated the protocol.

Notice what we said here. We described a different kind of protocol, a protocol that above we called a supporting protocol. We introduced a continuum of levels of a *controlled belief perception*. Control of belief is all that supporting protocols do. A supporting protocol is analogous to a function word in linear text or speech, in that it conveys little or none of the content of the Primal Message. It is a component of the protocol syntax. Function words convey little or none of the meaning of the text in which they appear, whereas content words without function words often suffice for communication.

Every protocol begins when the initiator wants to control some perception through the action of another controller, so the Analyst cannot distinguish protocol types at that point. But the readiness of the potential continuer to execute the protocol may be anywhere between having overtly displayed readiness — R-display — through acceptance of the invitation created by the initiator's disturbing influence, to rejection or failure to recognize that an invitation occurred. We call an invitation intended to produce R-display an "Interrupt". An Interrupt allows the potential continuer to perceive the initiator's intent to start a new protocol. We can ignore R-Display in the GPG itself, since it is merely a starting requirement to be satisfied before the Initiator starts sending the Primal Message.

Before considering this highest level protocol, the level at which the initiator is acting to control the perception that was initially in error, we should consider what we might call "function" protocols — protocols that implement various arcs in the General Protocol Grammar (Figure A8.1). Some arcs in the diagram are shown with heavy lines, some with light lines. The ones with light lines correspond roughly to function words in a sentence, while the ones with heavy lines correspond to content words such as nouns and verbs (Section II.12.3).

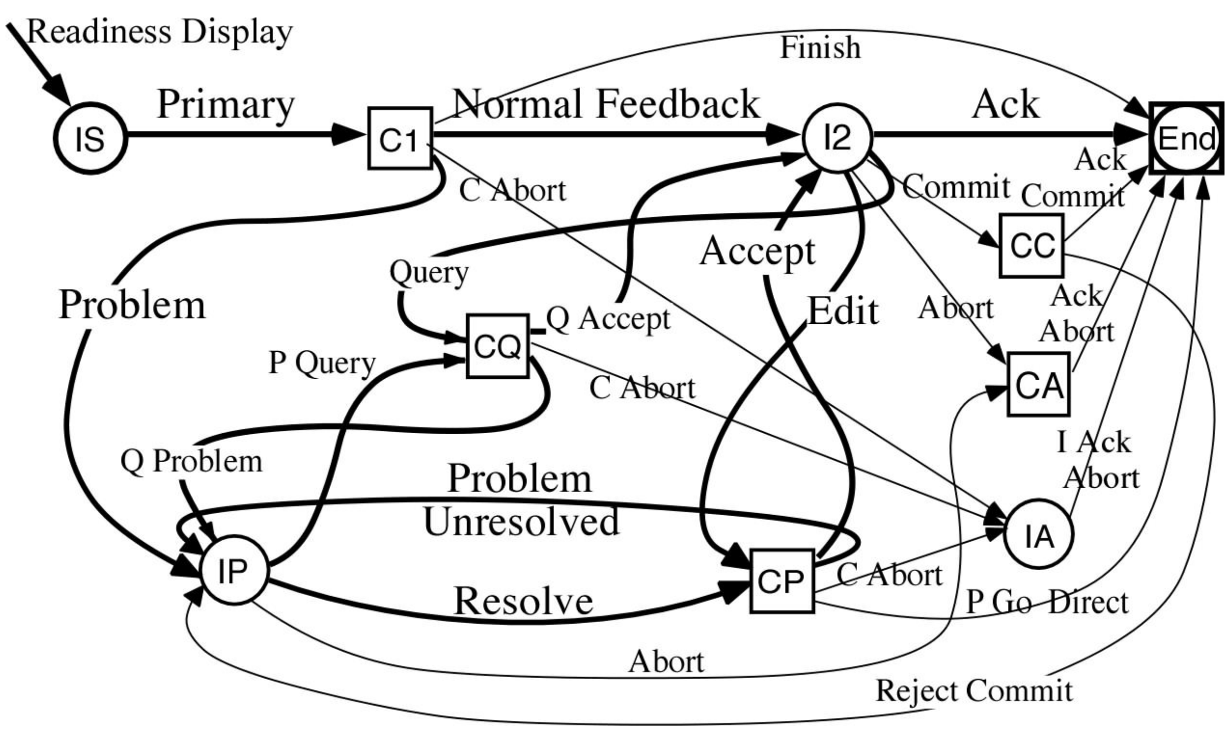

Figure qII.14.1 The General Protocol Grammar, distinguishing between arcs (light lines) that are purely functional and those (heavy lines) that might require arbitrary content to be displayed by one or other partner. Circles represent the continuer controlling a perception, while boxes represent the initiator controlling a perception. The arrows indicate direction of influence, not necessarily sequences of events, since both partners continuously control their perceptions.

Every arc in the GPG diagram represents one partner influencing some perceptual state in the other, but some kinds of function are common to several arcs. For example, there are many arcs labeled “Abort”. In each case, the supporting protocol consists only of a display by the appropriate partner that they do not wish to complete the protocol, while acknowledging that it is incomplete. Arcs labelled “Ack” on the other hand, are implemented by a display that the partner accepts and agrees with whatever the other displayed — either that the protocol should be aborted, or that it has been successfully completed. Other arcs, marked by heavy lines in Figure A8.1 might require the inclusion or modification of some specific message content.

Each arc is implemented by a complete lower-level GPG unless the arc has a “null” instantiation, and every GPG exists so that the initiator can control some perception through the actions of the continuer. Consider now the following three types of function, and how they might be instantiated: Abort, Acknowledge, and Accept. “Null” means that no action is required because the perception being controlled at that point is at its reference value. But if the instantiation is not “null”, what perceptions are controlled by the GPGs that implement these function arcs? They are all perceptions of the progress of the protocol itself, and not of the variable content, in the same way that function words in a sentence clarify the relationships among the variable content words. To instantiate them requires that the partner be able to perceive the intent of the one using the arc: “I want to stop”, “I acknowledge what you just communicated”, or “I accept what you just communicated”. The actual implementations of these messages must depend on the situation.

Just as “John gave Mary a book” has a surface structure similar to “The sunset gave the beach a glow”, so do all “Give me X” protocols share a structure, whether X be a “the ball” or “a smile” or “some indication you understand what I am saying”. The function arcs have no “X”. The content arcs usually do, in what is called an “Inform” instantiation of the arc. There are only a relatively small number of different distinguishable states and transitions in the GPG, just as there are only a small number of function words or morphemes in a language. On the other hand there is an indefinitely large number of content arcs, as there is a constantly changing inventory of content words of a language.

If function arcs are implemented by “Function GPGs”, then the content-carrying arcs might be said to be implemented by “Content GPGs”. Let us consider them briefly.

The “Primary” arc consists of a display by the initiator (call her Irene) that identifies the kind of protocol Irene wants to “dance” with Carlo, and may also display the perception being controlled and an action the continuer (Carlo) could take to correct its error. In a very simple case, Carlo might simply perform the desired action, which would be a case of “Normal Feedback”.

But suppose Professor Irene is controlling a perception of student Carlo’s understanding of tensor calculus, for which her current perception is that he understands vector calculus pretty well, but tensor calculus not at all. Irene is unlikely to initiate the protocol by delivering a full-semester’s lectures all at once. More probably Irene will provide little bits and give Carlo the opportunity to display how well she has understood the message so far. She will use Normal Feedback initially to show Irene something of what she knows, the content of which allows Irene to perceive her level of understanding. From there, they will probable continue a semester-long process of looping around the “Edit-Accept” loop until either

Carlo, in effect, uses the Accept arc to say "I've got it" and Irene uses the "Ack" exit from the protocol, or one of them uses the "Abort" exit from the loop, because it is the end of term or because one of them thinks she will never get it, at least not the way Irene is teaching it.

Of course, there are indefinitely many other possibilities, but these extreme cases might give a flavour of how the function and content arcs of the GPG represent the dynamically varying errors of the relevant controlled perceptions in the initiator and continuer.

Supporting protocols are structured exactly like main protocols, but their Primal Messages always are a degree of belief in one of the three propositions in the supported protocol, whereas the main protocol Primary Message carries most or all of the Primal Message content. If, however, the supporting protocol supports P1 (that the continuer understands the Primal Message), then the supporting protocol may well include some content as an indicator to one or other of the participants of the degree to which the Primal Message has been understood by the continuer (Carlos in the example case).

The only type of control loop that involves the perception of language by another person is a protocol loop. In a protocol loop, trajectories in feature space similar to those we discussed in the Story of Len and Rob and Len's family (Chapter 1) offer a major way one partner can perceive states in the other. In human terms, that means that the protocol works best when the parties display clearly to each other any issues, issues being deviations of controlled perceptions from their reference values. For the most part, language is a good medium for such displays.

The relation of the trajectory to the perception in question is initially completely arbitrary, as we discussed in connection with the development of the private language of mother Cora and Baby Ivan (Section II.11.3). Only through reorganization does a trajectory (e.g. a word or phrase) become a label display for the perception to be controlled, and only if reorganization is successful in producing a protocol loop that effectively controls the relevant perceptions in each partner will the trajectory be retained as a label to be recognized.

The "Give me X" protocol we used as an example in the "Story of Rob and Len" is one of many possible protocols that serve as atenfels for the control of various perceptions. However, if the protocol is to construct a feedback path for the control of a perception, the initiator must be able to perceive that a potential continuer is able to perform the desired action.

The initiator must also be able to disturb a perception in the continuer so that the continuer actually produces the action that reduces the error in the initiator's perception. This is true whether the potential continuer is perceived as cooperative, neutral, or even maliciously disposed toward the initiator. Indeed, one of the skills of a military commander or a judoka is the ability to get the opponent to do something that provides the commander with an advantage, as Sun Tzu noted 2500 years ago. Whether the Initiator is honest or duplicitous, the display should disturb the intended perception in the intended way if it is to be effective for the Initiator's purpose.

II.14.5 R-Display and Interrupt in the General Protocol Grammar

R-Display is a name I prefer for a construct introduced under the name of E-feedback to Layered Protocol Theory by Engel and Haakma (1993). E-feedback displays that the continuer is prepared to do something. Engel and Haakma used as examples the shapes of switches and knobs on electronic and mechanical devices. For example, a round knob, particularly one sporting a pointer, suggests that turning it might select among several possibilities or a continuous range of values that affect the way the device works, whereas a “switch-shaped” object suggests the possibility of changing between two possibilities, such as “On” and “Off”. The shape of the object seems to say something like ”I’m ready to help you” plus “Turn me and vary something smoothly” or “Flip me and something will change abruptly”.

When the continuer is a human, R-Display is more complex. In a shop, if you see a person wearing a uniform or a badge standing behind a cash register, they are perceived to display readiness to engage in a “Buy-Sell” protocol that is part of their “Cashier” Role. If someone is sitting with their back to you and you speak their name, they may turn to look at you or say “What”, “It’s you again”, or in some other way indicate readiness (“R”) to listen. You have executed the very first act of a protocol, an “Initiating Interrupt”, or simply an “interrupt”.

The kind of protocol for which the interrupted person is ready, if any, may be indicated by the form of the R-Display in the same way that the roundish shape of a knob suggests that it might be ready to be turned and the uniformed person standing behind the cash desk suggests they may be ready to take your money, or it may be more generic, awaiting the initiator’s display of the Primary Message in order to determine the nature of the protocol wanted by the initiator. Persons already showing R-display usually do not need to be interrupted before they continue with their side of the protocol.

