Talk #6 Pattern and Structure

> A pattern is a form, template, or model (or, more abstractly, a set of rules) which can be used to make or to generate things . . . The detection of underlying patterns is called pattern recognition.
>
> – Wikipedia definition
>
> Self-similarity is a newly discovered symmetry in nature by which parts of fractal objects relate to their wholes. That is, the overall pattern of a fractal is repeated at multiple size or time scales, from small to large scale. Sometimes this repetition is exact, as with a linear fractal. Most often, especially in natural fractals, self-similarity is approximate or statistical. This nonlinear property allows fractals as they appear in nature to embody irregularity, discontinuity, evolution and change.[1]
>
> – Fractal Dynamics of the Psyche, Terry Marks-Tarlow

Thea: You were going to talk about pattern tonight – pattern as opposed to chaos, presumably. That's a pretty basic distinction, isn't it? The Bible begins by saying that all was chaos at the beginning, before God divided the Light from the Darkness.

I was thinking that if the patterns in the world were God-given, you might expect them to be more reliable. Why are the patterns of life – the patterns we depend on – dependable much of the time, but apt to betray us when we least expect?

Guy: That question is timeless. In one form or other, philosophers and theologians have been asking it for millennia, while witch doctors propitiated gods and spirits to keep the patterns on track. We need to explain why the world is as regular is it is. We also need to explain why it is not more regular – why its patterns change and break down. And then there is the question of how our brains can track the world's patterns – and within what limits they can do so?

Thea: Do you have an answer?

Guy: The eD paradigm takes us some way toward an answer. To begin with, it teaches us to think about patterns as natural phenomena, without putting ourselves speciously at their center – without seeing them as designs intended for our punishment or benefit. It explains how patterns can arise by themselves, and why they tend to be fairly reliable but not perfectly so. It helps us understand why we are very good at recognizing some patterns, but much less good at others.

what is a pattern?

Thea: Let's start with the concept itself: I know what a dressmaker's pattern is, for example; but I've never seen a general definition. From your perspective, what is a pattern anyway?

Guy: The shortest answer might be that "pattern" is another word for re-suggestive structure – a source of intelligible suggestion. You could define the concept as a type, or class-designator, of which repeated instances can be recognized. Patterns are found everywhere in nature – and

---

1 *www.goertzel.org/dynapsyc/2002/FractalPsyche.htm*

at all levels, from sub-atomic particles and atoms to galaxies and galactic clusters. If we could not recognize underlying similarities amongst things, there would be no things – only hopeless chaos. So patterns are just recognizable similarities in the world – recognizable from place to place, and from time to time. Patterned "things" – instances of familiar patterns – make up the world as we know it. What's striking here is that a pattern is not just a feature of the world, but a match between our brains and the environments they deal with: an evolved "fitness" enabling us to notice and respond appropriately to more-or-less reliable regularities in our environments.

Words like "structure," "system" and "organization" refer to patterns of patterns. Order is a word for the quality of being patterned. Systems may be organized from the outside, by the transfer to them of some external pattern but, as we've seen,[2] some systems show a remarkable tendency to become increasingly patterned all by themselves.

Thea: Let's stick with the bare concept for now. You say a pattern is anything of which repeated instances can be recognized. But a pattern can also be used to generate repeated instances?

Guy: Of course. As you said, a dressmaker's pattern is one type. So is a recipe, or a blueprint for a house. I should have said that a pattern is anything of which repeated instances or iterations can be recognized or created.

Thea: So which comes first, a pattern or its instances? It can go either way, can't it?

Guy: Ahh . . . That's an interesting question, with a long history. Plato thought that the patterns came first; and, in some form or other, his doctrine of "natural kinds" – the metaphysical reality of patterns – held sway for almost two thousand years, though the issue was hotly debated. Only with Darwin and the rise of modern biology has the balance tipped the other way: Scientists now find it more fruitful to see patterns as abstractions by human minds from specific cases – in fact, from modal instances that are particularly good at reproducing themselves (or getting themselves reproduced) and then at persisting for a while.

Thea: Your ecoDarwinian paradigm.

Guy: Exactly. Darwin was the first to suggested a way in which certain kinds of patterns – the species of biological organisms – could come into being spontaneously, as a consequence of the life process itself. Since then, it has proven fruitful to explain many other patterns in a similar way. Natural selection (later generalized to other forms of self-organization) became a very successful pattern of explanation.

