Order and Self-organization

> Man follows the ways of earth; earth follows the ways of heaven. Heaven follows the ways of Tao, but Tao follows its own ways (*zi-ran*).
>
> *Tao Te Ching*, Verse 25

A few thousand years before Ross Ashby coined the term 'self-organization' in 1947,[1] Chinese Taoists had observed a tendency of systems to evolve toward a balance between centripetal and centrifugal forces, or between processes of intake and outflow (between the *yin* and *yang* as they called these processes respectively); and they already spoke of *zi-ran* – the self-so, that which happens by itself, as distinct from that which humans have to design and build. In the *Tao Te Ching*, the last lines of verse 25 are about *zi-ran*.

Today, self-organization is seen as a very general phenomenon, observed in physical systems at all scales from atoms to galaxies, and from organic molecules to multicellular organisms and whole societies of humans and some other creatures. There is a certain irony here: Where medieval thinkers saw the universe as a "Great Chain of Being," designed and called into existence by an almighty God, Oriental thought took a very different path. Having this concept of 'the self-so,' it did not need and did not develop our Western notion of a cosmos structured by divine law, and called into being by divine fiat. The irony is that our Western idea now seems to have been a very fruitful mistake. Given a concept of 'universal law' Western thinkers naturally inquired what that law might be; and seeking to know "the mind of God," they invented modern science. Only with Adam Smith and Darwin did our science begin to understand that order can sometimes happen all by itself, and it is now learning how this can happen. From a cosmos ruled by divine law, we are tending these days to think about a self-organizing cosmos in which the apparent 'laws' are themselves evolving patterns.

One reason, I think, why so many otherwise modern people still cling to the notion of a Creator God, is that people brought up in our traditional Western mindset – on ideas of conscious design and implementation whether in a Biblical world or a Newtonian one – tend to find the idea of causeless order (the 'self-so') not just counter-intuitive, but literally incredible. Some take on faith what the scientists are saying, but few really grasp how order can emerge spontaneously from chaos – how the self-organization of a universe was possible. And not just possible, but the most plausible creation story we have. It takes a fair amount of work to understand how order can emerge from chaos all by itself, but once this is done we can really let go of the idea of a designer God and consider implications of the modern worldview for human existence and identity. If

1 See also the FAQ on this subject at www.calresco.org/sos/sosfaq.htm.

so inclined, we can go on to consider possibilities for a spirtual and even religious life in a self-organizing cosmos.

In this appendix, we'll review the concept of order as such – apart from, or independent of, a material medium in which it occurs. We'll review what is known about the ways that order can increase spontaneously – all by itself, with input energy but without external guidance of any kind. And we'll discuss what the concept and mechanisms of self-organization can tell us about the evolution of life and mind in a purely physical world, when dualism, and the supernatural are put aside.

1. What is Order?

Physicists speak of '*order*' in a system when there are symmetries or other correlations amongst its separate parts. Conversely, they speak of 'disorder' or 'entropy' when there is homogeneity or randomness. The Second Law of Thermodynamics tells us that the homogeneity (entropy) of any closed system tends to increase with time. For example, hot coffee poured into cold milk becomes coffee with milk at a uniform temperature. The reverse does not happen – or happens with only negligible frequency. The molecules of air in a sealed box tend to fill up the available space, and do not cluster together in one corner, or segregate themselves by type in the eight corners.

But this 'Second Law' has two great limitations: First, it is only a statistical expectation, not an absolute necessity. Given sufficient time, the molecules of nitrogen, oxygen, carbon dioxide and other components might group themselves in the various corners, though they would not remain so for very long. Second, this 'Second Law' does not apply for *open* systems which receive energy (in whatever form) from outside their box. Given some input energy, the statistical tendency to homogenity and randomness can be reversed dramatically. For example, the complexity of life on earth is the gift of energy from the sun. In fact, as Figure 1 shows, our universe is filled with complex structures at all scales, from atomic nuclei to stars (stellar systems) and galaxies, with ourselves somewhere between. Above the level of a hydrogen atom, complex order is sparse in the universe, but definitely does occur. While the laws of physics don't exactly encourage complex order in the universe, they surely allow it to occur. They make it more likely than not under certain conditions.

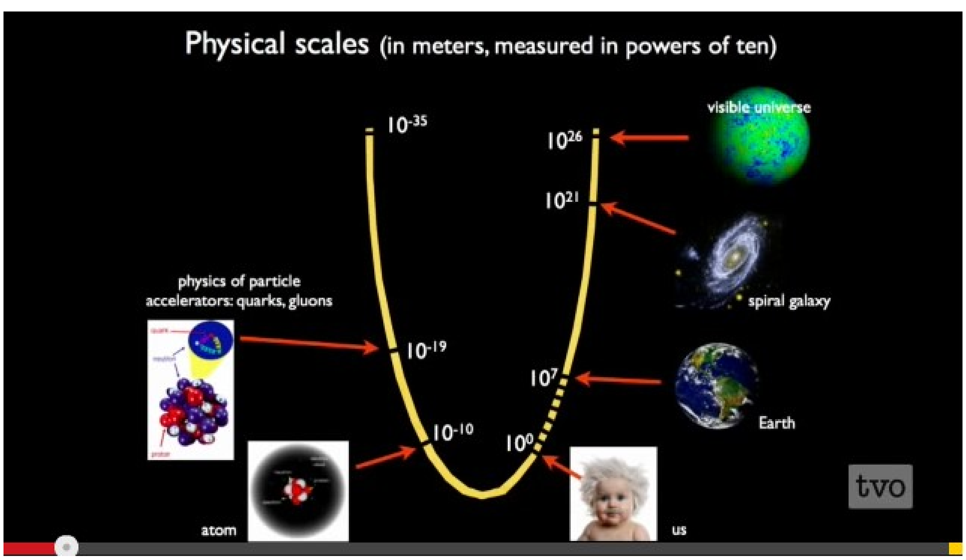

Pattern

We know order when we see it, but it is hard to define without circularity. In general, order is pattern, a discernible repetitiveness or regularity in nature or in a man-made design. In physics, one of the earliest studied examples of spontaneous pattern-formation were the so-called Bénard cells that form in a gently heated liquid, as shown below.[2]

Crystal Formation and the magnetic properties of compass needles are similar examples – long known but only recently understood. In every case, an ordered system has been driven away from homogeneity and thermal equilibrium by some input of energy, to display a regularity or pattern of which many types are known. Some types of pattern are reviewed in the remainder of this section. Some mechanisms of spontaneous pattern formation are briefly introduced in Section 2.

2 See http://www.eoht.info/page/B%C3%A9nard+cells

Symmetry and symmetry-breaking

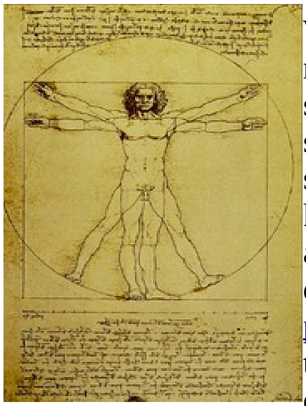

Whether it be the bilateral, twofold symmetry of a mammal, the fivefold symmetry of a starfish, the sixfold symmetry of a snowflake, the radial symmetry of splashing water, the spherical symmetry of a star or planet, or whatever other kind, very many systems in nature are approximately symmetrical – almost the same in different directions. (See the Wikipedia article on patterns in nature for examples.) And corresponding to the different types of symmetry, there are different underlying causes. Animals that move in one direction necessarily have upper and lower sides, head and tail ends, and therefore a left and a right. Evolution ensures that the head is specialised with a mouth and sense organs, so that eating occurs at the front end, and excretion at the rear. The body becomes bilaterally symmetric, though the internal organs are not.