Let us once again imagine Irene wants something of Carlo. Initially, Irene may or may not believe that Carlo is attending to her and is capable of understanding her communicative displays. Carlo may or may not know Irene wants something of him. However those may be, Irene does not believe P1 to be true. In the notation explained in Box 4, Irene(P1≠B). Since Irene’s reference value for P1 is Irene(P1=B) she has an error in a controlled perception, and acts.

Carlo may not have any beliefs about any of the propositions, but if he is displaying “R-Display”, Irene will perceive that he is ready to take the role of continuer and begin a protocol. If Carlo is not showing R-Display, a readiness for the kind of interaction she wants to initiate, Irene(P2<B). Her action, perhaps saying his name, is an “Interrupt” intended to disturb some perception Carlo may be controlling, in order to get him to act and show that he is attending to her, one of the necessary components of P2.

Carlo’s R-Display may be enough for Irene to believe both P2 and that Carlo also believes P2. In our notation, Irene(P2=B & Carlo(P2=B)=B)). If these beliefs are not at the level of “B”, then Irene will largely be at node IP in the GPG diagram of Appendix QQ. She will be largely there, but not necessarily completely there, because she may only be uncertain about her own or Carlo’s belief in the truth of P2. Rather than trying to bring her belief in P2 to a level of “B”, she may try to offer a Primary Message, perhaps using simple language if she perceives Carlo to be a foreigner, or speaking loudly if they are in a noisy environment. Carlo’s actions following this will show her the success or failure of her Primary Message. If it succeeds, then she will raise her level of belief in P2, but if it fails, she may then act from IP and try to fix the problem, perhaps trying a different language, or writing her message rather than speaking it.

At a rather high protocol level, a student might approach a professor and ask if she might offer a tutorial on quantum chromodynamics or special relativity. That would be the Interrupt. The professor might choose to provide R-Display by agreeing or might not, perhaps by pointing out that she was not a physicist, or did not give private tutorials. If she provided R-Display, the tutorial’s information transmission protocol might be initiated not by the teacher offering an initial introduction (a “Null” Primary Message), but by the student revealing his current state of knowledge on the subject, and would continue through weeks or months of using the Edit-Accept loop, until either they both decided he had learned what he wanted to learn, or that the protocol should be aborted, perhaps because the school term was over or because they agreed that the topic was beyond him.

If an overt interrupt is used in order to induce R-Display, it allows the continuer to perceive that the initiator wants to begin a new protocol, whether or not the participants are already conducting a protocol interaction. To start a protocol without an Interrupt may sometimes be considered impolite, because the potential Continuer may perceive the Initiator to believe the continuer to be always prepared to do what the Initiator wants, even in the absence of R-Display, and that perception might be controlled at a different reference value. If Jimmy wants to borrow the car, he is more likely to get it if he says "*Dad, can I borrow the car?*" than if he simply comes in and says "*Can I borrow the car?*". The "Dad" interrupt is functional!

After the Interrupt the protocol may get complicated. To describe what is likely to happen, the Layered Protocol theory employs a diagram called the General Protocol Grammar (Figure II.15.1) that looks like a state transition network but is actually more of a flow diagram. To say that the Initiator or the Continuer is "at" a particular node is to say that they believe the person identified by an "I" or a "C" in the node label is likely to act. At higher protocol levels, it is quite normal for both parties to be acting simultaneously, but at lower levels such as in an unemotional debate, often only one node is occupied, as both parties agree whose turn it is to act.

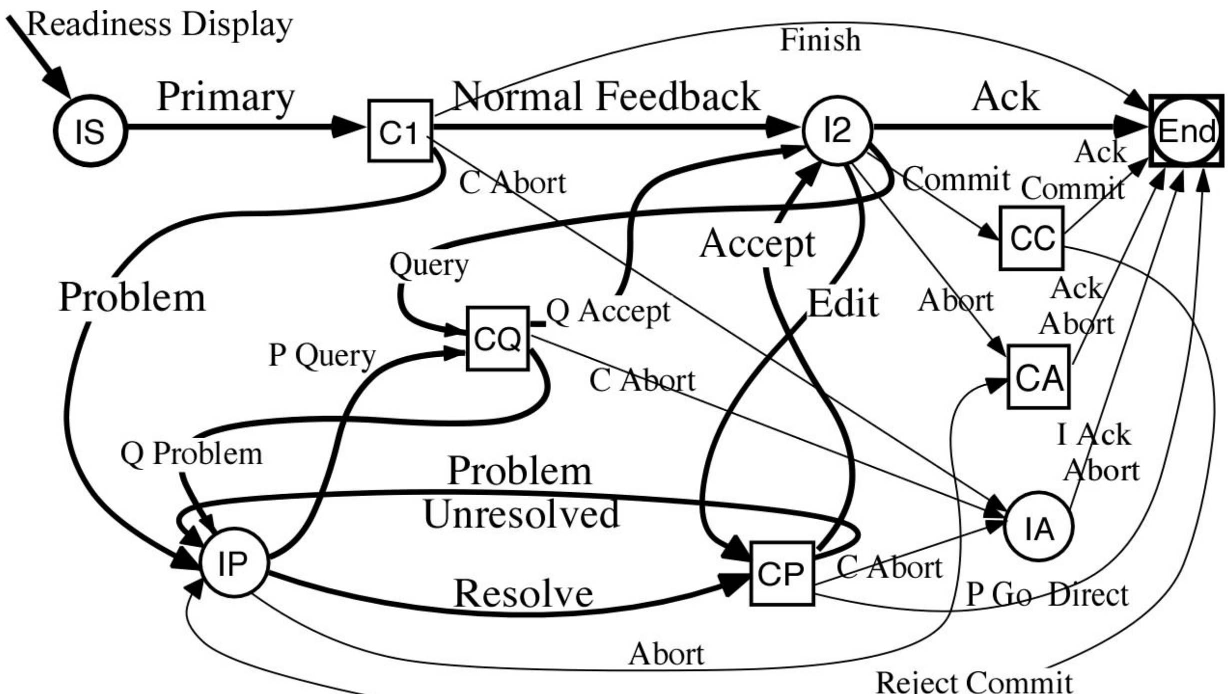

Figure II.15.1 The General Protocol Grammar, distinguishing between arcs (light lines) that are purely functional and those (heavy lines) that might require arbitrary content to be displayed by one or other partner. Circles represent the continuer controlling a perception, while boxes represent the initiator controlling a perception. The arrows indicate direction of influence, not necessarily sequences of events, since both partners continuously control their perceptions.

Let us examine a node in the diagram. we could use any node but I choose the one labelled I2, and will continue calling the Initiator and Continuer Irene and Carlo respectively. It is probably easiest to explain the diagram in the context of dialogue analysis, because there is little overlap in most dialogues, the parties for the most part taking turns to speak.

To say Carlo perceives the dialogue to be at I2 is to say he expects Irene to act in some way she chooses between arcs labelled “Ack”, “Edit”, or “Query”. All of these are shown with heavy lines, because they usually incorporate something relevant to her Primal Message. Carlo also knows she might indicate either a wish to Abort the protocol or to indicate that she thinks Carlo has adequately understood the Primal Message (the “Commit” arc).

If Irene also perceives the dialogue to be at I2, she has some perception that is not at its reference value and she acts to reduce the error. If she uses “Edit” she is telling Carlo that she understands something of what he understands of the Primal Message, and he is wrong or incomplete in that understanding, a problem that she may reduce by adding, explaining, or correcting something she perceives him to believe about the Primal Message. She would probably expect Carlo to perceive that the protocol should move to CP.

If she perceives herself not to understand what Carlo seems to perceive himself to understand, she will try to move the protocol more to CQ by including in her actions some pointer to where she is having a problem, but if she does understand and believes that Carlo believes P1 (and both other propositions) she will use the “Ack” arc, with content that shows Carlo what she perceives about what she believes. She could also use the Commit or Abort arc to signal that she considers the protocol satisfactorily completed, or that she wants to stop.

Most of the arcs in the GPG have different “instantiations”. To illustrate arc instantiations, let us consider the arc from C1 to I2, called “Normal Feedback”. In this arc, Carlo indicates to Irene what he understands of her Primal Message.

The simplest instantiation of Normal Feedback is “Null”, where Carlo does not act at all because the situation is such that he has a strong belief that Irene has a strong belief that he has understood. This is normally the case at low levels either of perception or of protocol. For example. Carlo will seldom confirm that he has understood each individual word that Irene speaks, but may confirm that he believes he understands what she has said.

The next simplest instantiation of Normal Feedback is “Neutral”. Carlo indicates that he has understood what Irene said but gives no indication of what he understands. Perhaps he simply nods his head.

Another kind of instantiation is “Inform”. Carlo may be unsure whether he has correctly understood, or may believe that Irene is unclear whether he understood correctly. He lets Irene know some or all (possibly in paraphrase) of what he did or what he things he may not have understood.

The final kind of Normal Feedback is “Correction”. Carlo believes he understood what Irene actually said and that he understands what she meant to say, which is different. A typical “Correction” instantiation might include the words “Don’t you mean” or equivalent.

Summed over all the different arcs of the GPG, Taylor, Farrell and Hollands (1999) counted 47 different arc instantiations that had a reasonable probability of being used in normal interactions, or fewer when one of the partners is electronic, at the present stage of development. Few of the arcs have as many potential instantiations as does the Normal Feedback arc, but many have two or three reasonable possibilities.

Each arc is implemented by a complete lower-level GPG unless the arc has a “null” instantiation, and every GPG exists so that the initiator can control some perception through the actions of the continuer. The initiator of this supporting GPG is the continuer of the higher-level, supported, GPG if the support is for an arc directed out of a Cx node. Typically, a supporting GPG has a Primal Message for which the primary content refers to some belief perception, but if it supports an arc containing some of the top-level Primal Message content (heavy line in the diagram) it may refer to that content itself.

Consider now the following three types of function, and how they might be instantiated: Abort, Acknowledge, and Accept. “Null” means that no action is required because the perception being controlled at that point is at its reference value. But if the instantiation is not “null”, what perceptions might be controlled by the GPGs that implement these function arcs? Their Primal Messages are all about perceptions of the progress of the protocol itself, and not of the variable content, in the same way that function words in a sentence clarify the relationships among the variable content words. To instantiate them requires that the partner be able to perceive the intent of the one using the arc: “I want to stop”, “I acknowledge what you just communicated”, or “I accept what you just communicated”. The actual implementations of these messages must depend on the situation.

Just as “John gave Mary a book” has a surface structure similar to “The sunset gave the beach a glow”, so do all “Give me X” protocols share a structure, whether X be a “the ball” or “a smile” or “some indication you understand what I am saying”. The function arcs have no “X”. The content arcs usually do, in what is called an “Inform” instantiation of the arc. There are only a relatively small number of different distinguishable states and transitions in the GPG, just as there are only a small number of function words or morphemes in a language. On the other hand there is an indefinitely large number of content arcs, as there is a constantly changing inventory of content words of a language.

If function arcs are implemented by “Function GPGs”, then the content-carrying arcs might be said to be implemented by “Content GPGs”. Let us consider them briefly.

The “Primary” arc consists of a display by the initiator (Irene) that identifies the kind of protocol Irene wants to “dance” with Carlo, and may also display the perception being controlled and an action the continuer (Carlo) could take to correct its error. In a very simple case, Carlo might simply perform the desired action (e.g close the window), which would be a case of “Normal Feedback”.

