Complexity Theory

Michael Strevens

For the *Oxford Handbook of the Philosophy of Science*,
edited by Paul Humphreys

Almost everything is a complex system: Manhattan at rush hour of course, but also, if you know how to look, a rock sitting in the middle of a field. Excited by the heat of the midday sun, the molecules that make up the rock are vibrating madly. Each is pulling at or shoving its neighbors, its parts shifting around its center of mass in a most haphazard way, their next move hinging on a multitude of minute details concerning the many atoms making up the surrounding stone.

The rock and the city—what do they share that makes each in its own way a paradigm of complexity? They are composed of many parts behaving somewhat independently yet interacting strongly. That I take to be the essential recipe for a complex system: sufficiently many parts, independent yet interacting.

Complexity's most salient consequence is intractability. Even supposing that the behavior of every part of a complex system, and thus of the whole, is entirely determined by the exact state of the parts and the fundamental laws of nature, there is little hope of building scientific models capable of representing the system at this fineness of grain, thus little hope of predicting and understanding the behavior of complex systems by tabulating the gyrations of their parts.

You might therefore wonder whether sciences of complex systems are possible. Complexity theory in its broadest sense is the body of work in science, mathematics, and philosophy that aims to provide an affirmative answer: to show how investigators inquire fruitfully into the workings of complex systems and to understand why they so often succeed.

And they do succeed. Our present-day science does not cover everything, but it covers a lot, and a lot of what's covered is complex: rocks, gases, organisms, ecosystems, economies, societies, languages, and minds. Evidently it is possible sometimes to predict, sometimes to explain, sometimes to predict and explain what is going on in a system without having to track the individual histories of its many parts.

Complexity theory looks both retrospectively at the successes, attempting to understand how they were achieved, and prospectively at possible future successes, attempting to formulate ideas that will engender the next generation of discoveries in developmental genetics, economics and sociology, and neuroscience.

This article will not discuss two important topics: attempts to define and quantify complexity, and the concept of complexity in computer science and mathematics. The latter I take to be outside my remit; as for the former, I am not persuaded that it is a fruitful way to begin the study of complex systems. I will therefore follow Weaver (1948) in relying on a rough and ready characterization—complexity as a consequence of numerous independent interacting parts—and a long list of examples from physics, biology, and the social and behavioral sciences. Definitions are left to complexity theory's seventh day.[1]

1. Two Grand Questions

Complex systems exhibit, on a grand scale, two kinds of behavior that call out for explanation: simplicity and sophistication. Sometimes it's one; sometimes

1. For a good overview of attempts to define complexity, see Mitchell (2009), chap. 7. Mitchell also mentions a survey suggesting that many complexity theorists believe that it is too soon to attempt such a definition (p. 299).

it's the other; oftentimes, a mix of both at different levels.

I call a behavior simple if it conforms at least approximately to a generalization that can be simply expressed—if, for example, it obeys a "law" containing just a few variables, as most gases roughly obey the ideal gas law $PV = kT$, or most organisms obey Kleiber's law according to which metabolic rate increases in proportion to the 3/4 power of body mass.

The simplicity in a complex system typically does not exist at the level of the many interacting parts that make for its complexity—at what I will call the system's *microlevel*. The reason is this: to track a complex system at the microlevel, you need sets of equations that for each small part describe its state and prescribe its interactions with the other parts. There is nothing at all simple about such a model. Simplicity in complex systems emerges rather at higher levels of description—at a macrolevel where variables or other terms in the simple description represent aggregate properties of many or all of the microlevel parts, as a gas's temperature represents its molecules' average kinetic energy or an organism's mass represents the total mass of all its cells and other physiological constituents.

The first grand question about complex systems, then, is of how macrosimplicity emerges (when it does) from microcomplexity, of how simplicity in the behavior of the system as a whole is sustained and explained by the convoluted action of the system's fundamental-level constituents.

Consider, for example, gases. Their low-level convolution is typically for all practical purposes intractable. We do not have the mental fortitude or the computational power to track the behavior of every one of even a small quantity of gas's heptillions of molecules, each colliding with the others billions of times a second. If that were what it took to build a science of gases, it would be beyond our reach.

Happily, there is another possibility. At the macrolevel, the behavior of the same quantity of gas can be characterized by a simple linear function of three variables—the ideal gas equation. The equation is no good for predicting the trajectories of individual molecules, but we are not much interested in that. Predicting changes in temperature, pressure and volume is good enough for many purposes—good enough to provide a starting point for a science of gases that rises above the intractability of the behavior of the gases' many parts.

The simple behavior of gases and many other complex systems is not only good for complex-system scientists, but essential for the rest of us. Gases' macrolevel simplicity means macrolevel stability: the air in a room tends to remain uniformly distributed throughout the available space, rather than surging to and fro from corner to corner.[2] So we can all keep breathing with a minimum of fuss. Numerous other stable behaviors of our environment—both its resources and its inhabitants—also play a critical role in making our continued existence possible. Macrolevel simplicity in the things around us is not sufficient for life, but it is essential.