In biology, the notion of species (a recognized pattern of living creature) gives way to clade[3] – a group of organisms that includes the most

---

2 In Talk #2.

3 From the Greek *klados*, meaning branch or twig.

recent common ancestor of selected individuals and all des-cendants of that common ancestor. For it turns out that the notion of species is incoherent: Not all the members of some species can interbreed, while interbreeding between species is sometimes possible.

Thea: Yet we do still continue to talk about different species. Cats are cats and dogs are dogs, despite the many breeds and mongrels of each. In general, when things are sufficiently alike, we regard them as being of the same kind. That's what it means to belong to a set, is it not? We recognize a set by the definition that all its members must satisfy.

Guy: In mathematics, yes. In ordinary language, no. In fact, there is evidence that the categories of ordinary language – even the common biological categories like cats and dogs and birds and trees – are not generated by definitions involving properties that all their members share, but in a different way entirely. Children learn to use such words correctly without being able to define them. Even an adult might be hard pressed to give adequate definitions of some very simple words like "glass" and "mug" and "cup." It seems that most word meanings are not given by definitions at all, but by prototypical mental images – mental patterns, in other words – to which current occurrences are assimilated as they are happening.

Thea: Then where do the prototypes come from? And how do we recognize what we see as falling under one prototype rather than another?

Guy: Our prototypes are taken from the culture that surrounds us, and from personal experience. That's why one child thinks dogs are cute and friendly while another is terrified of them. Their ideas of "dog" have been based on very different experiences. In your field, the point is a truism. At least since Freud's time, psychotherapists have explained to clients that their expectations and responses to other people are patterned on childhood experiences with "significant others" – people important to them in their earliest years.

In connection with these cognitive patterns of early childhood, it may be worth mentioning Carl Jung's idea that some prototypes (the so-called "archetypes") are shared by all human beings, perhaps because they've been hard-wired into our nervous systems. It's possible: Some birds respond with obvious fear to anything that looks like a snake. Vervet monkeys use three distinct alarm calls to warn each other of different kinds of predator. We humans may have our own hard-wired categories. Jung's idea may deserves more attention than anthropologists have given it.

Thea: But how does prototype recognition work? Why do we see something as one kind of thing and not another?

Guy: In the history of philosophy, that is another ancient question – the problem of similarity and difference. Given that things in the world keep changing, how different must a thing become before it is no longer itself but something else? When does it become a different kind of thing? Such questions are the basis for Aristotle's distinction between essential and inessential properties: Things that differ in essential properties must be considered different kinds of things. Things that differ merely in their

inessential properties are considered to be the same kind of thing but with different "qualities." But you have only to think of an Escher drawing, which turns one image into another, or the graphics software that can transform any image into any other through imperceptible changes, to realize that this simple distinction doesn't hold water.

In Nature, the boundaries of sets are not so well defined as they are in math. For this reason, Lofti Zadeh's "fuzzy logic" is very useful as an approach to such practical uncertainties of classification.

Thea: But fuzzy logic won't help if people are working from different prototypes, or different definitions?

Guy: No. And I don't think fuzzy logic does much for the problem of interpretations – the willful self-interested character of human perceptions. It was designed to apply in situations where there is agreement on definitions, but inadequate knowledge of whether their terms are met; and for this it is successful. It helps us recognize recurrences in space or time of things that are imperfectly described when we are agreed on what we are looking for. Or it helps us to recognize instances of a certain kind of thing whose particular specimens may differ greatly. Situations of this kind are common in Nature, where typical configurations recur spontaneously, but subject to variation. At different scales, snowflakes, hurricanes and stars are good examples. So are human beings, human families, human organizations. At every level of the cosmos, we see recurring similarities – not absolute identities. It is these similarities that we need to recognize and respond to appropriately; and it is these that our sciences hope to explain.

Thea: Recurring similarities, not absolute identities! We don't make enough of that distinction, do we? It's as if all Nature were like a great composer, writing dazzling variations on a few basic themes.

Guy: For a long time, that was the best theory of difference-in-similarity we had, and as a poet's dream it may still have value. One objection is that composers try to write only lovely and interesting variations. Nature writes all kinds, and then selects what is viable.