Sixfold symmetries derive ultimately from the properties of an equilateral triangle. With every angle precisely 60°, just six such shapes can pack in the 360° around a central point. The fivefold symmetry of starfish and most other echinoderms is rather puzzling – perhaps an evolutionary accident, and thus an example of the symmetry breaking that we should mention next.

Think of a pencil balanced on one end. It is in equilibrium for the moment, but this equilibrium is unstable. A puff of wind will make it tip and fall to a horizontal position. In the process, its radial symmetry is broken. Standing on end it pointed only at the sky; lying flat, it points in some particular direction.

Symmetry breaking happens unpredictably, randomly, or sometimes intentially, and it brings *history* into the world. Beforehand, a fall in any direction was possible. Afterwards, the pencil lies in some given way which is what it is, but might have been different. Various kinds of symmetry and symmetry breaking find many applications in modern physics, where they afford a structure and coherence to the laws of nature, and are used to guess at further laws.

Branching (or Converging) Trees

The tree structure – of limbs or nodes branching off a central 'trunk' is also very common both in nature and in the design of human systems (like computer programs and complex machines).

The tree pattern is widely used, both in man-made and evolutionary design, because its hierarchical structure allows fresh growth or building, in whatever direction, from an already fixed base. This enables a “divide and conquer” strategy in computer programming, for example. Conversely, it creates a possibility of convergence, in which things originally separate are joined together – as when streams of rain water join up into rivers and make their way to the sea.

One reason why hierarchical 'tree' patterns evolve is that they can be very stable. Perturbations at the lowest levels may cause damage and/or force changes while the order at higher levels is undisturbed.

Fractals

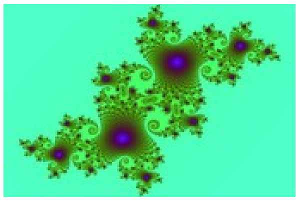

Fractal patterns are ubiquitous in nature, (in clouds, coastlines, ocean waves, mountain ranges, and certain vegetables, for example). They are widely used too in computer graphics, to generate artificial images that look like photographs of natural processes and objects. As a mathematical idea, fractals are repeating patterns which are self-similar on every scale – and are now understood to have a special kind of symmetry. The branching pattern, just discussed, *would* be fractal if could be repeated indefinitely.

Fractal symmetry is subtle because it does not look symmetrical. On the contrary, it allows for endless variation of detail in its features. Yet it results from the same rule (or underlying form) repeated over and over again, so that this structure remains the same at every level. This is why it is so common in both physical and biological processes.

Spirals

For example, spirals can be seen as fractal forms – an important special case of self-similarity. Many plants and some animals grow in spirals as they use their present state and shape as a template for further growth. The nautilus is a well-known, beautiful example. As it grows, each chamber of its shell is an approximate copy of the one before it, scaled by a constant factor and arranged in a logarithmic spiral.

Films, Bubbles and Foams

A film (see next page) is a very thin surface surrounded by air – formed, for example, by dipping a loop of wire into, soapy water. The film is held together by surface tension, and its surface as minimum area for the loop that forms it – due to the principle of least action, which will be discussed below. A bubble is a very thin film enclosing a hollow sphere – the smallest possible surface area for the volume that is enclosed. A foam is a structure of bubbles, formed by trapping pockets of gas within some liquid or solid. Foams of many different materials occur in nature, and foam patterns are common.

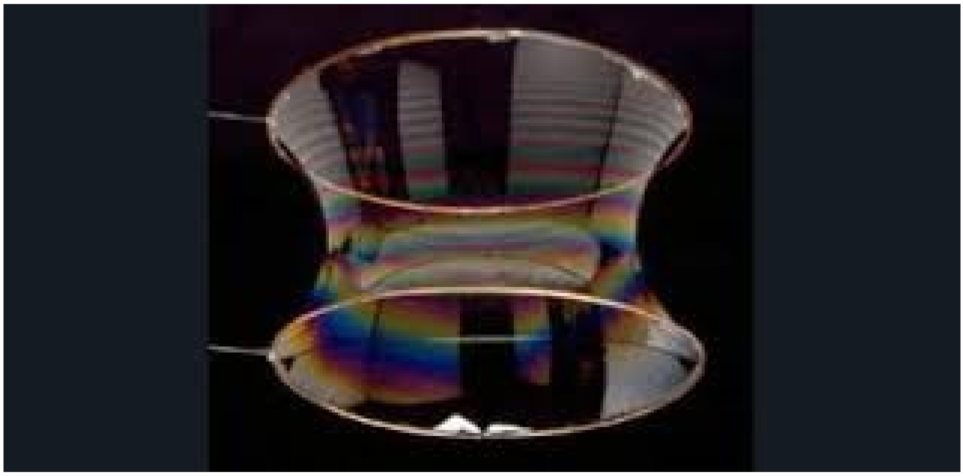

Waves

Waves are disturbances that carry energy as they move, typically through a medium like air or water, though light waves propagate through empty space). C**apillary waves (aka ripples) travel** along the surface of a fluid, with dynamics that depend on surface tension. Waves in the air, generated by speech or musical instruments, are interpreted by the brain as sound. Light waves, received by our eyes and brain are interpreted as shape and color. Light waves, in fact, are a special case of the electro-magnetic waves which are responsible (depending on wave length) for the phenomena of gamma rays, X-rays, radiant heat, visible light, and radio. In a vacuum, electromagetic waves propagate at a constant velocity (***c***) which is the same for every observer and one of the fundamental constants of the universe. Because of their importance in nature, the types and properties of waves have been extensively studied, and their math is highly developed.

Wind waves over the sea or lake make a chaotic pattern that will rock your boat or tip it. Wind waves over large bodies of sand create dunes, which form many different patterns, as will be discussed below in connection with the mechanism of self-organized criticality (SOC).

Crystals

Crystals are the patterns that make most solid objects possible – binding its atoms or molecules into some regular configuration. Examples of crystaline structure include snowflakes, diamonds, table salt, most metals, rocks, ceramics and ice. There are, however, some amorphous solids like glass, wax and many plasics whose atoms have no periodic structure whatsoever. The crystal lattice can be thought of as a 3-dimensional array of 'small boxes'; its symmetry is constrained by the requirement that these boxes stack perfectly with no gaps. A crystal's structure and symmetry play a role in determining many of its physical properties.

Temporal Patterns

Here we'll just consider rhythms – events that recur at regular time intervals. Other temporal patterns may be constructed, even without a rhythmic basis, yet rhythmic patterns are found everywhere; in fact, most of the systems that we perceive or try to analyze, are rhythmic or periodic. There is a simple reason why this is so: If the states of a system do not recur – at least approximately – they do not draw attention, recognition and study unless they are spectacular in some way. Their periodicity may fail in the long run, as our own bodies do, but until they cross some tipping point, they maintain some more-or-less regular orbit. The world is made of systems of this kind.

As the prototype of all periodic systems, consider the ideal case of a frictionless pendulum, swinging back and forth by converting some amount of potential energy (acquired when it was pulled up to a certain height) into an equal amount of kinetic energy at the bottom of its arc, which then returns to potential at the same height it started from, but on the other side of its low point. Many similar systems can be found, e.g. the motion of a child on a swing; the vibration of a spring or a guitar string, or of the air in an organ pipe; the swaying of a stalk of wheat or a tall building in the wind; the bouncing of a rubber ball; the elliptical orbit of a planet around the sun, or of a satellite around the Earth. The energy of a pendulum may be pumped in some way, as when you push a child on a swing. Or it may be damped, as when the child drags her feet in the sand to stop the swing and get off. Two or more such systems may be coupled together to build systems of practically unlimited size and complexity.