But suppose Professor Irene is controlling her perception of student Carlo’s understanding of tensor calculus, for which her current perception is that she understands vector calculus pretty well, but tensor calculus not at all. Irene is unlikely to initiate the protocol by delivering a full-semester’s lectures on tensor calculus in a single stream. More probably she will provide little bits and give Carlo the opportunity to display how well he has understood the message so far.

Carlo will use Normal Feedback initially to show Irene something of what he knows, the content of which allows Irene to perceive his level of understanding. From there, they will probable continue a semester-long process of looping around the “Edit-Accept” loop until either Carlo, in effect, uses the Accept arc to say “Now I’ve got it” and Irene uses the “Ack” exit from the protocol, or one of them uses the “Abort” exit from the loop, because it is the end of term or because one of them thinks he will never get it, at least not the way Irene is teaching it.

Of course, there are indefinitely many other possibilities, but these extreme cases might give a flavour of how the function and content arcs of the GPG represent the dynamically varying errors of the relevant controlled perceptions in the initiator and continuer.

When a student or a linguist analyzes, say, a sentence, marking parts of speech, relationships, and so on in whatever kind of grammar they use, they are doing consciously and deliberately whatever the person who produced the sentence did non-consciously. Nobody constructs a grammatical sentence in their native language by referring to a grammar book. They have reorganized so that the perceptions they control of relationships within their native language allow their interlocutor or reader to understand their intent. We speak grammatically, even though the grammar used among members of a street gang is unlikely to be the same as the grammar used among academicians using “the same” language.

This of usage versus analysis is equally true of the General Protocol Grammar (GPG). The participants in a dialogue interaction will not consciously act formally to change levels of belief in all the perceptions they control, but they will have reorganized so that the protocol interaction is likely to flow smoothly by moving perceived beliefs toward their reference values of complete belief in all three propositions. The Analyst is the one who can assign Irene and Carlo to different nodes, based on the degree of belief the Analyst knows each has in the perceptions they control.

When we say "*From Irene's viewpoint, she is mostly at Node IS*", it is in the same vein as a student saying that "Jack" is the subject of the sentence. The speaker doesn't think about subjects of sentences, but uses words in ways that allow Analysts to make consistent statements about them. When Irene is "at" or "mostly at" a node, the implication is that she will probably act in one of a selection of ways that the GPG assigns to that node. All of these ways are control actions to bring some of her perceptions nearer to their reference values, but Irene consciously knows nothing of that. She is just using her perceptual control hierarchy in the way she has reorganized to use it. She is "at" a node, but it is the Analyst who knows she is.

Most of the GPG is concerned with Carlo's increasing understanding of what Irene wants, and of his understanding of Irene's beliefs about his level of understanding, whether seen from his or from Irene's viewpoint. What Irene wants might be very complicated, requiring long sessions that could span hours, days, or years to complete the protocol, or it might be very simple, such as letting him know dinner will be at 6. It might be completed without a word being spoken, as might happen if Carlo were a cashier to whom Irene presented a purchased item, which Carlo rang up and for which she paid cash, or it might involve complex verbal arguments with explanations and questions as Carlo learns more of what she tries to get across.

No matter the complexity or duration of the protocol, the progress through the GPG is the same after Irene's acceptance of Carlo's R-Display (usually an "Ack" with a null instantiation), from IS to the point where all the belief perceptions about the three propositions have their reference values. At IS, Irene is uncertain about Carlo's understanding and about his belief in his understanding, whereas at I2 she knows more about both. At C1, Carlo's actions are to display to Irene his understanding and his level of belief in his understanding. At C2, they are similar, but rather than displaying the levels, Carlo displays the changes in his understanding and in his level of belief in his understanding, to allow Irene to perceive whether her actions are having the desired effects of reducing the error in her Primal Message.

If there are communication problems, Carlo may indicate that to Irene at C1, so that she moves to IP and begin sorting out the problem. For example, he may not be able to hear clearly what she said, and might show her his hand cupping his ear. She might talk louder, or might signal with her hand that they should move to a quieter place. Solving the communication problem has no categorical types in the way normal feedback has. At every protocol level the problems are different, though not understanding the language is a possibility at many different levels, from individual words to the dependence of special relativity on tensor calculus.

A protocol starts and finishes. It starts when the initiator disturbs some perception controlled by the continuer, and finishes when the initiator's perception is no longer in error and the continuer's perception is no longer in error after having been disturbed by the actions of the initiator. But for a protocol to work effectively as an atenfel for the initiator, both parties must be able to perceive that each is performing a role in the same protocol. The protocol format is likely to become a controlled perception, perhaps at the sequence level in Powers's hierarchy, so if either partner deviates from the other's reference for how the protocol should proceed, some corrective action might be expected.

The reference macrostate, or tolerance range (Section 5.1) in either partner for control of the format may, however, be fairly large, within which there may be many acceptable microstates, particular ways of performing the protocol. The requirement always is that both parties control for making the protocol achieve the result desired by the initiator, in other words to bring the error value in the initial controlled perception to zero.

II.14.6 Primary Message Content

Thus far, other than a few hints to the contrary, we have mostly treated protocols as though the perception being controlled by the initiator involves the partners acting on something concrete in the environment that both of them can observe, such as by handing the initiator a ball, or closing a window. Far more often, however, the perception the initiator wants to control is something about how the other perceives or acts upon the world — what the other knows or believes or feels, whether in order to change it or simply to discover it. Perhaps the initiator wants to change the continuer's World Model. Let us call this class of protocols, following Shannon (Shannon and Weaver, 1949), "Transmitter–Receiver", or, more to the point, "Teacher–Student".

Changing the identities of our protagonists, Isaac (the Initiator) wants Connie (the Continuer) to know something, which means that he wants Connie to perceive the world as being in a certain state or as working in a certain way. Perhaps he wants Connie to know it because Connie asked for the information and he is controlling for perceiving her to be content, perhaps Isaac is a politician who perceives that Connie currently sees the world in a way that differs from his reference value for how he wants to perceive her seeing the world (so that she would vote for him). In Volume II of this book, which concentrates on power and politics, we follow a few trails along which this latter possibility might lead.

Obviously nobody can perceive the complete set of memories and current perceptual values in another person. In fact nobody outside the person can perceive any of them. So how could Isaac control a perception of anything about Connie's background knowledge? The answer is that he can't, but he can perceive what Connie shows him about her knowledge, and can control that perception if his actions influence what she shows. To the extent that what he sees accurately reflects that small portion of Connie's knowledge, in principle, he could control it. *He can pass knowledge to Connie* (if she is controlling for accepting it). Passing knowledge, otherwise known as "communicating", is a class of protocol that requires no action on concrete parts of the mutual environment by either partner, but it is quite probably the class of protocol most widely used by humans.

The information Isaac passes to Connie alters some of her perceptions of the way the world is, whether they are based in imagination, in sensory data, or both. Connie also has a set of high-level reference values, which combine with her perceptions of the way the world is to create an equivalent set of error values. As her perceptions change, so do those error values, and consequently, so do all the reference values below the highest level at which a perceptual value changes, and the externally visible actions that depend on them.

If Isaac simply says that he has made a pot of tea, and Connie is not controlling for drinking tea, the changes might be very limited. She will not initiate the actions involved in accepting and drinking Isaac's proffered tea. She already perceived Isaac to be a nice guy, and his making tea does not change that perception, except possibly making him seem incrementally nicer. Or perhaps Connie has let Isaac know that she dislikes tea, and perceives his making a pot as an unfriendly gesture from not such a nice guy. Other changes might have far reaching effects on lower-level references. For example, if Connie was thirsty, she might change reference values all the way down to the muscles, and act overtly to accept and drink tea from a cup presented to her by Isaac.

One can never tell how far-reaching the effect of providing information might be in any particular situation, because one would have to know the other's current set of perceptual and reference values, as well as the functions that determine them. Nevertheless, people are often correct in providing information that results in the recipient of the information performing the desired action. If this were not so, no protocols could ever be developed. If Isaac were to say "Connie, it's getting late and we might miss the train if we don't hurry", Connie is likely to hurry, though if she were actually controlling for missing the train, she would not.

II.14.7 The meaning of "Meaning"

What might Isaac's comment about the train mean to Connie? There's no way of knowing precisely, but if she heard it at all, it changed some state within her. Perhaps what it changed was not at all what Isaac intended to change. For example, it might mean to her that Isaac is telling her that he believed she did not know the situation, and therefore that she thinks he thinks she is stupid. It might mean that Isaac is doing as she asked, by reminding her in case she might have been attending to something else entirely. It might mean to her that Isaac thought she did not know the time of the train. The possibilities are countless, but whatever they are in detail, fundamentally they are changes in her World Model, the Model that allows her to plan her choices of action. The world as she now perceives it is different from the world she perceived before Isaac spoke.

Suppose I were to tell you that King Isa Mei Salahud had abdicated the throne of Zingaria in favour of his younger son Zeo Mei Fanhuil, bypassing his eldest son Isa Mei Vidael. I would not be surprised if you were to respond "That means nothing to me."

Why would you be likely to say that? Everything I told you is in proper form. I told you about a very significant event in the political structure of a country, and possibly of its neighbours and allies, which might well change your World Model. Furthermore, it might tell you something about the structure of family names in Zingaria. Your World Model may well have changed, especially if you were familiar with Zingarian history, and yet you might reasonably say that my message is meaningless to you. Why?

Let's think again about the development of meaning for mother Cora and baby Ivan (Section II.11.1 and Section II.11.3). One pattern of Ivan's sound and movement may mean for Cora that Ivan is hungry, but a different pattern means he has a pin sticking into him. Cora may just notice that fact, and continue washing the dishes, and Ivan's actions will still have for her the same meaning, but we then ask how this meaning of the different effects of her different possible actions came into her mind. It was because she had been able to control some perception she had of Ivan by acting in a certain way, in the one case by feeding him, and in the other by finding and removing the pin. The effect of her actions in controlling her own perceptions was essential to the meaning.

Now, if Cora is controlling for perceiving her dishes to be clean, which conflicts with controlling for perceiving Ivan to be happy, and does not act on the latter, nevertheless Ivan's unhappy state becomes part of her World Model. In her model of "the way the world is", there may be many states that differ from the way she would like them to be, the politics and policies of Zingaria and its neighbours possibly among them. For many of those unfortunate states, she has no means of acting to bring them closer to her reference values for them. She cannot complete a control loop through the environment. But in the case of Ivan's state of happiness, she can, and when she perceives the dishes to be clean, she probably will attend to the state of Ivan's happiness.

Sir Arthur Conan Doyle provides another answer, which complements the answer we might derive from considering Cora and Ivan. In "A Study in Scarlet" (Doyle, 1887), he has Dr. Watson say of Sherlock Holmes:

> *His [Sherlock Holmes] ignorance was as remarkable as his knowledge. ... he was*

ignorant of the Copernican Theory and of the composition of the Solar System. ... "But the Solar System!" I protested. "What the deuce is it to me?" he interrupted impatiently; "you say that we go round the sun. If we went round the moon it would not make a pennyworth of difference to me or to my work."

and closely following,

"I consider that a man's brain originally is like a little empty attic, and you have to stock it with such furniture as you choose. A fool takes in all the lumber of every sort that he comes across, so that the knowledge which might be useful to him gets crowded out, or at best is jumbled up with a lot of other things so that he has a difficulty in laying his hands upon it. Now the skilful workman is very careful indeed as to what he takes into his brain-attic. He will have nothing but the tools which may help him in doing his work, but of these he has a large assortment, and all in the most perfect order. It is a mistake to think that that little room has elastic walls and can distend to any extent. Depend upon it there comes a time when for every addition of knowledge you forget something that you knew before. It is of the highest importance, therefore, not to have useless facts elbowing out the useful ones."