Thea: Well, we've talked about that, and will again. It's that randomness, that sublime indifference that people find so upsetting about your eD paradigm. Where thinkers once aspired to read the mind of God in the patterns of Nature, they now tell us merely to ask how the patterns copy themselves, and why some patterns fare better than others at doing so.

the propagation of patterns

Guy: Remember that when we speak of patterns being copied, or copying themselves, the metaphor is only roughly correct. Literal copying is done by an external agency – for example, manually by a scribe, or automatically by a photocopier. What happens in Nature is not always a literal copying process; so that term can give a false impression. It would be more accurate to speak of the propagation or transfer of patterns – remembering that in Nature several kinds of propagation are found, of which literal copying is only one.

Thea: Aren't you being a bit pedantic here?

Guy: Perhaps I am. But it seems important to focus attention on the process, rather than its outcome: a thing that more or less resembles its source or sources. To call that outcome a "copy" is bad enough – though nouns like "copy," "duplicate" and "replica" are hard to avoid. But to speak of the process as copying is misleading unless the process resembles literal transcription. And quite often in Nature this not what happens at all.

Thea: I feel a little old to ask for a lecture on the birds and bees, but how does this propagation take place, according to you? And before you get started, you might explain what you mean by the "propagation" of a pattern, if not the making of more-or-less accurate copies.

Guy: Any growth, extension, ramification or re-duplication of a pattern is a kind of propagation. Any process that leaves a pattern larger, or more intricate, or more numerous than it was before. The concept of patterns propagating (in this broad sense) points to a powerful and very general tendency in Nature that words like copying and even reproduction easily overlook. Patterns do something more than replicate, or get themselves replicated. They also tend to grow, vary, and swell their influence by branching out and specializing along different lines. In doing so, they come into competition with other patterns for whatever scarce resources they may need or want. In particular, to speak of an ecology of mind, we need to remember that the propagation of cognitive patterns involves something more than "reproductive fitness" – unless that concept is taken in a much broader sense than usual. Ideas and cognitive patterns expand and propagate in many ways.[4] Mere copying does not begin to describe the possibilities for doing so.

Thea: Yes, I can see that. The ways we influence each other, and what happens to us as we assimilate the influence of others is more complex – and, frankly, much more interesting. Theories of cognitive copying are dull, compared with the common-sense story of intentional influence and conflict.

Guy: You may be surprised that I agree with you on this. I think the new ecoDarwinian psychology has to do much better than it has to-date with the common-sense notions of beliefs, motives, intentions and influences. I think it has a potential to do this, but not in any simple-minded reductionist way. Replacing the concept of information with that of suggestion in the mind sciences may be one step in the right direction. Talking about the transfer or propagation of patterns, rather than simple copying, may be another.

Thea: It sounds to me like you are trying to have your cake and it eat it. On one hand you are applauding a biological, Darwinian psychology that treats the mind as a by-product of electrochemical processes in the brain. On the other, you hope to retain at least some aspects of folk psychology and the intentional stance. Don't you see a contradiction? Either we have some free

---

[4] See the discussion of meme theory in Talk #12.

will or we don't. Either we choose and tailor our ideas to suit our purposes, or they infect our brains like viruses, and run us like computer programs. Which is it?

Guy: Perhaps the contradiction is not as big as people fear. Perhaps all these metaphors are defective, and that our folk understanding has been approximately (but only approximately) true all along, though it is only now that we are beginning to be able to express it with any sort of precision. Or perhaps, the new psychology does shift our self-understanding in significant ways, but not to the exclusion of older views that still have their uses. The latter is the idea that I'd defend.

Thea: And good luck to you! But I'm not yet convinced it can be done.

Guy: Let's keep going. Perhaps a deeper understanding of pattern – and the propagation of pattern – will resolve some of these issues. I think it's already clear that patterns can propagate in very complex ways, and that cause-and-effect models of their propagation are inadequate.

In fact, patterns propagate by a variety of means, of which replication or direct copying is only one. Patterns also propagate by accretion as when an infant, or any young creature, grows and thrives on the food it eats. They propagate (in time, as an extension of life span) by replenishment as when resources are expended and replenished, and as component parts wear out and are replaced. They propagate by adaptive complementation, as when a participant in any relationship takes on the traits appropriate for dealing with fellow participants. Biological evolution is an example of pattern transfer from the environment to the species. Personal learning – facial recognition, skill acquisition, memories of events, and so forth – are pattern transfers to the individual creature.