Now think of a vast cloud of hydrogen atoms in almost empty space. As the hydrogen is compressed by gravity, its atoms will move faster and the temperature will rise, eventually reaching a point where nuclear fusion gets started. The result is not a sudden explosion but a self-organized and self-sustaining process: a star.

Much like a candle that melts, vaporizes and ignites its wax with the heat of its own burning, every star in the universe is a self-sustaining process which endures over its lifetime in dynamic balance between the immense gravitational (*yin*) forces pulling it together, and the (*yang*) forces blowing it apart. Around many stars a solar system of some kind evolves, with planets in stable orbit around their star, and moons in stable orbit around their planets.

Given a stable solar system, a "Goldilocks planet" may exist – not too hot, not too cold – and with ample carbon, water and other necessary materials – on which some form of life may evolve. And a living organism too is a very complex periodic system which maintains and replicates itself by alternately collecting and expending energy. We know now that much of the complexity of our universe, and much of its relative stability, can be seen as the effect of one simple dynamic: the tendency of things to self-organize rhythmically into a holarchy of periodic systems.

* * * * *

Our universe is estimated to be more than 13 billion years old. If the second law of thermodynamics were the whole story, we would expect it to have become much more uniform by now than it actually seems to be. Nor would there be creatures like ourselves to observe its complexity and be puzzled by it. But, in fact, we know today of several ways through which natural patterns tend to form, and can explain much, though not yet quite all of the world's complexity by appeal to these known mechanisms, and to combinations of these. Still other pattern-forming mechanisms, one suspects, are waiting to be discovered – especially at scales to tiny or to vast for humans to probe experimentally:

- the Planck scale (down at 10^{-35} meters) at which the fabric of space and time is knitted by relativistic and quantum effects working in combination;
- the cosmic scale (above 10^{26} meters, the current size of the visible universe) at which whole universes bubble into existence.

Here I will describe eight order-forming mechanisms, roughly in the order in which they were noticed and described – in some cases, long before self-organization was perceived and studied as a general phenomenon.

2. Mechanisms of Self-Organization

Under various names – e.g. self-organization, complexity theory, and synergetics – spontaneous pattern formation in nature is now intensively studied, not just in numerous specific cases, but as a general phenomenon. While each case is unique in some respects, it is now believed that certain general principles of self-organization apply across the board, shaping the patterns of nature at every scale and wherever found. Bishop Paley's famous watchmaker argument was not stupid in 1802 when it was written, but it is obsolete today – partly because philosophers can show that the argument is logically invalid and fails to prove what it purports, but much more because we now know of several ways in which complex order may come to pass all by itself, with no need for an intelligent watchmaker. Indeed, it now seems far more likely that all 'intelligent watchmakers must themselves be products of natural mechanisms like the following:

Equilibrium or Balance

We've already mentioned the ancient Chinese Taoists with their concept of a natural balance between the (female) *yin* processes of absorbtion and accumulation and the (male) *yang* processes of radiance and effusion. The original perception may have been the contrast and mutual interdependence between sunshine and shadow on a mountain, or on the roof of a house. It was observed that in the long run, *yin* and *yang* had to be in balance: excess of *yin* produced an accumulation, a surplus, that would discharge somehow, while excess of *yang* produced a deficit, a partial vacuum, that would be filled. This idea of necessary balance underlay a whole cosmology, a political philosophy and a school of medicine – nearly the whole of ancient Chinese civilization. The Greeks too, with their concept of *sophrosyne* (moderation), understood very well that extremes were dangerous and that "too much of anything is toxic."

Under whatever name, this dynamic principle of self-organization provides a tendency toward equilibrium, and maintains a system in loose equilibrium, or periodic (possibly chaotic) oscillation, once such a state is found. Through friction, the energy of a pendulum gradually leaks away until that system comes to rest. Through the balance of torques about a fulcrum an ancient balance scale measures weight. Through the balance between centripetal and centrifugal forces, a planet in orbit around its star remains so, despite small perturbations from other masses. Through feedback, by maintaining the balance of temperature, acidity and various chemicals in your bloodstream, your body keeps itself alive.

Self-similarity

Another simple way for a system to self-organize is to get itself into a loop. When this happens, we see a tremendous amount of activity going nowhere in particular, but creating a loosely stable state. One example, already mentioned in connection with equilibrium, is the orbital motion of our planet around the Sun. More interestingly, however, we often observe repetitive activity that *does* goes somewhere – for example, the turning wheels that move a wagon; the cycling engine that powers a car. Such a system is stable insofar as it repeats itself; but at the same time it is *unstable*, or only loosely stable, insofar as small changes accumulate until a 'tipping-point' is crossed, at which some very large change results. The general form looks like a spiral: basically cyclic, but with a tendency to expand or contract or rise in level. The spiriform mollusc has already been cited as an example. Already Aristotle, writing around 350 BC, had noticed this pattern of self-similar growth.

For any growth process – of plant from seed, of chick from egg – today's growth always begins where yesterday's left off. In bio-chemistry, self-similar cycles of this kind are called autocatalytic or re-entrant, or autopoietic loops, and thought to be the origin of life itself.

In many such systems, a small residue or "gain" accumulates in the successive repetitions. Also, the system may experience small perturbations making the growth pattern somewhat irregular. As the result of an unusually large disturbance, or eventually, after enough repetitions, the system may undergo what we observe as drastic, qualitative change of state, crossing a pass into a different 'basin.' Provided only that some basins of attraction are easier to get into, and/or harder to get out of than others, a form of self-organization will result. As Gregory Bateson put it, "Longer lasting patterns last longer than patterns which last not so long." So understood, evolution is no more than a tautology – a necessity of logic.

Least Action

The principle of least action, first discovered by Maupertuis in the 18^{th} century states that a certain mathematical property (called the *action) of any dynamical system* is 'stationary' (at a minimum, but sometimes a maximum or a saddle point) for any natural trajectory. Many systems, including those Bénard cells, configure spontaneously to minimize one of their parameters. Soap films on a loop of wire configure to minimize their surface area. Below a critical temperature, bars of iron (and other materials) spontaneously magnetize themselves because alignment of the magnetic fields of their particles represents a "least energy" configuration for the bar as a whole. A beam of light, passing from a fish in the water to your eye is refracted (bent at some definite angle) to make its trip in the least time.

Think, for example, of a chain hanging between two posts or pillars to form the curve called a catenary, which looks similar but is mathematically quite different from a parabola. If you lift the chain at any point and then drop it, it will fall back to this exact same shape. It will always do so – thereby demonstrating the ability of a physical system to self organize in the simplest and most obvious way. The hanging chain is a classical, Newtonian system; but this same principle of least action is crucial in quantum mechanics where it leads directly to the Schrödinger equation.

Lobster Trap Effect

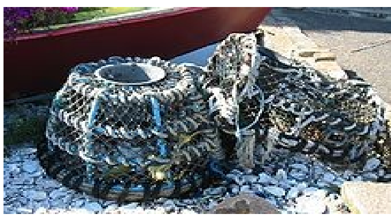

I don't know when the lobster trap effect was first recognized and applied theoretically. The trap itself must be ancient: just a simple box with bait inside and a funnel leading into it, so that the lobsters can easily find their way into the box but cannot find the small opening leading out. The general principle is that a change process can be one-way only, or overwhelmingly more probable in one direction than the other, so that random arrivals produce a steady accumulation of order. Through this mechanism, atoms and molecules (or whatever components) can *self-assemble* into complex, highly ordered structures – intricately folded molecular chains, and sheets and lattices, for example. This is the explanation for most teleonomic processes in biology and chemistry. It helps explain why an acorn develops into an oak tree, and why a fertilized human egg cell develops into a baby.