Holmes is in both quotes saying that it is not the content of one's World Model that matters, but how that content influences one's ability to control one's perceptions. To him, though it may be of use to an astronomer to know the Copernican Theory, to know it would be unlikely to assist in his ability to solve crimes. One might dispute his claim, but whether one does or not, he is saying that what is not "useful" is meaningless, at least to him.

We often attribute meaning to observations that have no human cause. To see snow on the ground "means" that it is cold outside. To see green shoots poking through a thin layer of snow "means" that the weather will soon be warmer. In ancient times many natural omens were interpreted (usually by specialized interpreters we call shamans or priests) as meaning that certain events were soon to happen if the enquirer were to take some action. Among the oldest known writings are the Chinese "oracle bones", on which were inscribed questions before the bones were cracked by heating. The form of the cracks held the meaning of the answers to the questions (often as ambiguous as those of the verbal messages produced by the Oracle of Delphi much later).

The answer to a question might be accepted as adequate by the questioner, but it might not. Under what conditions would it be likely to be accepted as an adequate answer? I suggest that the question would have been asked in order to control a perception that some aspect of the World Model is uncertain as to possible actions that might reduce some perceptual error, and the answer should reduce the uncertainty. If it doesn't it will not be accepted as an answer. But then one asks, in typical hierarchical PCT (and small child) fashion, "Why does it matter that this is uncertain?" The answer is often, perhaps always "because I can do this if the answer is X, but if it is Y I will do something else." Again, we come up against the notion that meaning has to do with present or anticipated control of some perception.

If the "meaning" of an observation is that it changes one's World Model, the role of protocols as communication channels for information, rather than as a means of getting someone to do something, becomes clear. The provider of the information controls some perception that would be closer to its reference value if the recipient uses it to change her World Model in a way that increases the likelihood that her future actions facilitate the provider's future ability to control some perception or perceptions.

More simply put, knowing that the recipient has the information, the information provider can do some things that he could not otherwise do so easily. To provide the information is to gain power, or as we discussed in Section 6.2, "worth". This fact is central to of the effectiveness of "The Big Lie" that we discuss later, mainly in the latter half of Chapter III.7. If information provided by the "Big Lie" is accepted as truth, the recipient's knowledge becomes one of the stabilities described in the extensive quote from McClelland in the introduction to Chapter II.4. There, we used the quote to emphasize the energy of the work required for the maintenance of structure in the face of normal entropic degradation. Here we use it to consider the effects of what is being maintained by the expenditure of this energy. The relevant extract from this quote is:

> *..work activities produce some kind of environmental stabilization, the creation of some atenfel, molenfel, or molenex, which can then be used in controlling other perceptions. Manual workers create stable feedback paths by manipulating physical objects; they build things, make things, and clean things up. Agricultural workers produce fields of crops and confinements full of animals to be used as food. Transportation workers move truckloads of products from factories to stores, where sales workers make those products available to customers in exchanges with predictably structured protocols. Service workers manipulate and stabilize the immediate physical environments of individuals, including their dwellings and even their physical bodies, as barbers and hairdressers do. Healthcare workers attempt to stabilize the physiological functioning of people's bodies. Educators strive to turn out classes of graduates with predictable abilities and skills, people who can then be hired to put their skills to work creating various kinds of feedback paths for others. Government workers maintain stability and order for the community in a wide variety of ways, from removing trash to providing and enforcing laws designed to regulate commercial transactions and maintain public order, and thus preventing large disturbances that would make control of other perceptions difficult.*

To this list of examples, we have added the stability of what we believe someone else knows. They may know it because someone they believe to be trustworthy told it to them, and we have observed their reaction to being told. We may use that stability in control of other perceptions, in the same way that we can use the knowledge that the sun comes up in the morning and sets in the evening to schedule the day's activities, despite being unable to control the timings of the sun's rising and setting, or even to certify that the sun will rise at all tomorrow.

The collectively controlled (Chapter III.1) stability of a formal ritual (Chapter III.5) is what gives meaning to the performance of a rite. The performance allows the observers of the ritual to know the status of the performers more stably than they might without the ritual, and can use that stability in control of some perceptions. The holder of an office through ritual appointment has a changed relationship to others who perceive that the ritual has been performed. A couple is perceived to be married because a ritual has been performed and someone perceived to be playing the role of the appropriate office with the appropriate authority has used the stabilities conferred by the office to say that they are. The performance of the ritual has meaning insofar as it influence the abilities of both the performers and the observers to control some perceptions.

Think again about my message that told you of the abdication of King Isa Mei Salahud of Zingaria in favour of his son Zeo Mei Fanhuil, which I presumed you would say was meaningless to you. If I now tell you that Zeo Mei Fanhul is aggressive, that, too, may be meaningless to you, but if I add that I want you to go to neighbouring Fahlistan to retrieve something before the new king tries to invade, my earlier meaningless information now means much more. It influences not only what you might control for in support of your control for perceiving yourself to please me, but how you would do it. By incorporating possible action related to the information, the meaningless "information" has become meaningful.

"Meaningful" differs from "meaning". How? Above, it is argued that an observation or a message has meaning to you — is meaningful — insofar as it affects at least one of your purposes. I suggest that "meaningfulness" is a weighted measure of how many purposes it apparently influences, and how strongly. Your life is meaningful to the extent that it influences for better or worse the perceptions controlled and controllable by yourself and by other people (and non-human living things).

"Information" is a quantifiable measure of possible "meaning", while "meaningfulness" is a potentially quantifiable measure of the range of purposes those meanings affect. The more strongly an observation influences the way you are able to satisfy an intention, the more meaning it has for control of that perception, whereas the more intentions the observation influences, the more meaningful it is.

In quick summary, the argument I am proposing is that the meaning of "meaning" is entirely in how an observation or message influences your perceived ability to control some of your perceptions, and is not inherent in the observation or message. The meaning to the sender of a message is in the perceptions the sender, and only the sender, controls differently when the recipient receives it. That meaning may be very different from the meaning of the message to the recipient, even when both parties agree that the message has been correctly transmitted and understood.

To understand how a message can be transmitted to the satisfaction of the sender, we must look a bit more closely at the operation of protocols, and at the detection and correction of errors in the message transmission, as a process of perceptual control of key variables that are common to all protocol interactions. These expand the two control loops of the protocol to include a few more of the fully elaborated motif of 19 control loops.

II.14.9 Protocol Failure I: Error Correction

Failure to execute a protocol "properly" might have any of several reasons, one of which is that one partner misperceives the other's state. Sometimes we call this "talking at cross-purposes. Such misperceptions can be corrected, but only if they are perceived as happening. So, in a protocol between Ilya and Cassie, not only must Ilya (the initiator) perceive Cassie's actions, which he hopes will appropriately influence his perception that was in error, but in many cases he must somehow perceive Cassie's perception of what he wants.

In Ilya, the error that leads to his own action can be induced either by an environmental disturbance to the CEV of some controlled perception or by a change in the perceptual reference value, whereas in Cassie, the continuing partner, the error is necessarily produced by a disturbance caused by Ilya's actions. If she wants to cooperate (and often if she does't) Cassie must be able to perceive not only that Ilya wants her to produce some kind of action to correct a perception he controls, but also whether Ilya perceives her to be doing the right thing.

A protocol "works" if it serves to reduce the error in a perception controlled by the initiator, and by virtue of that effect, reduces error induced in a perception that was disturbed by the initiator and is controlled by the continuer. Protocols usually do not achieve their final result immediately following the initiating action, though sometimes they do, as in the classic "*Close the window*" example. That does not mean the protocol has failed, any more than control in a tracking task has failed if the error induced by a disturbance remains uncorrected after only one loop delay.

In a tracking task, control takes its time. Likewise, in a protocol loop, it is quite normal for the errors in the two controlled perceptions to be corrected over a time that lasts several loop delays around the "Edit-Accept" loop of the GPG. As with a tracking task, both legs of the Edit-Accept loop can and at higher protocol levels do function simultaneously and continuously. There is no more need for turn-taking in the operation of a protocol than for stopping to check the effect of each individual output action in a simple tracking control loop. Error correction on both sides can usually proceed simultaneously and continuously.

Cassie, of course, is likely to have several ways of countering Ilya's disturbance, depending not only on what perception(s) Ilya's action disturbs, but also on the menu of means she has to correct any error he has induced. We may, for example, assume that she is controlling a perception of Ilya's state with a reference value that leads her to act to reduce Ilya's error and thereby her own. Something of the kind is almost essential for her to use the protocol to Ilya's benefit, unless Ilya has surreptitiously contrived a situation in which she cannot control the perception he disturbed without her action benefiting him. If she is controlling a perception of Ilya's pleasure, she has to be able to perceive what he wants, and she will use the protocol that she perceives Ilya to have initiated. Ilya's "display" allows Cassie to discriminate what Ilya wants, among the many possible protocols that they both know.

However, even if Cassie is controlling her perception of Ilya's state with a reference to perceive Ilya to be annoyed or uncomfortable, she might use the same protocol, but instead of using it to reduce the error in Ilya's controlled perception, she might intentionally act so as to increase his error. Ilya's action to continue the protocol might then actually serve to reduce Cassie's error, as she perceives him to be getting more and more annoyed. This use of the protocol loop might, but does not necessarily, lead to an escalating conflict, because Cassie's error is being reduced while Ilya's is being increased. But even an innocent misunderstanding can sometimes result in an escalating conflict.

A misunderstanding persists when Cassie is controlling for perceiving Ilya to be content so long as Cassie acts to control the perception Ilya disturbs in a way that increases Ilya's error. This could happen in a variety of ways, even when both parties have the best of intentions.

1. Ilya's actions may disturb a perception in Cassie that is not the one that "should" have been disturbed in the protocol he intended to use. In other words, he has a false perception of Cassie's current state.
2. Ilya may have disturbed the perception in Cassie that he intended to disturb, but he may have misperceived the current value of that perception in Cassie.
3. Ilya may have disturbed the intended perception in Cassie, but his actions may also have disturbed another perception Cassie controls, so Cassie's control actions are not as Ilya modelled. In other words, Ilya initiated or performed the protocol improperly.
4. Cassie may have misperceived what perception Ilya was trying to control through her actions, and acted to influence a different perception that Ilya also controls. Ilya's actions to counter this disturbance might disturb another perception in Cassie, leading to a loop with positive gain through the two partners, and an escalating conflict.
5. Cassie may have correctly perceived what perception Ilya was trying to control, but not Ilya's reference value for that perception.
6. Cassie might simply have executed her side of the protocol improperly.
7. Ilya's and Cassie's actions may have sufficient timing delays to make the loop gain turn positive. We will return to this problem when we deal with the differences between formal broadcast and informal interactive use of language in Chapter III.1.

***II.13.10 Protocol Failure II: Learning to Fix It

When can we consider that a protocol has failed? Ilya, the initiator, will consider it to have failed if Cassie, the continuer, does not act to influence the perception he is trying to control in the direction that reduces the error. The reason might be that Cassie does not at the moment want to do what Ilya needs, but it might be that Cassie does not perceive what Ilya wants her to do. For her part, Cassie will consider it failed if her actions do not reduce the error in the perception that Ilya disturbed in her.