We are far from a complete understanding of how patterns propagate, and how they can originate by themselves. But it seems clear, at any rate, that the general concept of pattern is sufficient to account for the richness of the world we know – given the further thought that propagation of patterns can be conceived as an interplay and ecological stabilizing of competing suggestions. The idea that science can and must try to reduce all explanation to mechanical cause is obsolete even in physics. In biology, it's very shaky. In the cognitive sciences, it makes no sense at all.

Thea: Bravo! But then what takes its place? A science of pattern?

Guy: Hopefully, yes – a general systems theory, complete with a general theory of pattern, influence, self-organization and ecological stability.

Thea: Do we have such a theory?

Guy: We begin to have one. Translating the concepts of biological and cultural pattern into suggestion language may help a little by encouraging us to think beyond the physical level of messages and information to the more abstract level of form and meaning – what a message means to the system that receives it, what it suggests that system become or do in response.

Thea: So your point is that an ecoDarwinian language about the propagation of

patterns should replace the classical language of cause-and-effect?

Guy: That's it. You have to begin by looking at the different ways that patterns can propagate, and by seeing growth, replication, behavioral influence, and so forth as an uptake of suggestions as well as raw materials. But the point was Gregory Bateson's not mine. [5]

Thea: What about the growth of crystals in a super-saturated solution, or the processes of star or galaxy formation? Many growth processes are driven by purely mechanical forces. It seems odd to speak of suggestions when the recipients are not suggers.

Guy: Perhaps we should think of molecules, atoms, and even more fundamental particles as very primitive suggers, moved by forces that we regard as trivial suggestions. I don't know enough physics to judge whether such an approach would be useful in quantum mechanics, for example. But it's clear that for complex chemical, biological or social systems it is hopeless to think of observed effects as produced by a single "cause."

Clearly, the notion of suggestion – and of organization as suggestion ecology – applies most naturally to complex production systems that possess significant autonomy. The mechanical systems you mention are troublesome borderline cases – like viruses, which may or may not be thought of as living things. My claim, once again, is that the notion of "copying" does not begin to describe all the ways that a given pattern may seek to organize its environment as an extension of itself.

Thea: "To organize the environment as an extension of itself." That's a suggestive phrase! Therapists are often concerned with what you'd call the transfer of patterns between our clients and their life-worlds. Family therapists are concerned with pattern transfers between family members, and from the family's social environment.

Guy: We need to look at psychological causation as a propagation of cognitive patterns subject to various constraints. People and groups are autonomous systems, with considerable but not unlimited freedom to organize adaptive, sometimes really creative responses to the suggestions we receive. Our cognitive adaptation is never purely passive; on the other hand, it is not completely unconditioned either. In general, it will be an outcome of cues and nudges from all quarters – all indicative (in the literal sense of "pointing"), none definitively controlling. At one extreme, the totality of suggestions acting on a system might be what we think of as a vector sum of forces. At the other, it might be some vague policy by which specific actions are guided.

Thea: Could you give an example?

Guy: Easily. Why are we having this conversation? Why are we living together, for that matter? What it comes down to is that for both of us, the

---

[5] See Bateson's *Steps To an Ecology of Mind* and *Metapatterns: Across Space, Time and Mind* by Tyler Volk.

suggestions that we stay together are stronger than the suggestions that we break up. Or take a great historical event like the outbreak of World War I. Probably a few million pages have been written about how and why the assassination of an Austrian duke by Serbian nationalists led to general war in Europe, but there is no clear answer to the question of what caused the war. There had been serious tension amongst the European powers since . . . well, since the end of the last major war; and there were any number of diplomatic crises over one thing and another. What was special about this particular crisis? Why was it not resolved or contained as the crises before it had been?

Thea: Tell me.

Guy: Not just now. My point is that the general answer to a request for explanation must be a loosely stable patterning of suggestions. In the case of World War I, what's needed is not just an account of the events that led to war, but of the meanings of those events (that is, the suggestions they conveyed) to two dozen statesmen and their political audiences. In the same way, suppose you want to understand why a client is having a certain problem; or why the fertilized egg of a cat develops into a kitten and not a puppy. Questions like these cannot be answered in causal terms, nor yet in statistical terms; they demand some other kind of answer entirely.