Lock-and-Key Mechanisms

Similar to the lobster trap effect is the so-called lock-and-key mechanism, exploited by the genetic code, and also by the body's immune system, by its sense of smell, by neurochemical triggers and inhibitors in the brain, and elsewhere. Again, the basic idea is very simple: A complex molecule may be folded in such a way that only a complementary molecule folded in corresponding fashion can bond with it. Molecules drift around at random, or over a fixed surface, until a 'key' molecule approaches its 'lock' at something close to the correct angle. Then, like the lobster it is 'trapped,' and some effect is triggered.

Swarming

Swarming (aka flocking, schooling or herding, depending the species at point) is collective motion by a large number of autonomous creatures (or robots) which direct themselves with reference to their near neighbours – for example by moving in the same direction and remaining close to your neighbors, while avoiding collisions with them. Swarming behavior can be modelled or programmed with very simple rule sets; and it has evolved independently, with characteristic variations and elaborations in many different species. In combination with other mechanisms of self-organization, notably stigmergy (introduced below), systems comprised of very limited individual components (e.g. ants, bees or termites) are capable of very sophisticated behavior in groups.

Stigmergy

'Stigmergy,' from two Greek words 'stigma' (sign) and ergos (work) is basically a mode of communication by writing on the environment. It was discovered by the French zoologist Pierre-Paul Grassé in connection with his work on termite colonies; and also used by ants, and even by bacteria which differentiate into different cell types and participate in group behaviors controlled by their own, collective chemical secretions. Building houses, tools and infrastructure, making pictures, maps and diagrams and writing texts of all kinds, our own species makes extensive use of stigmergic communication – as all our artifacts suggest their own use and carry other messages, whatever else they do. Just think of guns, for example, which suggest that you can shoot someone. Then think of buildings, streets and roads, all with signs on them. Then think of newspapers, magazines and books.

As a mechanism of self-organization, stigmergy promotes order through a social sharing of information. As each component of the system (each cell, termite or human) goes about its business, it leaves traces in its environment which are then read and responded to by others on a "to-whom-it-may-concern" basis. In effect, each entity makes its contribution to a context which all share together, thereby enabling collective learning and knowledge, and collective action too, as will be seen below.

The Matthew Effect

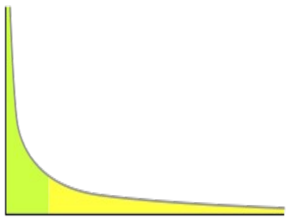

The Matthew effect is named from Matthew 25:29: "For to those who have, more will be given . . . but from those who have nothing, even what they have will be taken away." It explains why rich people tend to grow richer, and why celebrities can be famous just for being famous. Most succinctly and generally it says that whatever has attracts more).[3] One of its consequences is the Pareto principle, aka 80-20 rule, that "roughly 80% of the effects (in many areas) come from just 20% of the causes." This distribution was first noted in connection with the ownership of land in Italy, but it is valid also in such fields as income and profit distribution, occupational safety, quality control, and the fixing of software bugs. The Matthew effect and Pareto principle are represented mathematically by a 'power law' distribution in which the magnitude of an effect falls off very rapidly against the number of cases at that magnitude. Many physical, biological, and man-made phenomena roughly follow a power law curve: the size of earthquakes and avalanches, wars, criminal charges per convict, lunar craters, Google search results and the frequencies of words in most languages – for just a few examples.

Natural Selection

The mechanism that Charles Darwin called ***natural selection*** (briefly outlined here) put the notion of self-organization on the map of scientific thought. Some idea of biological evolution was already in the air in 19th century England and Europe but no one knew how it could work. Darwin and Alfred Russel Wallace proposed the first real theory of how it could.

The basic idea is simple: Variation in genetic makeup occurs at random in any population of interbreeding organisms. Most of these changes will not affect survival or will affect it adversely, but some changes may improve an individual's chances of survival and reproduction. A rabbit that runs a little faster than others may be more likely to escape from predators. Green plants that are more efficient at extracting energy from sunlight will grow faster. Ultimately, what matters is a concept called inclusive fitness – the ability of one version (*allele*) of a gene to promote the survival and reproduction of other individuals with that same version.

Opponents of evolution often find it inconceivable that natural selection acting on random genetic change could produce the complex designs that living things reveal. This incredulity may have been justified a hundred years ago, but it is out-of-date today. Natural selection has been demonstrated over-and-over again in the laboratory – in fruit fly populations and in bacterial resistance to antibiotics. Computer simulations have shown

3 See http://en.wikipedia.org/wiki/Matthew_effect_%28sociology%29

how complex anatomy (the eye is a favorite example) can evolve over time through random genetic drift. And finally, in several ways, it is known by now that natural selection is not quite so random as it first appears, and that its operations are supplemented by other mechanisms of self-organization, and by at least three effects: endosymbiosis, the Baldwin effect and sexual selection that I leave for readers to follow up on their own.

In his book *Chance and Necessity* (1971), Jacques Monod justly complained, "*Another curious aspect of the theory of evolution is that everybody thinks he understands it.*" This comment is important. Conceptually simple as it is, Darwin's mechanism of natural selection is full of subtleties, even without the other mechanisms of self-organization that I'm describing. The most important, perhaps, is that evolution does not occur in a vacuum: While the rabbit is evolving to escape the fox, the fox is also evolving to catch the rabbit. For this reason, we need to recognize that all evolution is really '*co-evolution*,' an evolution of the ecosystem as a whole. I like to speak of the whole concept of a bottom-up, self-organized world as *the ecoDarwinian paradigm*. Ecology and natural selection are really two sides of the same coin.

Self-organized Criticality

Self-organized criticality (SOC) can be defined broadly as the tendency of systemic order to maintain itself at a critical point by breaking down when it goes past that point. The effect was first demonstrated and named in 1987 by some physicists, Per Bak and his associates, with the following simple experiment:

Imagine that grains of sand are dropped at random onto the flat surface of a kitchen table. At first, when the grains land, they bounce (essentially at random), and then lie where they have fallen. In this way little piles build up, merging together into a single large pile which gets higher and steeper as the grains continue to fall, bouncing and rolling until they come to rest. The result so far is a cone of sand, almost perfectly symmetrical, building up toward some particular height and steepness (a *'criticality'*).

From this point on, sand falling at or near the top of the pile tends to trigger avalanches. The cone's steepness remains the same, but its radius at the base keeps getting wider and wider. The avalanches are of varying size, following a power-law distribution, with small avalanches exponentially more probable than large ones. Like fractals, they are scale-invariant. The steepness of the cone is a self-organized parameter that can be changed by using a different material, but not by the rate or distribution of where they are dropped. When sand does not fall randomly but is blown around by the wind, the dunes created are still limited by SOC but also has some features of wave motion.

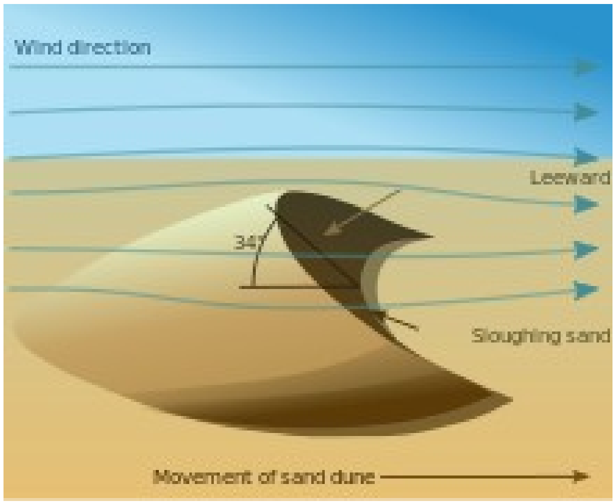

Disease epidemics, mass extinctions in the biosphere, shake-outs in the economic marketplace and many other phenomena show a similar pattern of breakdown once a critical point is reached. Even human societies and individual human identities may be limited in this fashion and organizations and institutions build up until they collapse to something simpler. In general, the suggestion of SOC is that order should be seen as a balancing act against the second law of thermodynamics – the tendency toward homogeneity and chaos. All organization is fragile.