Whatever the reason, when the error in a controlled perception remains large or continues to increase, the rate of reorganization increases — people try other actions. We discussed this process at length in the examples of baby Ivan and mother Cora, and of Len, Sophie and Dan's development of a language of syncons.

The stereotypical "misunderstood English tourist" in a foreign country is reputed to keep saying the same thing ever louder, as though the "foreign" person simply cannot hear properly but would understand if the words were pronounced clearly and loudly. An actual tourist is more likely to try to find some words in the foreign language, or to try to find another person who could translate or who could serve as a protocol partner or intermediary. The old adage "If at first you don't succeed, try, try, try again" sometimes works (it is, after all, the process of error reduction in control), but is often replaced by "If at last you don't succeed, try it a different way" (which is the basic principle of reorganization).

Chapter II.15. Protocol Implementation Issues

II.15.1 A Hierarchy of Protocols

Protocols exist in a hierarchy, as do perceptions in the individual control hierarchy. A complex protocol may be implemented by passing through many levels of simpler protocols before the interactions actually are perceived through the shared environment. A "General Protocol Grammar" (Taylor, Farrell and Hollands, 1999; MMT in LCS IV) provides the kind of linkages among levels of a protocol hierarchy in the way that perceptual and output-to-reference signals provide linkages among the levels of a simple control hierarchy.

In the world of Len and Sophie's family, the lowest controlled perceptual level is that of features, above which are syncons and trajectories. Above these are the lowest level protocols such as acceptance and rejection, as well as a variety of simple protocols that can often be completed without overt dialogue sequences, such as "Please" and "Thank you", "Open the X", "Pass the Y" and similar, all of which can be used as components of more complex protocols. They are implemented through trajectories among the syncons emitted by one partner and perceived by the other.

A protocol can be learned or developed in at least three ways. In Chapter II.11.1, Cora and Ivan reorganized together so that Ivan would tend to do whatever Cora had latched onto as his signal for a particular need, and Cora developed the perceptual function that would allow her to identify Ivan's need and thus to control most efficiently her own disturbed perception of his well-being. Co-reorganization is probably the only way novel protocols are developed, as we saw when Sophie unknowingly divided the pre-existing "BA" for a roundish object into "BA" for a spherical object and the previously unused "VA" for a less symmetric smooth roundish object (Section II.13.4).

A second way to learn a protocol, which applies to a protocol that is collectively controlled within the culture, is for a naïve person to observe the protocol in action between mature users. Learning by watching seldom leads to good control, but it can lead to sufficient control that practice can make perfect by further reorganization.

One of the mantras of PCT would, however, seem to pose a problem: "*You can't tell what a person is doing by watching what they are doing*". Is a person seen knocking on a door testing the acoustic qualities of the door, feeling how knocking affects painful knuckles, selling cosmetics, visiting a friend who lives there, deceiving someone else by making them else think he is interested in that house when he is actually watching another, or what?

The mantra does not apply so strongly when one observes a protocol being executed, because if they are not trying to hide some private matter, the players each use displays intended to make it easy for the other to see something of their internal state, probably including what perceptions they are trying to control at that time. Provided the observer has the necessary perceptual functions and the performers are not deliberately hiding their purposes from an observer, the nature of the protocol usually makes clear what the initiator is doing.

This is especially true when the actions of the continuer work to bring the initiator's controlled perception near its reference value. A child can learn by watching that, for example, when someone not in the family comes to the door, the visitor and the person opening the door are likely to say something like "Hi" or "Good Morning" to each other (if that is appropriate in the child's culture). Later, if the child is in that situation, she may also say "Good Morning".

The third way to learn a protocol is by being corrected when misusing it. The learner is exposed to the initiating display and produces an action different from what the initiator or a bystander deems appropriate, or the learner wants to achieve a result for which there is a protocol and uses an inappropriate initiating display. For example, if the child answering the door says “Hi” to a dignified guest, she will probably be told “That’s the way you may greet your friends, but not grown-ups”. A child might say “Gimme ball”, to which mother might refuse the ball, but instead say “Say ‘Please’”. Or the child might be asked by a brother “Please could you give me the ball” and not respond, following which mother might say “He asked you nicely, so you should give it to him or at least tell him why not”.

The form of a protocol may be controlled by two individuals, but most of the protocol forms in a culture are between two roles such as between a buyer and a seller, or a doctor and a patient. These “cultural” protocols are stochastically collectively controlled (Chapter xII.15). In any culture more complex than a nomadic hunting or herding groups of a few families, there are a lot of roles. If the culture allows for N roles, there are potentially N^2 interaction types between them, each of which uses its own collectively controlled format. This number is unrealistically large. A person while playing a cashier role does not normally interact with a person who is at that moment playing the role of dentist or farmer, for example. Usually, the cashier role interacts only with the customer role.

The situation is numerically similar to the autocatalytic situation in which the basic elements are chemical elements, only some pairs of which easily react together to form molecular complexes. In Chapter II.15 we will talk more about protocols and roles, and the way that “necessity is the mother of invention” in the way that the interactions implemented by protocols can catalyze other interactions to form a different kind of autocatalytic, inventive, loop. For now, however, we stay with single levels of single protocols and their properties.

II.15.2 Protocol loop dynamical considerations

A protocol loop is a negative feedback loop, somewhat more complex than the simple loop in which the controlled perception is of some property of an inanimate object, but a negative feedback loop nonetheless. As such, it is subject to all the dynamical constraints to which a simple control loop is subject, but it has additional constraints that need to be addressed, because they go to the heart of language and other cultural conventions.

Returning to the world of human conversation, we start with some approximate timing facts about spoken language, starting with the fastest component, the acoustic waveform, and ending with the apparently successful completion of a protocol, which might take as long as the communication of a complex message over a period of months or years, or as short as the performance of a simple action in seconds. We will assume that the initiator wanted to produce a particular level of understanding in the continuer, for which the shorthand is “passing a message”. By “apparently successful” we mean that both parties perceive that the listener perceives whatever the talker intended should be perceived, such as that dinner will be at 6 0’clock, not 6:30. Correctness at this “message intention” level is assured by feedback:

Robert: Did you say dinner was at 6 or 6:30?

Terence: I said 6, but you can come earlier if you want.

Robert: I didn’t hear you properly. You mean on the hour, not the half hour?

Terence: That’s right.

Robert: OK, I’ll see you a bit before 6.

At the lowest levels, feedback cannot be used. Phonemes go by at maybe 8 per second, much too fast for feedback to check every one of them, and even words pass at something like 2 per second. To say “OK” or “Missed that” for every word would be impossible. Even if the listener were capable of doing it, the mere act would interfere with hearing the next word, and the delay would be enough that the talker would probably not be able to perceive which words referenced by "OK" and which by "Missed that". At these levels, the internal structure, which we generically call "syntax", creates redundancy that the listener can use for error correction, as in the discussion of "syncon trajectories". At these lower levels, the protocol shades into pure perception.

One technical problem is that Shannon (1949) proved that perfect error correction requires infinite time even with perfectly matched encoder and decoder pairs, so it is highly probable that some errors (misunderstandings) would remain even at levels slower than words. These slower changing macrostates are amenable to feedback correction in face-to-face communication, but not in written text (perhaps apart from "texting"). Written text still requires syntax (structural redundancy) for error correction, even at the highest level. Table II.15.1 and Figure II.15.1 offer an informal list of the preferred methods of error correction (syntactic redundancy or feedback) for speech segments of different time-scales in face-to-face interaction.

Table II.15.1 Academic discipline, approximate timings and method of introducing redundancy to reduce error rate for different segmentations of speech in a face-to-face interaction (Based on Figure II.2-10 of Taylor, 1989)

Discipline	Abstraction	Approx. Sample Rate	Redundancy Method
Psychoacoustics	Acoustic	II.6,000/sec	Harmonics and formants
Speech Recognition	Phonetic Segment	II.6/sec	Syntax
	Phoneme	II.13/sec	
	Syllable	II.4/sec	
	Word	II.2-3/sec	
Psycholinguistics	Phrase	II.1 second	Mainly syntax
	Sentence	II.4 sec	Mainly feedback
Dialogue Analysis	Speech Act	II.4 sec	Feedback
	Topic	several seconds	
	•	•	
	Goal	Possibly days or years	

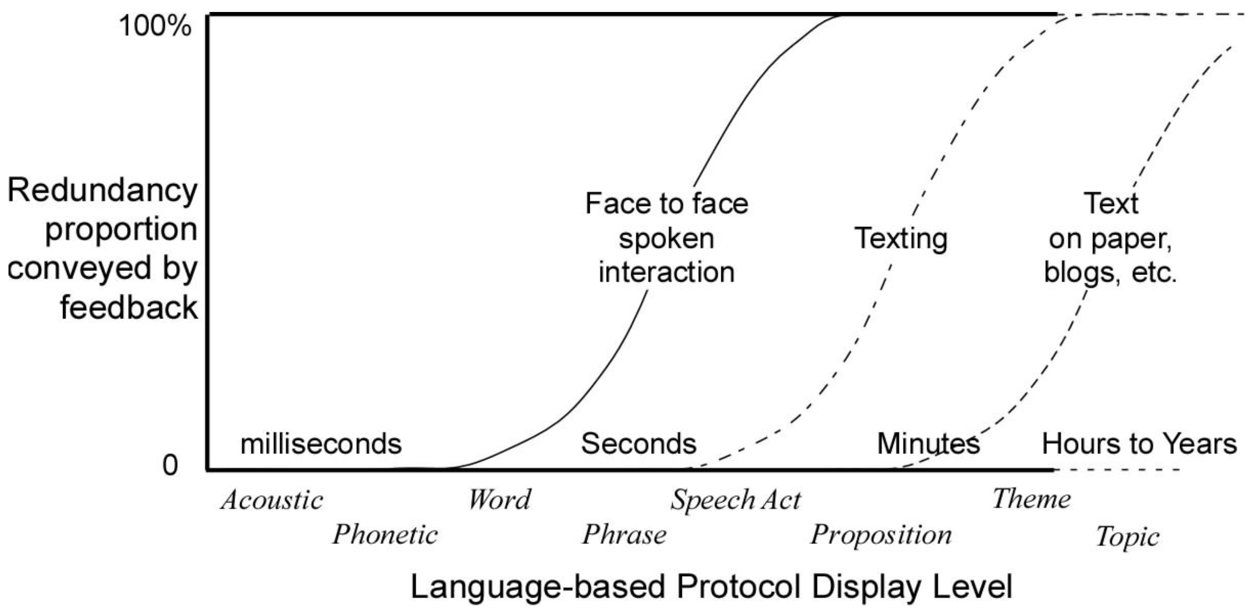

Figure II.15.1 A suggestion of the proportion of error-correction performed by feedback at different time-scales and characteristic protocol levels for communication with a partner familiar with the speaker/writer's pattern of language use in the particular method of communication. The remainder is corrected by structural redundancy (syntax).

When we were dealing with simple control loops, we discussed the effect of loop transport lag on the stability of the loop. If the loop gain is too high at a frequency where the loop delay returns the signal in phase with its outgoing version, the loop can become unstable, with the signal amplitude increasing exponentially until some component of the loop saturates and limits the amplitude.

Now we are dealing with a more complicated loop, in which the signals are not simple waveforms but are syllables, phrases, and topics. Nevertheless, the same considerations apply. If the loop gain is too high (the partners try too hard to control their protocol perceptions with little tolerance for error) and the loop delay is long enough, their joint control efforts will result in the overall loop feedback becoming positive, with results more like those of an ordinary conflict than of the cooperation implicit in the use of a protocol. Communication fails.