It will be fruitful, I believe, to approach such questions through a language of ecology and suggestive guidance. Human activities – the activities of living things in general – are shaped in response to a myriad of cues and nudges, toward outcomes that we cannot predict in any strict sense, but that we can sometimes anticipate with fair reliability. For example, we can anticipate that a pregnant cat will have kittens if she has anything at all. We know she will not have puppies; but we do not know for sure that she will bring her kittens to term; and we don't know what kind of kittens they will be. We can't predict their markings or their personalities, for example.

Thea: That's an interesting distinction you draw between prediction and anticipation. Certainly, when it comes to human development or behavior there are all kinds of things we reasonably anticipate that we can't actually predict. Yet I don't see how you would draw the distinction except on a statistical basis. Even that would be awkward. How probable does something have to be to graduate from a mere anticipation to a real prediction?

Guy: I don't believe the distinction is a matter of probability at all. Astronomers predict when the sun will rise tomorrow morning and at which point on the horizon it will appear. But I anticipate that event in general terms without ever looking at their calculations. I anticipate another sunrise mostly on the basis of hearsay and my own experience, and partly from a scientific story I've been told about the Earth's rotation. But I have no idea how to calculate when the sun will come up. Though I'm familiar with Kepler's Laws and with Newton's Theory of Gravitation, I don't normally think about the sunrise in astronomical terms. Similarly, at the beginning of the last century, many people anticipated a European war sooner or later. Because they anticipated war, they prepared for it; and their anticipations

and preparations were among the factors that made it happen. Their anticipations were factors suggesting war, I would say, against competing suggestions to keep the peace. All this is common knowledge, but no one has any idea how to calculate the probabilities involved or the details of the war's coming, so we cannot really speak of prediction.

Thea: After the fact one can usually explain what happened in terms of suggestion and anticipation. But what can we say that's useful about correct anticipation before the fact? About what we should anticipate, and what our anticipations should suggest we do?

Guy: We may never be able look at an ecological situation, collect data on it, and compute its likely outcome. Ecological systems are usually too complicated, with too many layers and variables. But we may sometimes be able to anticipate a rough outcome, or range of possible outcomes, based on our understanding of a system's dynamics, and of the factors influencing it.

The best approach we have at present is to build computer models of a system, set up initial conditions that interest us, and watch the models run. At present, what we are mostly learning is how to build better models. But that may change as we learn to trust our simulations – as these approximate more closely to what is observed, and our anticipations become more reliable.

Thea: In effect, the models themselves are evolving.

Guy: Yes. And in time, may be able to do so by themselves, based on data supplied or gathered by them automatically, without explicit human programming.

Thea: Such evolving simulation programs would be just as difficult to understand as the world itself.

Guy: True. But if we can write programs that evolve then, possibly, we could write them to document and explain themselves. These simulations would be discussed on the World Wide Web, of course, so we could anticipate a Net gain in our understanding.

Thea: Instead of responding to that atrocity, I want to go back to your statement earlier that the general answer to any request for explanation will be a pattern of suggestions that compete and reconcile. I have to agree with you about this. Therapists know that the features of our clients' personalities do not have straightforward, simple causes, and are not readily predictable from parenting and early experience. Whatever risk factors can be found are seldom either necessary or sufficient. I fully agree that for human beings, and probably for any system of sufficient complexity, the whole idea of cause is more confusing than otherwise.

But it's such a difficult idea to give up! We are completely accustomed to think of ourselves as agents who cause things to happen – sometimes intentionally, sometimes not. It's hard to see how we could do otherwise.

Guy: Here's a simple example – well within your and my and every adult's experience. In the phase of life called adolescence, young people explore a

larger world – beyond the boundaries set for them as children by their immediate families and authorized teachers. In doing so, they are subjected to a wide spectrum of suggestive influences – some wholesome and potentially life-enhancing, others not. The problem they face is, first, to evaluate all these suggestions, accepting some and declining others; then to survive and learn from the experiences they undergo; and then to shape their personal manifolds of experience, competence and relationship into some kind of adult life.

All of us remember the difficulties of this period. People our age remember watching and trying to be helpful as our children went through it. This is a perfect example of the problems of anticipation, explanation and pattern ecology in a system no bigger than an individual life.