Top-down Causation

In sum, all of these mechanisms can be seen as instances of '*top-down causation*,' the relationship that obtains when some event or structure is explained with reference to the whole of which it is a part. Died-in-the-wool reductionists insist that there is no such thing – that global systemic behavior can only be explained with reference to the activities of their component parts, yet there are many cases where the reverse seems true – where the parts and their behaviors are explained only with reference to the whole system in which they exist and function. For example:

- The periodic motion of a pendulum bob depends on the geometry of whole system – the length of its string and (somewhat) on the angle of displacement – but not at all on the mass of the bob.
- A wheel rolling down hill stays balanced on its narrow rim, and does not fall to the more stable, horizontal position until comes to a stop. The motion of a given atom in that wheel is constrained by that of the wheel as a whole.
- The magnetic properties of a compass needle is again a top-down collective effect, resulting from the coherent alignment of all the

magnetic fields of its individual atoms.

- In the swarming behavior of birds, fish, wolves and many other creatures, the path of each individual seems constrained and guided by that of the group as a whole – though the reverse is also true.
- When two young lovers agree to get married, their simple jointly-held decision triggers a flurry of activity in each, and in their friends and families.

In each of these examples, the reductive strategy fails because a whole is something more than all its parts together. It includes too the relationships among those parts, and the context that they produce together – a logical abstraction *with causal powers!*

Our individual human choices are the prototypical examples of top-down cause in ordinary experience, leading us to associate such causes with ideas of volition, intention and purpose. Yet top-down cause is not an effect of conscious purpose but of context, as we see from the examples of pendulum, rolling wheel and magnet, which are purely physical, without a trace of intentionality. In a much more complex way, human intentionality itself can now be seen as a top-down effect, resulting either from the whole pattern of neural firing, or from an alignment of thoughts (cognitive events) – depending on a choice of viewpoint, which is itself a causal context.

*** * * * ***

A fully general science of self-organization does not exist at this time; but it is already a good bet that Nature and the universe, viewed as a whole, is a single self-organized phenomenon, revealing self-similar geometry and logic in every area and at every scale that we can study. If and when we have a general science of biology, studying the evolution of life anywhere possible in the universe, if and when we have a true science of mind, with human identities and cultures properly understood, it's very likely that SOC, natural selection, and the other mechanisms of self-organization will be central features of such disciplines. It's likely too, at the foundations of these subjects, that concepts of *coevolution*, *context*, *ecology* and *emergence* will replace the evolution of individual patterns at the focus of interest; and it is to such holistic concepts that we now turn.

3. Life

Given the mechanisms of self-organization that we've just reviewed, we begin to understand how life and mind are possible in a purely physical universe – without qi (*ki*), *prana*, *élan vital*, or supernatural substances of any kind. Such concepts may still find use as figures of speech – as they do in aikido, t'ai chi ch'uan and yoga – but they have no standing as meta-physical realities. In place of vitalist and dualist ideas, it is now becoming possible to understand life as a physical phenomenon, with *minding* as a crucial function that all living entities perform, at least to some primitive extent. Still, we want to distinguish physical systems that are alive, from many other systems that are not. How should we draw this distinction?

A Definition of Life

There is probably no one right way to define 'life', because our conception of it includes several functions which are mutually independent – so that a system in question may have only some of them, while lacking others. Are viruses alive? The robots of Isaac Asimov or Karel Čapek? The epsilons in *Brave New World*? These cases are not so clear. The virus is only a complex molecule until it usurps the metabolic and reproductive machinery of a host cell. The robot characters in Asimov's stories and Čapek's play can think and feel, but cannot reproduce themselves. The epsilons in Aldous Huxley's novel also cannot reproduce, and are too stupid to do much thinking or feeling. They are cloned and trained to perform repetitive tasks, and are fit for nothing else. They seem alive, but not fully so. They are basically artifacts, manufactured by stunting them as embryos.

For our purposes, and following Stuart Kauffman's treatment, we can play the definition this way: We'll say that a system is *living*, or *alive*, if it

a) collects and expends energy;
b) makes copies (replications) closely similar to itself; and
c) acts autonomously 'on its own behalf' by maintaining itself in these ways.

The fact that life collects and expends energy means that it is limited by available supplies of energy, and by the laws of thermodynamics. A living system performs thermodynamic work cycles, transferring heat and work into and out of the system, (varying pressure, temperature, and other state variables), and eventually returning to an intial state. For that reason, it must be displaced from thermodynamic equilibrium – as a swinging pendulum has been displaced. It must maintain itself in this displacement from equilibrium.

Just for this reason, our definition of life entails eventual death: Without death, life as we know it would not be possible, because no living creature can stay alive indefinitely, or becomes alive all by itself. The Urban Dictionary defines 'life' as “a sexually-transmitted, terminal disease,” and this is more than a joke because it is eventual death (along with replication) that makes evolution possible. Then too, due to the conservation energy, animate life is only possible by killing and eating other living things. Only green plants don't rely on death to make a living. Once we understand life in ecological and planetary terms, life and death appear to be inseparable.

All living species – here on Earth, and probably in the universe – exist within some ecosystem or biosphere in ecological relationship (of predation, competition or symbiosis) with other living things. For this reason, they don't so much evolve as co-evolve – with and against each other – not just through natural selection but through a very fancy type of SOC (self-organized criticality) that we take up below.

4. Mind

One defining feature of living systems is that ***they can act on their own behalf***; and this capability raises a logical distinction between *doings* and mere *happenings*. Doings, the actions of a living system, have a purposeful, intentional quality that non-living systems lack. In the world of physics and chemistry there are only events or happenings; but in the life sciences, intended *doing* has emerged somehow, and this calls for explanation.

Take this laptop, for example, the one I'm typing on as I write this. It is a fairly complex system. Things can happen to it, and it can transit from one state to another as I work with it, according to the logic of its own hardware and software. But if I took a hammer to it for some reason, or decide to replace it with a newer model, it can do nothing to defend itself. It cannot act on its own behalf, because it has no interests to pursue. It cannot act at all, because its changes of state do not result from intentions of its own. But even the humblest one-celled creatures have enough mind to sense a chemical gradient and swim in one direction rather than another.

To have intentions, to be capable of action on your own behalf, requires that you have interests and a mind, however primitive or rudimentary. It makes no sense to attribute interests to a stone, or to a drop of water, or even to a laptop, but we can attribute agency and purpose to the simplest living things, trying to stay alive and reproduce as they evolved to do. They may not be *aware* of their interests. Even humans do not always know what is good for them. But they have enough mind to form intentions and to act upon them.

Kinds of Minds

In his book, *Kinds of Minds* (1998), Daniel Dennett suggested the following classification of relative sophistication for the minds in nature:

- The most basic, *Darwinian minds* learn only through natural selection, by paying with their lives for their mistakes. One-celled creatures, beetles and plants are the most obvious examples. Plants, for example, do not have nervous systems or any sort of consciousness, but they do have rudimentary minds including roots which grow toward moisture and leaves which turn toward the sun. Time-lapse photography shows how they distinctions that matter to them and act on their own behalf.

- Skinnerian minds, so-called because they have been extensively studied by behavioral psychologists like B.F. Skinner, can learn by trial-and-error from individual experience. We can meaningfully attribute skills to such creatures as their activities become more successful or more efficient with practice. Pigeons are one example. These birds can remember which behaviors got rewarded or punished in a given situation, and can modify their choices accordingly. What they conspicuously lack is any ability to imagine counter-factual alternatives.