We come to some competing pressures that do not exist for the simple control loop, but are important for communication using language. The main difference between them is that the operation of a protocol loop depends on the initiator's disturbing a perception in the continuer. The slower the change in the disturbance, the easier it is for the continuer to control. A slowly changing disturbance is one that has a low information rate. But this particular disturbance is in a section of the larger protocol loop, and stable loops need fast transport. So there are competing pressures for the speech transmission from the initiator to the continuer to be slow and for it to be fast.

Fast communication is possible, at the cost of increasing error rate. But the initiator wants the continuer to know what is being asked, at least accurately enough for the continuer to produce the appropriate action perceptible to the initiator. In other words, the initiator would like to control his perception of the continuer's understanding. Here we have a separate pair of competing pressures on loop stability, as high gain demands low loop transport lag, but low error requires long delay (Shannon, 1949).

Generally speaking, more speed allows for higher gain without loss of stability, but induces more error at the levels where error correction is provided by syntax. Errors that propagate upward to the levels where they can be corrected by feedback cause delays at those levels, so co-reorganization needs to find a balance between speed and accuracy across the levels when communicating between participants familiar with each other. As we shall see later, a "familiar participant" might be an arbitrary individual acting in a familiar cultural role.

Self-organizing processes such as reorganization often produce results at or near a critical boundary. We first saw this in connection with "The Bomb in the Hierarchy" (Section 6.4), which can develop even in a hierarchy that controls perceptions only of non-living things. The situation here is very similar. Speed helps loop stability, but if the control is poor because of transmission errors, stability is compromised. The evolution of the language, including culturally equivalent gestures that increase the available channel capacity by using a visual channel in parallel with the auditory, will necessarily work to reduce the transmission of redundant material. On the other hand if communication fails because of insufficient redundancy, the tendency will be for redundancy to increase. The two opposing tendencies lead language, including the accompanying gestures, to converge to the critical boundary.

Combining these tendencies leads to the observed pattern of what Nevin describes in his Chapter of LCS IV as "reductions", the apparent omission of speech elements likely to be automatically replaced by the listener. The elisions in the frequent American pronunciation of "President" as "Pres'n" or "Pre'n" offer an example. Such elisions are correctible by feedback (which is speeded by their omission) whenever their absence leaves the initiator's intent uncertain to the listener. The same pattern was seen in the syncon trajectories used by Len and Sophie (e.g. Figure II.13.17 or Figure II.13.18), which can allow a syncon to be perceived even if the trajectory actually bypasses it by a substantial margin on the way between its predecessor and successor.

In Figure II.15.1, some different levels of language structure were shown, each of which might constitute a level of a protocol, the faster and more "syntactic" supporting the slower and more "feedbackish" and vice-versa. The analogy between the layers of protocols and the levels of the control hierarchy should be clear, but one specific aspect of the analogy might be less obvious. In both the HPCT hierarchy and the protocol layers, the lower levels *stabilize* the inputs to the higher levels, reducing the likelihood of high-level error and allowing for slower and more powerful correction processes without inducing loop instability. As we have seen, this can be overdone, as too much stress on low-level precision can slow the higher loops, leading to instability despite the increased precision of low-level control.

Reorganization allows the total system to be fast, accurate, and with a high loop gain, but the resulting system is on the edge of instability. Such systems are often the most effective in a variable world. For example, a bicycle and a fast fighter jet are designed to be on the edge of stability so that the rider or pilot can change course very quickly if need be. Such systems, whether mechanical, evolved, or self-organized, are liable to errors, small, frequent and easily corrected as well as larger and hard to correct, or even very rare but catastrophic.

The possibility of error avalanches is the price paid for *toughness*, a combination of flexibility and strength, as we saw in the case of the physical tensegrity structures that we proposed as analogues to human individual control structures, and, we shall argue, to social structures as well.

II.15.3 Multiplexing and Diviplexing

When discussing protocols, we usually use the word "message" in place of the cumbersome "action to influence a perception in the partner". They mean essentially the same. Protocols both facilitate and use communication, the process of producing a desired perception in another organism. If the originator of a communication successfully brings the recipient's perceptual profile to a particular macrostate, we say that a "message" as been correctly transmitted.

Teri wants a particular blue box to be on a table and says or gestures something to Hal, who is near several boxes. When Hal places the blue box on the table, Teri's message has been both communicated and acted upon. We know this from our Analyst viewpoint, but Teri never knows whether Hal independently decided to control his perception of the relationship between the box and the table, having not even noticed Teri's existence. We look more closely at their interaction later in this Section.

The reference value of the perception that the initiator of a protocol wants to control by using the protocol is given a special name: the "Primal Message". That is what the initiator wants the continuer to understand. Within the protocol, the initiating communicative action is called the "Primary Message". If the Primal Message is very simple, the Primary Message may constitute the entire protocol, but usually there are other messages, questioning, confirming, editing, and so forth, until the initiator and continuer both believe that the Primal Message has been correctly received — the perception in the continuer has been influenced as the initiator intended, whether or not the continuer's action has actually reduced the error in the perception the initiator was originally trying to control through use of the protocol.

In telecommunication, two or more independent messages are often multiplexed onto the same channel and split apart by the receiver. The same can happen in human face-to-face communication. For example, the word "Yes" spoken with different intonations can convey acceptance simultaneously with a mood such as joy, surprise, or doubt (e.g. Allwood, 2000 for the Swedish version).

The word "diviplexing" is, on the other hand, not common. It was introduced by Taylor and Waugh (2000) to describe the situation in which one message is sent by two parallel channels, as suggested above for verbal and gestural language. Figure II.15.2 illustrates both multiplexing and diviplexing. At a very low level, every phoneme is diviplexed across many features. In the syncon world, the speakers diviplex across three features (and later, more, when trajectory types begin to act as independent features).

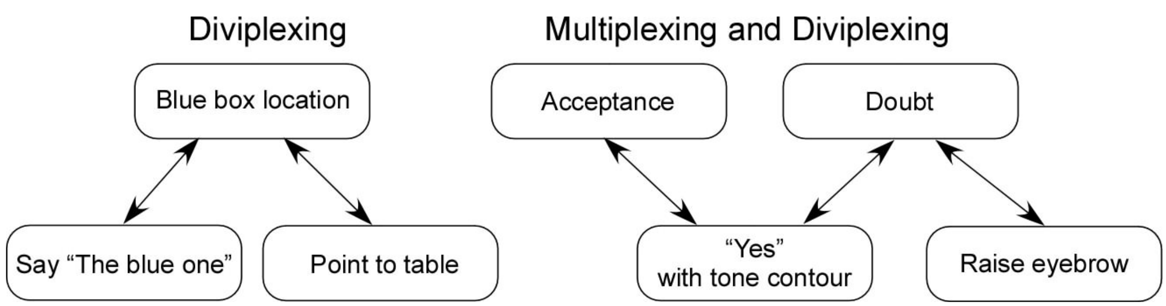

Figure II.15.2 Diviplexing and Multiplexing. Diviplexing both speeds and adds an orthogonal kind of redundancy to the communication, whereas Multiplexing uses an otherwise available part of the communication channel to convey a separate message. Both save time, which is important for the stability of a protocol loop.

When you speak, your voice always has a tonal pattern, whether or not that pattern is used to communicate a message. In some languages, as for example Chinese, the tone contour conveys as much of the meaning as does the sequence of phonemes that constitute the syllable. The syllable "ma" in the four different tones of Mandarin can mean "*mother*", "*hemp*", "*horse*" or "*to scold*", no two of which are even in the same semantic domain. In most languages, however, the tone contour does not change the meaning of a sequence of phonemes, and is therefore available to convey an independent message, often a mood such as joy, doubt, or sadness, or to change the meaning from a statement to a question.

Diviplexing is illustrated both by the message of "Doubt" in the right-hand part of Figure II.15.2 and by a classic example of a protocol initiation that was publicized by Bolt (1980) as an example of something difficult to implement in a computer interface: "Put that there", in which the initiator is controlling for seeing a particular object in a particular place, but does not specify in words either the object or the place[68].

For example, if Hal enters a room and looks expectantly at Teri, Teri might simply wave a hand toward an empty table and Hal would then put a blue box onto the table. No words were spoken, and none were omitted, but each of them brought their controlled perceptions nearer their reference values through the action of the other, because of what each perceived the other to be perceiving of the situation. Teri might, for example, have previously called the maintenance office to ask for someone to come and move the blue box, and perceived Hal to be that person when he showed up at the door. Hal displayed to Teri that he did not know where she wanted the box, and she gave him a display that allowed him to correct the error in his perception of box location.

Here are a few more implementations of this same protocol, in which the overt actions change according to what each can perceive of the other's current World Model in slightly changed circumstances. In all cases, Teri wants to perceive a blue box to be on the table and Hal wants to perceive what Teri wants done so that he can do it. In examples 1 through 4 Teri perceives Hal to have the information that is not explicit in what she says, whereas in examples 5 through 8 the complete message is diviplexed between voice and gesture. Figure II.15.2 (left panel) illustrates yet another possibility. These examples of what Teri might say or do to get Hal to put the blue box on the table are quoted from Taylor and Waugh (2000), only changing the original "talker" and "hearer" to Teri and Hal respectively. All of the italicized quotes are from Teri.

II.1. *The blue one*

Teri believes that Hal believes a box is to be placed on the table, but does not know which of several boxes to move.

II.2. *On the table*

Teri believes that Hal knows what is to be moved, and that it must be placed somewhere, but does not know where. Perhaps Hal has just entered the room, carrying the blue box, and is looking for a place to put it down.

II.4. *Put it on the table*

Teri believes that Hal knows something is to be done with the blue box, but not what is to be done or where to do it.

II.5. *Put it down*

68. McCann, Taylor, and Tuori (1988) used the inspiration of Bolt's example to demonstrate a working interface in which voice, gesture and keyboard were used in a freely interchangeable combination to query a miniature military situational database. For example the operator might say "*Show me these*"while pointing at a symbol such as a battalion, and continuing "*such that strength*", while keying "> *80%*" to get a map display of battalions that were up to at least 80% of full strength. This project was the inspiration for the Layered Protocol Theory of communication (Taylor, 1988a, 1988b) which led me to PCT.

Teri believes that Hal knows what is being talked about, and if given the action to carry out, will perform it in the proper location.

II.8. *There* (pointing)

The situation is equivalent to that of number 2, but Teri chooses to gesture toward the table rather than to identify it verbally. The word “There” signals that the required information is being provided in a visual mode.

II.12. *Put the blue box there* (pointing)

Teri believes Hal not to know what to do, which object to do it with, or where to do it. The object and the action are conveyed verbally, and “there” signifies that the location information is to be found in the visual supporting message.

II.13. *Put that there* (point to the blue box and then to the table)

The “classic” example, from Bolt (1980). Both “that” and “there” signify information to be obtained elsewhere, presumably by way of the visual gestural message.

II.12. *Put it there* (point)

Both anaphora and gesture are used to complete the information required. Teri believes that Hal has the blue box in focus[69], perhaps because it was recently mentioned, perhaps because both parties are looking at the box and each knows the other is doing so. Anaphoric reference need not be to previous dialogue items, but they must be to items Teri believes Hal to have in focus.

And finally:

II.1. *Teri waves at the table*

This is the interaction mentioned at the start of this Section. Teri perceives that Hal is carrying the blue box and is looking for the right place to put it.