Thea: If the mind is an ecology, then adolescence is a jungle of weird growths struggling for existence. It's scarcely possible to get at the causes of adolescent life-choices. All we can see are the various influences on a young boy or girl – what you call the patterns of suggestion. And yes – it seems much more hopeful to ask how these patterns propagate and interact with one another than to ask the causes of the choices made. But what can you say about those spreading patterns? That's the root idea of suggestion, isn't it? – how patterns spread from one system to another?

Guy: You probably know as much or more about such transfers than I do. The cybernetic concepts of symmetry and complementarity, calibration, feedback and feedforward, triangulation and so forth are pretty well understood by family therapists. All I want to add is that we should think of communication as a propagation of suggestion rather than information. But that is just a refinement of concepts, not a change of direction.

My point is that a suggestive influence does more than inform, but less than control. In combination with other suggestions, perhaps competing or conflicting ones, it guides the receiving sugger in constructing a pattern of response from its available repertoire. Such response patterns in turn send messages that are received and parsed as suggestions by other suggers, as they've been prepared by their life-histories to do. Instead of causes, we should imagine ripples of influence that cross and re-cross until a holon sub-system finds some kind of loose balance.

Thea: What will we say about human agency, then, if events in the world (including other people's actions) are not directly caused by what we do, but only more or less successfully influenced?

Guy: Beyond the local level – what we can directly manipulate with our own bodies – it doesn't always amount to much. Sometimes the ecological effects can be neglected and we achieve our intended results. Sometime they can't be neglected, and we encounter systemic effects that may run completely counter to our wishes.

Thea: "Man proposes but God disposes," as the saying goes. That's such a difficult lesson to learn.

Guy: Well, ecological relationships can be weird, but at least they are not supernatural. Though we can't usually predict how they will play out, we can

sometimes anticipate – get a feel for their possibilities through model-building and direct experience. They are at least potentially intelligible.

Thea: The most confusing feature of these ecological relationships is the combination of inter-dependence and ferocious competition. We try to know our friends from our enemies; and biologists try to distinguish between predators and symbionts; but so many people and species seem to be both at the same time. Some plants rely on the animals who eat them to distribute their seed in the stool. I remember reading that Arctic caribou follow the Inuit peoples who hunt them, as much as the Inuit follow the caribou. Human workers depend on bosses who exploit them to give them jobs.

Guy: Remember what we said a few days ago about the politicious quality of ecological interaction? It's characteristic of ecological relationships that they are usually cooperative and competitive at the same time. Patterns in general must thrive with each other's help, but also at each other's expense. They sustain themselves in balanced configurations of mutual dependency and encroachment to the extent they endure at all.

Thea: Dare I ask, configurations of what? The things and beings we find in Nature always turn out to be configurations of patterns which are themselves configurations of patterns. Nothing is solid. When you look closely at anything it dissolves into components which themselves dissolve when looked at.

the re-combinant holarchy

Guy: Here, Arthur Koestler's notion of holarchy is useful: The cosmos can be seen as a hierarchy of entities that he called holons because they exist as wholes-in-themselves at the same time that they are parts of something else. The coined word serves to emphasize a group of properties that all discernible entities share:

- Holons are Janus-faced: they exchange suggestions with holons above, below, and on the same level in the hierarchy; and they are bound into relationships by the suggestions they exchange.
- Holons are "greater than the sum of their parts," in the sense that they typically show emergent properties that their components do not possess. Lower-level holons are typically more mechanical, routinized and predictable in their behaviors. Higher level holons are usually more flexible, and more abstract.
- Holons are mutable: They can preserve an integrity as wholes even as their components are altered or replaced.
- Holons are re-combinant: They can serve as component parts of various holons above them in the hierarchy; and they can transfer their participation and service from one superior to another.
- Holons are inter-permeable and open-ended: It may be convenient to analyze them separately, or to treat them as interwoven with other holons. Also, while it may be convenient to think about and study holons with a lowest and/or highest level, we can always decide to analyze them further, or consider them as parts of a still larger whole.

To summarize: The cosmos reveals itself as a re-combinant hierarchy of Koestler's holons that we are calling "suggers" to emphasize how they respond to and are guided by each other's suggestions. Each holon, as much as the wholes of which it is part, is a hierarchically-structured, multi-layered system. To the extent that a holon subsists in quasi-stable equilibrium with itself, we can regard it as an ecology. Whatever we call such entities, they are the systems and structures that we encounter at every level of being. We ourselves, our groups, human society, and the universe as a whole are systems of this kind.