- Popperian minds (named for Karl Popper, the philosopher of science) are much smarter, because they can form *ideas* (virtual images) of how their situation might be better or worse or merely different than it actually is. They can progressively refine such ideas, adhering to those that are confirmed by experience, and abandoning those that fail. Like primitive scientists, they frame and test hypotheses (i.e. speculative conjectures) about the world. Most mammals do this, at least to some extent. Our hominid cousins (orangutans, gorillas and chimps) are very good at it. But human babies are much better – and their parents get to watch, in real time, how childish conjectures converge toward something like reality.

- Gregorian minds – named after Richard Gregory, one of the founders of cognitive psychology – offload substantial portions of adaptive intelligence and knowledge onto their environments. This is how older children and adults function in the world. Our tools suggest and remind us of how we can use them. Our houses, streets and towns set up constrained environments in which some behaviors are suggested while others are either impossible, or strongly contra-indicated. Pictures, diagrams and texts present ideas about the world – summing (in libraries, and now on the Web) into a collective knowledge environment more vast than any individual could acquire.

- Finally then, it makes sense to think of human families, business firms, church communities, nation-states and some other institutions as collective minds, composed of numerous individual minds that interact and write on their environment. Many human group minds of this kind must be considered insane – dysfunctional and strongly divided against themselves – but, to the extent that they are capable of any collective agential behavior, they are minds nonetheless. Anthills, termite colonies and beehives are examples of collective mind – much studied and now well-understood. (Human organizations are much harder to study because more difficult to experiment with. Also, their human components can and do lead

> private lives – outside of, different from, and sometimes in contradiction to their public ones.) Even human individuals can be seen and studied as collective minds, comprised by the interactions of our neurons and other body cells.

These different types of entity are recognized as minds because they act on their own behalf. At different levels of complexity and sophistication, they are autonomous agents, formed by natural selection and the other mechanisms reviewed above. Human intelligence, extraordinary as it is, is only a remarkable extension and specialization of intelligences that preceded us. Without Darwinian, Skinnerian and Popperian minds, our Gregorian and collective ones could never have evolved.

The Self-Organization of Mind

A volume of Gregory Bateson's collected papers bears the title, *Steps to an Ecology of Mind*.[4] With this title, Bateson suggested the paradigm of mind as a self-organizing ecosystem. There's really no such *thing* as 'a mind.' What we observe is a *minding process*, performed within and amongst its substrate brains. What we call a human mind is better understood as a self-organized *identity*, imagined, by crude analogy, as software running on the 'wetware' of a human organism. My purpose in this appendix, has been to make it plausible that the abstract patterns of our minding are wrought by the same mechanisms of self-organization which shape the patterns in physical matter. When neuroscience addresses "the hard problem" of conscious mind, that's where it is going.

Self-similarity

Let's begin with the fact all learning can only stem from, and base itself upon mindset already in place. We learn by extending, refining and sometimes changing drastically the patterns of mindset that we already hold. As our attention, values and assumptions are always based on these existing patterns, the outcome of learning is a metaphorical spiral, analogous to that spiriform nautilus shell. Identity today is always a product of identity yesterday, and a template for identity tomorrow.

For this reason, identity is rather stable most of the time. We cycle round and round in our routine arrangements, repeating these with only minor gains and losses from one round to the next. Only at tipping points do we see drastic, discontinuous changes.

Lobster Trap Effects

Like the lobster trap, most of the identity-forming games that people play are easier to sign up for than to drop out of. Joining the military or the mafia are obvious examples, but there are many others. Nearly all children

4 Available on the Web in full here.

must do the best they can with the parents and siblings they've got. Despite possibilities of abortion, pregnancy by default is a one-way street to parenthood. Living in a certain neighborhood, going to a certain school, having certain friends (but not others) will be life-determining, for most people.

Lobster traps need not work 100% of the time. If most of the lobsters who go in stay in, that is more than good enough to ensure a reliable supply, whether for the market or just for dinner. Similarly, human contracts and relationships can be amply stable without always being escape proof or permanent. In the government and many private companies, it takes much longer to fire an unsatisfactory employee than to hire a new one. Marriages may end in divorce, but they take longer to get out of than into; and some relationship between former spouses often persists, especially when there are children involved. Doctors, lawyers and other highly trained professionals tend to remain in their fields for life, given the costs of becoming qualified in the first place. True that professionals may lose their licenses, or burn out, but most of them stay with it.

Natural Selection

Much of our learning works by trial-and-error, and therefore by a kind of natural selection. Roughly analogous to “survival of the fittest” in a natural ecosystem, the behaviors that work well for us become more frequent, while behaviors that fail or get punished become less so. Also, behaviors that work well for one person, especially for a prestigious person, tend to be copied by others but then extinguished if they don't work.

As with biological selection, there is a Baldwin effect in the sphere of mind if we have some choice in the games we play. In both cases, by choosing our games we also choose the selection criteria that work upon us. As in the natural world, this will lead to differentiation and specialization of identities – all the more so when possibilities of cooperation and reciprocity encourage concentration in areas where we have most to offer – that is, where some comparative advantage already exists.

Matthew Effects

The rich get richer. Celebrities are famous for being famous. Websites that attract a lot of traffic, appear earlier in the search results, and tend to attract even more. But Matthew effects work at the level of identity as well. Work flows toward people who are already busy. People who made friends easily as children acquire good 'people skills' and tend to make friends more easily as adults. The evidence that you notice (and don't notice) tends to confirm the beliefs you already have. In the development of identity, Matthew effects combine with natural selection and with the principle of comparative advantage to ensure that we practice harder and keep getting better at things we are already doing well. The upshot is a trend toward increasing specialization in every area of personality.

Least Action

Economy of effort shapes the way we use our bodies, and then our habits as these patterns develop and change. In satisfaction of our needs and in search of pleasure, we seek a maximum of 'bang for the buck' in the way we spend our energy and cash. In such ways, a principle of least action shapes our lives from the beginning, and puts a premium on efficiency and low cost in the production of goods.

But the principle of least action is most interesting where it is apparently violated – in all the gratuitous activities of art, religion, sport and sexuality which cost a great deal of time, energy and money with no direct return in biological fitness. Why do we do such things? For meaning and self-expression we say, but that answer only pushes the riddle one step back: Why are meaning and the expression of identity so important to us, that we expend so much for their sake?

The answer seems to be that individual self-expression and quest for meaning are central to our nature as ultrasocial animals and crucial for the self-organization both of culture and of personal identities. We expend energy on projects of meaning and self expression by instinct, not by rational calculation of self-interest, because such projects are vital for the self-construction of our own identities and for Bateson's 'ecology of mind.' Insofar as we try to conduct such projects economically, the principle of least action is upheld. Though the projects themselves may appear wasteful, they contribute immensely to the memetic repertoires of our groups and whole societies.

Equilibrium

Much as with physiological parameters like blood pH and so forth, certain cognitive pairs like fun/comfort, work/play, dominance/submission also 'ask' to be kept in balance. Equilibrium in human identities is not necessary or automatic, but we can see a tendency to balance work with play, and discipline with self-indulgence. We often compensate for weakness in one area or relationship by asserting power in another, and vice versa – e.g. by compensating for power and dominance in the work world with domestic or sexual submissiveness.

The ancient Chinese sought a balance of *Yin* and *Yang* in every sphere of life – not less in politics and human relationships than in medicine. Greek thought too stressed the importance of moderation in human ambition and striving as in drink and diet. Both saw a principle of equilibrium as central both to human mindset and human affairs.

Self-organized Criticality

Finally, we can see self-organized criticality (SOC) in the breakdowns and conversion experiences that shift identity quite radically sometimes. The build up and breakdown of identity structures is strikingly analogous to the build-up and eventual avalanche of snow fields on a mountain.