All these possibilities implement exactly the same protocol, whether Teri’s observable actions use explicit language or no language at all. When language is used, it may be used in partial sentences or even isolated words. All depends on what Teri perceives to be likely to disturb Hal’s controlled perceptions in such a way that the blue box appears on the table.

The examples assume that Hal will understand Teri’s intent and will act to correct what he perceives to be the error in her controlled perception of the box location by putting the box on the table, but he may not always understand correctly. Suppose, for example, Teri uses implementation 3 above: “Put it on the table”, incorrectly perceiving that Hal will perceive “it” to be the blue box. If he does not, either Hal will ask for clarification, using feedback to reduce his uncertainty if he cannot guess what Teri wants on the table, or Teri will use a corrective display if Hal puts a red ball on the table. Either way, because they are face-to-face, they are likely not to use full sentences, but instead to use forms such as: Hal “What on the table?”, or Teri “The blue box”.

69. . “In focus” is a loose term that implies that the blue box is a component of some perception Hal is controlling that is not currently at its reference value.

II.15.4 Protocol Identification

Every protocol exists because there are occasions when a person wants someone else to perform the actions that reduce error in a controlled perception, even if the action is only an eventual display of comprehension of something the initiator wants to perceive the continuer to understand. Reorganization within an individual's control hierarchy creates the possibility of controlling myriads of different perceptions, but because of the limitations on the individual's means of acting on the environment, reorganization also develops a structure that minimizes (but does not eliminate) conflict among all those control units. The initiating and continuing halves of protocols are among the control systems developed by reorganization, and as with the rest of the structure, reorganization tends toward reducing conflict.

Within an individual, conflict can occur only among the control systems inside that individual, but when we consider the co-reorganization of pairs of individuals, the opportunities for conflict increase dramatically, since twice as many control systems exist in two individuals as in one. The two reorganization processes that create effective protocols must operate in such a way that the level of conflict involved in using protocols is not appreciably worse, and with luck is better, than the level of conflict that remains within the individuals.

When a newcomer such as Dan tries to use an established protocol, it is the newcomer who will do most of the reorganization, because the existing users of the protocol, Len and Sophie, live in a world in which much of the conflict-reducing co-reorganization has already occurred. The protocol exists in a world of protocols that work together reasonably well — a "culture" — and it is unlikely that Dan can introduce a new one or greatly change an existing one without inducing some conflict in the process. Len and Sophie's existing protocol structure may provide yet another example of an analogue to a tensegrity structure that is resilient against moderate stresses.

The difficulty of co-reorganizing relatively conflict-free protocols is diminished by the fact that only a small proportion of all the perceptions controlled in an individual are protocol-related, but it is enhanced by the fact that every protocol requires that the participant display effectively, so that each partner correctly perceives what the other wants them to perceive of their internal state. This requirement puts limits on the number of distinguishable protocol components, though as we discussed above, the more features enter into the identification of a perceptual category, the larger is the space into which discriminable categories can fit. Protocol displays thus include features such as the many dimensions of facial expression, intonation contours, and body language as well as words and word sequences.

Even though the initiator might use many different features to help the continuer to perceive that he needs her actions to influence one of his perceptions, nevertheless there are a lot of different things he might want her to do. Just as they both will have developed trajectory recognizers, so they should be expected to develop protocol recognizers.

Only a small proportion of the possible feature patterns define actual syncon locations, and relatively few of the possible syncon sequences define trajectories (words and phrases) that are actually used. Likewise, few trajectory patterns across the protocol partners specify actual protocols that the partners can perceive as entities. However, in the same way that different words and phrases can be used to specify the same thing in the environment, so also can different sets of trajectories define the same protocol, as in the natural language example of the blue box, which uses words and gestures interchangeably. The macrostate that defines one protocol contains many microstates that describe different ways to implement it.

II.15.5 Protocol as the syntax of interaction

We have tended to treat protocols as though they were fixed and immutable. They are not. Like the syntax of a language, they provide a framework within which the variable elements are readily interpreted. Consider the common English word order "Subject-Verb-Object". If we scramble the order, we have to rely on the pragmatic or semantic properties of entities labelled by the words. "Cow grass eat" is easily interpreted because we know cows eat grass, though that interpretation might be wrong. The speaker, for whom English is a foreign language, might be commenting on the height of the grass, which is sufficient to hide ("eat") a cow. A listener observing the same field of very tall grass would probably interpret the phrase as intended, but someone hearing it on a phone might not.

One of the mantras of PCT is "many means to the same end". But this does not imply that all means will lead to the same end. Although the low-level details may be highly variable, there are only a few ways of acquiring an ice-cream cone, and prominent among them is to offer an ice-cream seller some money to exchange for the cone. This is the initiating Primary Message of a "buying" or "commercial transaction" protocol. The same "buying" and "selling" components of the "commercial transaction" protocol are used for acquiring things other than ice-cream, though they may be implemented very differently if the object bought and sold is a house, a chocolate bar, a million shares of some company, or an on-line movie.

One aspect of the "selling" component of the "commercial transaction" is that the buyer who initiated the transaction must have been able to perceive that the seller is likely to have available the object of the transaction. You do not go around randomly asking people on the street whether they would sell you an ice-cream cone. You see a truck or a shop that looks like a place where ice-cream might be sold.

Engel and Haakma (1993) called a display of openness to an initiating protocol "E-feedback" for "Expectation feedback", though it is not truly an example of feedback any more than seeing a heavy plank by the side of a stream is feedback when one is looking for a way to cross the stream dry-shod. Their examples include knobs that display by their shape that they can be rotated, as opposed to switches that display by their shape that they can be flipped. As with the plank, E-feedback is rather a display of readiness to be used in the control of different perceptions — in other words, a display of at least one atenfel. We use the term "Readiness display" ("R-display").

R-display may broadcast to all and sundry, as does the advertising display on the ice-cream truck or a switch on a wall. On the other hand, the initiating and continuing components of a protocol are directed mainly at the partner, as for example are the initiating offer of money in one direction and its continuation by transferring ice-cream in the other. Of course, actions can simultaneously control more than one perception ("killing two birds with one stone" or multiplexing messages), and this applies equally to protocols, as may happen, for example, if Andrew deliberately ensures that a third party, Tom, hears a protocol transaction between Andrew and Beth. In such a case, Andrew separately controls a perception of a state in Tom by the way he implements the protocol with Beth.

When analyzing language, we can look at the detail of the words actually used, or we can take a more generic view, looking only at common forms. For example, we can judge that "The cow ate grass" makes sense, but "the cow ate glass" does not (in most circumstances), though there is only one phoneme difference between the two sentences. However, if "the cow ate …" had been preceded by a veterinarian having asked "Why is the cow sick?", suddenly the second sentence makes more sense than the first. Such contrasts illustrate three different levels of analysis: Syntax, semantics, and pragmatics.

Syntactically, both "the cow ate grass" and "the cow ate glass" are of a form that could be labelled "actor-act-object"; Semantically, they differ: "mammal ingestion digestible/indigestible". If the cow is sick because it ate glass, the semantic incorrectness is accommodated by the pragmatic context, because although cows do not normally eat glass, this one actually did and might die as a result. In a similar way, semantic information can accommodate a violation of syntactic structure, as might happen if a non-English speaker might say “cow grass ate”, which would be understood correctly by most English speakers.

These three levels of analysis (syntactic, semantic, and pragmatic) are a long way from the refinement described by Nevin in his chapter in LCS IV, but they will suffice as a parallel for a similarly crude analysis of protocol patterns. Syntax provides a frame for a single protocol within which content must be provided, semantics describes the normal relationships among content items within that protocol, whereas pragmatics (and only pragmatics) takes account of the surrounding circumstances of the moment, including protocols at supporting and supported levels.

II.15.6 Protocol versus Ritual

The intended effects of an ordinary protocol action by one partner are on perceptions controlled by the other. A ritual is different. We use the word to describe some habitual sequence of actions by one or more people. The significant fact of such a ritual is the behaviour that could be perceived by non-participants, not the effects on the actions or world models of the participants. There is a level of perceptual control at which the reference pattern cannot be altered without making this an ill-performed ritual, or perhaps turning it into a quite different ritual.

Making a regular visit to leave flowers on the grave of a departed loved-one can be a ritual. The actual movements involved will differ, but the act of laying flowers and then making some gesture of remembrance will not. Nobody else’s perceptions are deliberately influenced by these actions, but a chance observer would probably understand that the actor was performing a remembrance ritual.

We also use the word “ritual” to describe a public performance in which the participants are controlling for a change in other people’s perception of the status or properties of at least one of them. The important point of such a public ritual is that it be observed. Swearing someone into office would be pointless if nobody knew it had happened. A religious ritual may, to a participant, seem like a protocol executed with a God as a partner, but to the observer it identifies the performer as intending to be perceived as belonging to a particular group of like-minded people.

Public rituals may have several or even many actors, but the important concept is “witnessing”. The actions of a public ritual are performed so that people may witness them and change something about their perceptions of one or more of the performers or of their own place in the culture (I think here of witch-burning as an example, which is intended to show the power of the Church over the witnesses).

In some cultures, a couple who stand before a clergyman who says a defined set of things, and each say “I do” will thereafter be perceived as having a new property, “married”, whereas before the clergyman says a particular formula such as “I now pronounce you…”, they are perceived as “unmarried”. Generally, public rituals lead witness members of the public to change their World Models. Those changes in World Models are the reference values for the perceptions controlled by the participants. Thereafter, witnesses might change the protocols they use when interacting with the performers, one might almost say “actors” in the theatrical sense, of the ritual.

Protocols can be rituals, though most are not. Buying and Selling is performed as a protocol, but the result is a public perception that ownership of something has changed. The giving of a receipt allows the purchaser to show to members of the public who did not witness the protocol performance proof of the new ownership, just as a marriage certificate shows that the couple have properly said “I do” at the appropriate moment. A normally public ritual that changes the physical environment in a commonly interpreted way can therefore be performed in private, as, for example, a secret wedding. A secret wedding does not lead other people to perceive the couple as being married, but the signatures on a piece of paper allow the couple to choose who will so perceive them.

There is an overlap between protocol and ritual, and between ritual and ceremony, but if we take the viewpoint of perceptual control, there is also a critical difference. In a ritual or ceremony but not in a protocol, the actions themselves are the important part. In a protocol, the perceptions that are controlled are perceptions of each partner by the other. In a public ritual such as a wedding or a swearing-in, the important perceptions that the performers control are their perceptions of the witnesses.

Some protocols involve more than two participants, as do many rituals and all ceremonies. If a protocol does have more than two performers, the controlled perceptions involve all of them, either serially or in parallel. An analysis of a triadic protocol is in (Taylor, LCS IV). Protocols involving more participants are quite possible, but for the purposes of this work we do not need to be concerned with such complex protocol forms.

II.15.7 Deceit and camouflage

"One may smile, and smile, and be a villain" (Shakespear, Hamlet)

The multimodal proliferation of display features that help a protocol partner to determine what the other wants not only serves to identify the intended protocol, but also serves as a deterrent to deceit. A protocol user who attempts to deceive must ensure that not too many of the displayed features that the partner will use to produce the desired perception differ much from the pattern evolved during honest use of the protocol.

It is hard to learn to act on stage or on camera well enough that the lines written by a playwright come out with timing and intonation patterns that lead the audience to feel that the fantasy situation is real. The skill of an actor is often said to be to "get into" the person portrayed, and feel what the fictional person would feel. In the language of this book, that skill is the skill to use protocols as the audience would use them.