Thea: My head is spinning. I think I need a few aspirins. More than two, probably.

Guy: I don't blame you. J.B.S. Haldane once remarked that the universe is not only queerer than we suppose, but may be queerer than we can suppose. It's not at all clear to what extent human minds can encompass such complexity. Our brains, after all, evolved for survival on a savanna, not for contemplation of the cosmos. It's amazing that humans have managed to learn and understand so much.

Thea: Amazing and frightening. I doubt that a world view of such complexity can gain acceptance by more than a fraction of the world's people. It's not surprising that the wars of religion are coming back, and likely to get worse before they get better. Making wise use of this kind of knowledge is yet another question. What only a few can understand affects the lives and prospects of everyone. We may very well destroy ourselves before we learn how to live with it.

Guy: I'm forced to agree. I think there are legitimate concerns on both sides, but the quarrel today between the naturalistic and traditional world views is pursued so ignorantly, for the most part, that the result is unedifying, to say the least.

Thea: Well, I hope you and I can do better with those concerns than the global conversation is doing. What I take from this holon concept is that the patterns in our lives combine and re-combine with other patterns, and that our world is built like a set of Russian dolls – level within level within level – with the difference that there are many dolls, not one at each level, and that the components of the dolls keep changing. The possible combinations of patterns into larger patterns and the potential for chaotic change are practically infinite; however the patterns are stabilized to some extent by the principles of self-organization that you've described.

Guy: Very good. Basic chemistry affords the clearest and best understood example: Just a little over a hundred different types of atoms can produce a practically unlimited number of chemical compounds, each with characteristic properties that could not have been predicted just by a knowledge of their component atoms. The same principle can be seen in human social arrangements – for example, when various job skills are brought together in a business organization. Similarly, within the brain itself, patterns of neural firing re-combine with one another to produce the mind's sensations, thoughts and feelings. Adaptive intelligence is possible because the recurring patterns of our environment are matched somehow by

patterns in the brain and mind.[6] Wherever we look in Nature, we find this same meta-pattern of re-combinant holons, arranged hierarchically and guided by each other's suggestions.

Thea: Why do holons have to be structured hierarchically? Why not some more egalitarian relationship?

Guy: Actually, there are good reasons why hierarchical arrangements are favored by evolution. It's basically a result of the power laws that we discussed in our second talk, and the principle of re-combination that we just discussed. Evolving systems must deal with change selectively. They must resist some changes but accept others; and they must preserve their patterns in some respects, while accepting change in other respects. Hierarchical structure is a simple way, and perhaps the only way, to meet this requirement for relative stability in the face of change. As Robert Michels and Herbert Simon recognized some time ago, hierarchy allows for an efficient blend of local adaptability with over-all coordination.[7] It allows potentially disruptive changes to be contained and localized. It allows for easy re-combination of the modular components when their high-level organization is disrupted.

Thea: Why would a self-organizing system require over-all coordination? Why not rely on the "Invisible Hand" of evolution to provide coordination?

Guy: That is a very good question, and a vexed political one, which needs an evening to itself. I suggest we put it aside for now, and come back to it later when we are ready to talk about the relationship between society and government.[8]

Thea: Just one more question then: You seem to be using the words "system" and "structure" almost interchangeably – describing them both as networks or hierarchies of suggestion processing holons. What's the difference between a system and a structure?

Guy: So far as I can see, those words just invoke different perspectives: A structure is an assembly of parts. The concept of structure is synchronic, pointing to the static relationships amongst the parts at a given point in time. The concept of system suggests a diachronic perspective. It asks us to look at the configuration as a dynamic, functioning entity – subject to change and interaction with its environment. Normally, we'd think of an office building as a (relatively static) structure, and a living cell as a (dynamic) system; but the building includes systems for heating, vertical transportation, and other functions, while the cell contains various structures: a cell membrane, a nucleus, and so on. The usage of the words is not always consistent, but so far as I can tell they are synonymous, except for these different perspectives on time and change.

---

6 As will be discussed further in Talk #7.

7 See Robert Michels' book *Political Parties* and Herbert Simon's *The Sciences of the Artificial*.

8 In Talk #13.