Organizations and states expand but then prune back. On occasion they collapse altogether. Something similar seems to happen with personal identities as we try this and that, expanding our arrangements, our webs of inter-relationship, and our structures of worldview and mindset. The tendency (as with Per Bak's sandpile) is to maintain identity as a dynamic equilibrium, hovering on the edge of chaotic breakdown.

As the *Tao Te Ching* warns:

> Stretch a bow to the very full,
> And you will wish you had stopped in time;
> Temper a sword-edge to its very sharpest,
> And you will find it soon grows dull.
> When bronze and jade fill your hall.
> It can no longer be guarded.
> Wealth and place breed insolence.
> That brings ruin in its train.
> When your work is done, then withdraw!
> Such is Heaven's Way.[5]

Beyond a certain critical point, a system becomes unstable, and some compensation, cut back or breakdown is only to be expected.

5. Coevolution and Emergence

We turn now to the biosphere that emerges when self-organization has produced a diverse population of living things. At the beginning, perhaps, there may be only one form of life, drawing its energy from sunlight, and/or from the residual heat of the core of a cooling planet. But in due course, diversity and specialization occur by essentially the same logic that drives the specialization of economic systems – a diversity of local conditions and resources, and the opportunity to gain competitive advantage through mutual interdependency and the perfection of certain capabilities at the expense of others.

At this point, we're no longer speaking of the evolution of individual life forms, of separate species adaptating to their respective environments , but of the coevolution of this whole 'ecosystem' – comprised not just by its various species, but by their web of interrelationships, by the niches that they afford for one another. And we see at once that full understanding of such a system will require explanation of its dynamics both in relation to itself (its coherence, heath and prospects), and also to all relevant abiotic factors – the rhythms of day and year, the input of solar energy and energy from the planetary core, the bombardment by meteorites, the geographical variations of climate, and so on. Such a system may be perturbed, disrupted, in any number of ways, and its health cannot be taken for granted.

5 *Tao Te Ching* #9, Arthur Waley translation.

Systemic Health

In fact, the health of Earth's biosphere today has become an urgent political issue as well as a scientific one, and it's bound up with some huge doubts about the health of current human society – itself a kind of ecosystem as we are learning. To think clearly about such matters, requires that we understand the concept of systemic *health* much better than at present. As applied to the Earth's biosphere and its global human society, the definition of this term by the World Health Organization (WHO) is both vague and circular.[6] Let us see if we can do better.

I have in mind the memory of my daughter in a playground, when she was just over a year old, but already walking very well. I would sit on a bench and watch her as she bumped around, exploring this and that. She would go off some way and come back to me, go off in some other direction and come back; and I remember reading somewhere that this 'petal pattern' is common behavior, and that a barking dog, or anything a bit frightening will cause the child to stay closer. Too bad I can't find a reference to such behavior now, but let's assume that the pattern holds up – that the distance a child wanders from mom or dad and the time she spends away from her parent on each cruise away, are measures of her confidence.

Now, this petal phenomenon may be very general, and may measure systemic 'health' – the capability of a system to self-organize – as a function of parameters analogous to the confidence and exuberance of a healthy child. When a system is safe and healthy it behaves boldly: exploring its world, and testing its opportunities and limits. An ailing system, threatened in some way, huddles down and circles its wagons – opting for comfort and security. In fact, we can observe such tendencies not just in children but in adults and groups and whole societies. There is a dialectic, it seems, between moods of fun, expansion and exploration, against moods of safety-seeking and self-comforting. Epochs of expansion and exploration tend to encounter threats. Successful defense will bring a dividend of confidence – with opportunities, perhaps, for further exploration and expansion. In this way a system may exist and grow, alternating between these poles until destroyed either from within by its own deterioration, or from without by a threat it cannot meet.

Suggested here is that the same dialectic of exploration and self-protection applies to the whole biosphere as it does with any self-organizing system. When conditions allow, such a system explores boldly, expanding into what Kauffman calls “the adjacent possible” – the states immediately accessible – thereby increasing diversity (i.e. the possibilities for what can happen next). When conditions are hostile, the space of accessible states will shrink, reducing the diversity of specializations and encouraging the expansion of 'weed species' which can thrive in a wide

6 WHO defined 'health' very broadly in 1946 as “a state of complete physical, mental, and social well-being, and not merely the absence of disease or infirmity.”

range of conditions. Of, course, if the conditions become too hostile, the ecosystem will collapse and die, and idea of 'health' is just a memory. In fact, life cannot even begin when conditions are too extreme. Planet Earth is still the only place in the universe that we know for sure allowed self-organization at this level of complexity. But our astronomy has now reached the level of observing not just galaxies and stars, but the planetary systems around some stars; and the search is on for 'Goldilocks planets' in a star's habitable zone: “not too hot and not too cold, but just right.”

Emergent Order at the Edge of Chaos

From the mechanism of self-organized criticality, it is conjectured that an ecosystem tends to evolve toward the edge of chaotic breakdown. Beyond this limit, avalanches of extinction prune back the burgeoning complexity below its critical limit. Order builds up to a point where it becomes unstable, in a manner analogous to Per Bak's sand pile – where the *yang* of dispersion and dissipation overcomes the *yin* of further growth.

But how do we get from the organization of sand piles and snow slopes to the organization of minds like our own? It seems impossible or miraculous; and it is no surprise that humans have sought and lived with supernatural explanations until our own time. David Chalmers has argued that no wholly physical explanation of consciousness is possible. But, at the state of current knowledge, it's a good bet that he is wrong, though our idea of physics itself may well be stretched along the way. Terrence Deacon, a professor of anthropology and cognitive science at Berkeley has identified five levels of emergence – five levels of ordered complexity – culminating in minds like our own, emerging bottom-up through purely natural processes:[7]

- At the first level of emergence, the bottom-up processes of statistical physics lead to generic, global properties like the wetness and temperature of a tub of water. A statistical, coarse-grained perspective on the system hides information (e.g. about the motion of individual molecules) which might be retrieved in principle, but not in practice. At this level, time acquires an 'arrow,' a definite direction. System entropy tends to increase, as the time-symmetry of basic-level physics disappears.
- At a second level, boundary conditions accumulate, so that bottom-up action creates new structure. Examples would include Per Bak's sandpiles, convection patterns like those Bénard cells and gravitational structures like solar systems and galaxies. These

7 This account of Deacon's levels of emergence is based on a paper by George F R Ellis, *On the Nature of Emergent Reality*. Deacon's recent book, *Incomplete Nature: How Mind Emerged from Matter* (2012) discusses Chalmer's “hard problem” of consciousness in detail, from the idea that absence itself can have causal properties, as introduced in Section 6 below. Perhaps something can come from nothing after all!

systems are an important step in the development of biological complexity, but the purposeful, goal-seeking of living systems is not yet present. At this stage, we are in a realm of self-organized criticality, chaos theory and molecular structure, but there is no goal-seeking, information-directed activity as yet.

- At the third level, informational learning becomes possible – on an evolutionary scale over the history of the species, but not by individuals in their own lifetimes. For example, genetic learning occurs through natural selection. All plants and self-reproducing cells show emergence at this level. A degree of teleonomy is found, based on pre-determined 'instinctive' rules which may enable feedback controls and coordinated top-down responses to current conditions and disturbances. The living cell, it has been said, may be seen as "a sort of working hypothesis about what information will survive and thrive most effectively in a given environment." At this third level, purpose emerges, yet this purpose is also a kind of absence, like the emptiness of a cup into which some liquid might be poured – or better, like an empty stomach waiting for some food. A pull of unrealized possibility opens the way to biological function and to the purposeful action of individual creatures.
- At the fourth level, individual learning occurs. Feedback controls come to be directed by explicit goals ('values' if you will) based not only on the Darwinian learning of a species but on the remembered experiences of the individual creature. We observe such learning in cats and dogs, and in all mammals probably. In social animals, it may allow for significant communication, based on a transmission and response to learned signs, as in Pavlov's famous experiments.
- At level 5, we find true symbols in addition to mere signs (like the bell that Pavlov's dog learned to recognize as a sign of food). The difference between signs and symbols was elegantly discussed by Deacon in a previous book, *The Symbolic Species* (1997). Where signs point to and correlate with events in the world, symbols point to and contrast with other symbols to form complicated webs of association. With symbols we gain the possibililty of language, and other complex modelling or representation of a physical and social world. We arrive at minds like our own, with a capacity both for conscious self-expression and for social participation.