An actor deceives, but the audience is complicit in the deception, whereas a con-man uses the same skill to deceive, but the audience does not have the perceived situational context of "fiction" to allow for a relaxed acceptance of slightly mis-used protocols.

What Ernest wants Connie to believe or perceive about the world is part of Ernest's control of his own perceptions, not Connie's. If Ernest is controlling for perceiving a state that damages Connie's ability to control, he may use protocols in the normal way, but with content that is likely to alter the way Connie will act in controlling her perceptions later. Such actions on Ernest's part may be called deceit, and be considered morally wrong, but they parallel what evolution has done in many parts of the biological world by developing mimicry and camouflage.

PCT does not offer moral judgment, any more than does any other theory of the natural world. If Ernest's deceit is a way that Ernest can reduce error in a controlled perception, then the Theory of Perceptual Control suggests that Ernest may well use deceit. It also suggests that if others perceive the deceit, their perception of Ernest will change, as will their use of protocols in interaction with him. This may create conflict between different control systems within Ernest, perhaps increasing the rate of reorganization in Ernest's control hierarchy so that he tends to use deceit less. If, however. Ernest's use of deceit has helped him control more perceptions than honest use of protocols would (in his perception) have done, he is unlikely to reorganize to act more honestly.

Conflict internal to Ernest might occur if he controls a self-image perception with a reference to perceive himself as honest. To control a perception by using deceit would then conflict with controlling his self-image perception. Which control system wins the conflict would depend on such things as whether Ernest has an easy atenfel for controlling the other perception without deceit, and whether his self-image controller is working with a wide or narrow tolerance zone; would it permit a "white lie?"

Of course, if what Connie displays is not the faithful representation of her World Model that Ernest perceives it to be, then Ernest's communicative actions are unlikely to have the effect he thinks they do. If she feigns understanding, saying "Yes, dear", because she wants to continue reading the newspaper rather than because she really does understand something, Isaac would not be able to know that his communication failed unless Connie allowed him to test it (using the TCV in the shape of different protocols that disturbed other perceptions Connie might be controlling).

Feigning understanding is a kind of deceit, but Ernest is unlikely to know about it immediately. Any problems it might cause become manifest only later, when Ernest relies on a misperception of Connie's state in initiating some other protocol, or when Connie does something she would have been unlikely to do had she understood. "But, Connie dear, I told you about that only a few minutes ago."

Deceit may be unintentional, but here we are concerned with the intentional kind. In our discussions of honest protocols, we assumed that at least the person continuing a protocol was controlling some perception of the other with a reference to perceive the other as satisfied with the effect of the protocol. An intentionally deceitful initiator or continuer will control for perceiving the victim to be similarly satisfied.

A deceitful Andre controls some perception relating to Beth, but what he shows Beth need not allow her to see what perception he is actually controlling. He can show a naïve Beth anything at all to disturb a perception Beth controls, and Beth's actions will allow Andre to show her whatever he wants until her actions are what he truly wants to see. At that point, he can allow her to see that his falsely presented wish has been satisfied. Beth may have been damaged in ways she does not immediately perceive. For example, her bank account may have been emptied, which she will discover only when she tries to withdraw some money.

For Andre's deceit to work, he has to use camouflage. As suggested several times in this book, many perceptions are based on association, otherwise known as context or redundancy. Even at very low levels, the same physical surface may be seen as quite different colours in different contexts, as Figure 9.1 suggests in grey-scale. Camouflage depends on context; the camouflaged object has to seem to be consistent with its background. Andre presents a background set of perceptions; he may "smile, and be a villain", the smile being what Beth would perceive (wrongly) as showing his goodwill toward her.

Of course, there are many other elements of the context that would lead Beth to perceive "trustworthy", and if Andre is a good actor, he would use them; Beth would not have any of the relevant perceptions disturbed except those Andre intends to disturb. But if Andre creates a pattern of actions that Beth can perceive not to fit properly with patterns she knows to correspond with "trustworthiness" (Figure II.15.3), she may not act as Andre wishes when he initiates the damaging protocol.

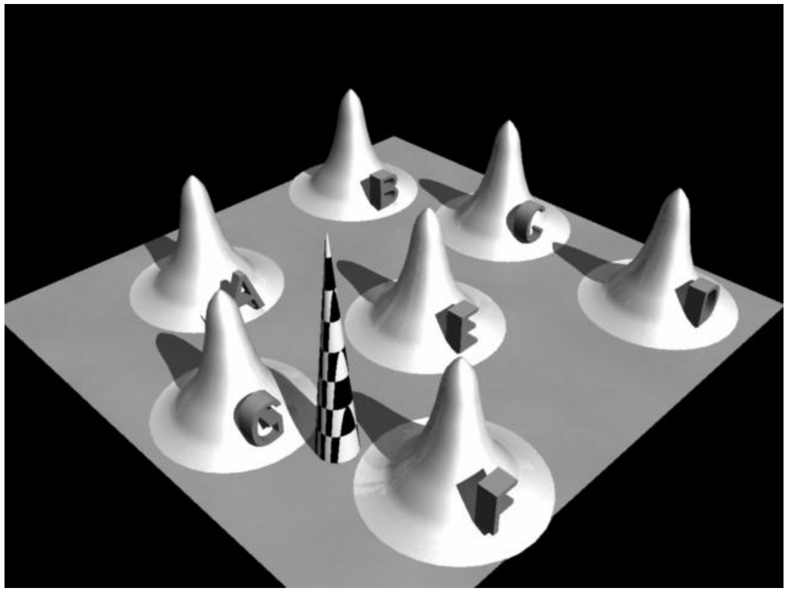

Figure II.15.3. The "Seven Syncon" patterns of Figure II.13.4 illustrate in 2-D a possible failure of Andre's camouflage. He (patterned cone) has mimicked one feature of patterns A, F and G, but not their other feature. Andre's deceit probably will be unmasked, but perhaps his pattern will be perceived as a variant of F or G, allowing Andre to continue his deceit.

We will delay further consideration of deceit and treachery until we start talking about politics, though it is a very important issue when dealing with the interactions among people who are in conflict because of a lack of environmental degrees of freedom. In that kind of conflict, two or more control systems try to set fewer environmental variables to distinct values, which means that not all the contesting control systems can reduce their errors to near zero.

More simply put: "*All's fair in love or war*" and war definitely involves deceit, as Sun Tzu made very clear two and a half millennia ago. But most of our protocol interactions are not intended to be deceitful. We trust the displays that contribute to our perceptions of their progress.

II.15.8 Counter-Control

When Sun Tzu advocated acting so that the enemy perceived the world as other than it was, he intended the enemy to act in a way advantageous to his own side. He did not actually control the enemy's actions, but the effect Sun Tzu wanted was as if he had been. Indeed, Sun Tzu could not reach into the enemy general's control structure and set reference values for certain perceptions, which would be actual control. He had to act through the environment, deceptively.

In the framework of a protocol, his deception would be the same as Andre's if the enemy general realized that Sun Tzu was "displaying" the situation. But that would not lead to success for Sun Tzu, because the enemy general would have perceived that Sun Tzu would try to mislead him. Instead, Sun Tzu would try to hide the fact that he was manipulating what the enemy general perceived, unless he was engaged in a "double bluff". If the enemy general perceived that Sun Tzu was trying to get him to act in a certain way, his most probable actions would be quite different, possibly the direct opposite. Of course, if Sun Tzu had been engaged in a double bluff, that would have been what Sun Tzu wanted, but the enemy general would not have known that to be the case.

What we just described was the likelihood of the enemy general performing what has been called "counter-control", the tendency for someone who perceives they are being controlled by another to do the opposite of what the other appears to be trying to get them to do. "Tendency" usually implies that a statistical test has shown that the "tendency" will show up more often than would be suggested by chance. From a PCT viewpoint, that is a matter of peripheral interest. PCT is a theory that suggests why such effects as counter-control might or might not occur in any specific case. The statistical "tendency" is a side-effect depending on conflicts within the individuals concerned, of whom there might be many or perhaps only one.

We have presumed that a high-level pair of perceptions controlled by many, perhaps most, people are self-self-image and other-self-image (the self as perceived by the actor, and the self as the actor perceives others to perceive him or her). One possible controllable property of either self-image is autonomy, the ability to choose freely what to do under all conditions. If a person is controlling for either self-self-image or other-self-image to have a level of autonomy in respect of some perception they control, and they perceive that some other person is controlling for them to act in a certain way when controlling that perception, then at least the self-self-image is likely to be disturbed, leading to counter-control action.

Perhaps more important is the other-self-image, because it forms part of the social relationship networks that we will discuss in the rest of this book, beginning in the next Chapter with the various forms of Collective Control, but becoming prominent in Section III.5.8, entitled "Quarantine, Nonconformity, Schism".

II.15.9 Arguments: Truth-Seeking and Political

Thus far, we have considered protocols as controlling the belief structures of two individuals, one of whom wants to convey a Primal Message to the other. Whether this Primal Message has any relation to what the initiator believes about Real Reality (RR) or about the relationships among the White Boxes in Perceptual Reality is irrelevant. What has mattered to this point is the degree to which the nine beliefs in each party about the three-level nestings of the three propositions have become sufficiently strongly held to be true.

To refresh, the three nested levels of belief are (1) the degree to which the initiator believes the proposition, (2) the degree to which the initiator believes the continuer believes the proposition, and (3) the degree to which the initiator believes the continuer believes the initiator to believe the proposition. The three propositions are, again:

- P1: *That the continuer has understood the Primal Message.*
- P2: *The quality of the communication process is sufficient for an adequate interpretation*.
- P3: *That it is not worth continuing to transmit this message.*

In an argument, these three propositions are insufficient to describe the protocol, because the two parties disagree about the truth of the Primal Message, not about whether the continuer understands what the initiator is trying to get across. The truth of the Primal Message is the degree to which it matches what each party believes about something outside the protocol, and in a Truth-Seeking argument those beliefs differ. The argument might be collaborative, both seeking to discover some pattern of relationships that has not been perceived by one or the other or both. Non-political scientific arguments are of this form.

What do I mean by "political" here? In a political argument, at least one of the parties does not control for finding a match between their two perceptions of an external reality. Instead, at least one of them controls for other people to believe that person's initial position to be true. A political argument almost necessarily results in a conflict if person X controls for other people not to believe what person Y asserts to be true.

A truth seeking argument is, metaphorically, a voyage of discovery. Both parties have a self-self-image of controlling for reducing their uncertainty about the truth of what they believe, at least about the topic of the argument. In joining the argument, each has an other-self-image in relation to the other party that corresponds to this facet of their self-self-image.

The protocol structure we have considered so far has considered the degree of belief each party has in the three levels of the three propositions. Those beliefs change as a consequence of the performance of the supporting protocols within the protocol hierarchy. That hierarchy is self-contained, consisting only of beliefs in the syntactic variables of the protocol. At the top is the degree of belief, not belief by either party in the truth of the Primal Message, but in the degree to which the continuer understands the Primal Message, whatever it might be that the initiator, deceptively or cooperatively, wants the continuer to believe.

In an Argument class of protocol, this self-contained processing network is opened to beliefs about the external world, both physical and social, in much the same way that a control loop is open to the effects of disturbing influences from the world outside the loop. These beliefs are likely to be controlled perceptions of something about the external world, but they may not be controlled. If they are controlled in both parties, the difference between their reference values or the embodied perceptions