Through mechanisms of self-organization, each level can evolve and emerge from the one before, as stable loops of action grow increasingly nested, entangled and complex. Over-all, it becomes intelligible how abstractions like agency and consciousness can evolve – how 'minds' can perceive and exert causal power, without supernatural add-ons, in a purely physical world.

6. Conclusion: *Ex Nihilo Fit*

Self-organized agency, the idea that mind can emerge and supervene on ordinary matter, is profoundly counter-intuitive. It seems so obvious that nothing can come from nothing. (“*Ex nihilo nihil fit,” as* Parmenides and Lucretius said.) But that ancient maxim is turning out to be wrong. We know today that a whole complex cosmos can emerge from nothing, and we begin to understand how this can happen. Vacuum energy is just about the hottest topic in physics today; and given the vacuum energy and sufficient time (going on 14 billion years, as we estimate), it seems that a complex, richly structured universe can emerge – including creatures with minds.

Thoughts are pure abstractions, but they have physical effects. An intention can throw a stone – or build a spacecraft and send it to the moon. In a recent book,[8] Deacon directly tackled the central meta-physical puzzle of self-organization: How can we think of pure abstractions as having physical consequences and causes? How can thoughts cause things to happen, and (conversely) be caused themselves by physical events? What makes the difference between a physical *happening* and the purposeful *doing* of a living creature?

The answer he proposes is that all thoughts and intentions can be seen as physical *constraints* – as much so as Einstein's constraint on the speed of light. As the desk constrains my coffee mug not to fall to the floor, as it keeps my arms on its surface when I rest them on it, so my intention to throw that mug constrains my arm to move and send it flying through the air. Like intentions, perceptions, feelings and every kind of thought are also constraints of a sort: constraints on what my brain is doing, and on how it will constrain itself next. When the nurse sticks a needle in my arm I cannot *not* feel the pain. Feeling it, I cannot *not* respond somehow: like a young child by crying and jerking away, or like an adult by willing my arm to hold still.

What enables living creatures to act as agents on their own behalf is the legacy of intricately nested constraints (constraints on constraints), evolved in the course of time. Let me spell this out a little. The first move is to see intention or purpose as a form of constraint. Buridan's poor donkey, starving to death between two equally attractive bales of hay, is in a symmetrical state of indecision. Whatever breaks that symmetry (however subtle or random) in effect constrains him to move toward one pile rather than the other. Any law or constraint on any physical system whatsoever can be seen abstractly as a breaking of *a priori* symmetry. Unconstrained, it might do anything or nothing. Given the constraint, we know that it will act in some particular way.

8 *Incomplete Nature: How Mind Emerged from Matter* (2012)

But a physical law or constraint is not a 'something.' The second move is to understand that information is not a physical event or 'thing,' though it is always carried by some physical signal. Deacon's central idea is that all such abstractions are better seen as patterns of emptiness, analogous to the empty space that makes a cup or a house useful, as the *Tao Te Ching* said long ago):

> We put thirty spokes together and call it a wheel;
> But it is on the space where there is nothing
> That the usefulness of the wheel depends.
> We turn clay to make a vessel;
> But it is on the space where there is nothing
> That the usefulness of the vessel depends.
> We pierce doors and windows to make a house;
> And it is on these spaces where there is nothing
> That the usefulness of the house depends.
> Therefore just as we take advantage of what is,
> We should recognize the usefulness of what is not.[9]

In Deacon's view, our beliefs, desires and intentions are seen most clearly as constraints, breaking the symmetry of what we can think, feel and intend, and thus of what we can do. Consciousness itself is a constraint, a breaking of *a priori* symmetry on attention and/or subvocal speech. Matter and energy are needed to impose or transmit constraints; but they are not the constraints themselves.

To make sense of agency, we must recognize that intention is a type of cause, different from other (more directly physical) causes in at least two ways: first, because its force stems from an absence *now* of something needed or desired; and second, because it anticipates an imagined future and is therefore (in a sense) directed *backwards* in time. The present intention could not exist if that absent future had not already become imaginable.

A third crucial point is that this backward-directed, teleological causality – does not require conscious purpose, nor anything else that we would recognize as mind. The humble thermostat already has some *telos*, endowed by its human designers – to turn a thermostat on and off, and so maintain its room at a calibrated temperature.

But any living cell or multi-cellular organism already has a mind of sorts. As an agent it must gather and expend energy, and keep itself from being eaten by other agents, who are also in search of energy. Also it must replicate. To do these things it will have evolved some primitive 'values' – some evolved distinctions between '*yuck*' and '*yum*' (as Stuart Kauffman puts it), with a propensity to avoid the first and search out the second. Consciousness, which evolves much later, can then be seen as the self-referential property of an advanced brain: its capability to monitor, pay

9 *Tao Te Ching* #11, Arthur Waley translation.

attention to, and purposefully direct its own faculties of attention.

Deacon speaks of *absential* phenomena, and compares the new recognition of their causal power to the introduction of a concept of zero in mathematics. Though the idea of zero as a valid number seemed scandalous at first, it proved so useful and made possible so many other advances that it is now taken for granted and seems completely obvious. Similarly, Deacon argues, the introduction of a "placeholder concept" of absential cause will prove immensely useful in bridging the conceptual gap between mind and matter.

* * * * *

Finally, like the glass that is either half-empty or half-full, what one sees in all this is largely a matter of temperament. Kauffman contemplates self-organization and sees a naturalized idea of God, much as Spinoza did before him. Others think more as I do, that the idea of God, even as a handy name for the cosmic processes of self-organization, raises more problems than it resolves and might be better forgotten. That word still tempts people to think of a top-down source in all the order that we see around us, while the central point of modern thought has been that order bubbles up out of the void, through the subtle interplay between chance and necessity that Jacques Monod described. But I too have written a long essay called Theology For Atheists, seeing God-mindedness and religion as necessary, and sometimes useful features of human thought.

The take-away point from all this, the point I hope that this appendix leaves with you, is that we live at a time when three tremendously influential worldviews are finding themselves challenged, and even supplanted by a fourth, still-unfamiliar conception based on ideas of self-organization like those developed here. All three of the more traditional worldviews are put in question: The possibility of self-organization suggests to us that the world we know – the physical universe, the biosphere, our human society – were not ordained by a providential God, and were not by-products of a mechanistic, predictable, Newtonian world. It suggests too that our world is not so barren or disenchanted as modernism suggested to thinkers like Friedrich Nietzsche, Max Weber and Sigmund Freud. Though the traditional (and non-traditional) religions, reductionist science and various forms of nihilism still have plenty of adherents, a more edifying and pragmatically useful worldview is now available at the cost of some imagination and intellectual work.

Unfortunately, that required effort is considerable; and for this reason, it is impossible to say how pervasive this new dispensation will become in human affairs, and what its long-term effects will be. I wish it were easier to hope that in time an ecoDarwian paradigm of self-organzation and co-evolution will either supplant or acceptably complement the older world-views I 've mentioned, and that the culture wars of this epoch will abate.