Many complex systems behave in striking ways that go far beyond the simplicity that provides a stable background for life to flourish: they exhibit what I will call *sophistication*. That label embraces many things: the orchestration of biological development, in which an elaborately structured, highly functional organism is built from a single cell; the organism's subsequent exploitation of its habitat; the operations of the minds of intelligent organisms directing the exploitation; and much more. Frequently, the sophisticated behavior of such systems is evident in the kind of plasticity that we call goal-directed. But there are other varieties of sophistication, to be considered in due course.

The second grand question about complex systems, then, is how sophisticated behavior emerges from the interaction of relatively simplistic parts. Both grand questions are posed by the appearance of something at the macrolevel that appears to belie what is at the microlevel: simplicity or stability from

2. Simplicity in the sense of compact describability does not guarantee stability—as chaos theory shows, simple laws can generate behavior that looks to us to be highly erratic—but the two go together often enough (and wherever there is stability, there is more or less by definition simplicity).

the convoluted interaction of parts; sophistication from the interaction of simplistic or simply behaving parts.[3] The emergence of simplicity makes complex-system science possible, by creating behaviors that scientists can reasonably hope to capture in a generalization or a model. The emergence of sophistication makes complex-system science interesting; indeed, essential.

The two varieties of emergence may, and often do, coexist. Inside a nerve cell's axon, a great horde of molecules bounces this way and that, the trajectory of each irregular enough to seem utterly random. But the movements of these molecules taken as a whole are regular enough that the cell exhibits a fairly simple, fairly stable behavior: upon receiving the right stimulus, a wave of electric charge runs up the axon. Examined at the microlevel, this wave consists in the individually unpredictable movements of vast numbers of charge-carrying ions; in the aggregate, however, these movements add up to the reliable, predictable pulses by which information is transmitted long distances through the body and brain. The building blocks of thought, in particular, are in large part these simple behaviors that emerge from the convoluted molecular chaos of the cell.

Yet thought itself is not simple; it is sophisticated. Somehow, the relatively simplistic mechanics of individual neurons are, when harnessed together in a brain, capable of producing behavior that is purposive, intelligent, even rational. From the simple but numerous interactions between molecules, then, comes the convoluted behavior of axonal ions; from this convolution comes a higher-level simplicity at the cellular level; and from the interactions between neurons behaving relatively simply, something extraordinarily sophisticated from the assembly that is the brain.

---

3. The second question concerns a kind of complexity emerging from a kind of simplicity; the first question, a slightly different kind of simplicity emerging from a completely different kind of complexity. I keep the two complexities well apart, hence my different names for the two kinds of complex behavior: *sophistication* and *convolution*. Simplicity in the sense of a dynamics' compact describability is not quite the same thing as simplicity in the sense of a dynamics' lack of sophistication, but the two are close enough that it seems unnecessarily pedantic to encode the difference in separate terms—hence I use *simplicity* for both.

The same complementary arrangement of simplicity from convolution and sophistication from simplicity is found in the balance of life in a mature ecosystem. Individual organisms, especially animals, interacting with each other and with their environments can be as haphazard and unpredictable in their movements as colliding molecules. The intersecting life trajectories of such organisms, then, make for a highly convoluted whole.

Within such a whole, however, simple high-level patterns emerge. The success of ecological modeling using equations tracking only aggregate populations suggests that the convolutions in individual lives add up to stable rates of reproduction, predation, and death. While it may be effectively impossible to predict whether an individual hare will get eaten in the course of a month, it is relatively easy to predict the rate at which hares in general will get themselves eaten: it seems to depend on only a few high-level variables, such as the number of predators in the ecosystem. From convolution, simplicity emerges.

This simplicity in turn gives rise to sophistication: the regularity in rates of birth and death is what makes for stable evolutionary fitness in a trait—for the fact that possession of a trait can result in a determinate increase in the rate of reproduction, in viability, and so on—and thus makes it possible for natural selection to operate with a certain consistency over time periods long enough to result in evolution and adaptation. The adaptedness of things is of course a kind of sophistication, surely the most important kind that we have yet explained.

This essay is organized around the two grand questions, focusing on attempts to provide answers to those questions of the widest possible scope, that is "theories of complexity" that attempt to give very general conditions for simplicity's emergence from convolution and sophistication's emergence from simplicity.

Notions such as explanation and emergence can become themselves a topic of debate in these endeavors, but I have put such arguments to one side, as they are covered by separate entries in this handbook.

2. From Convolution to Simplicity

In a convoluted complex system it is almost impossible to model, and so to predict, the movements (or other changes in state) of one of the system's parts—the trajectory of a gas molecule, a lynx's success in hunting rabbits, a soldier's fate on the field of battle. In some cases the chaos is replicated at the macrolevel: might the outcome of the Battle of Waterloo have hinged on any of a number of small-scale deployments or split-second decisions? But often enough, the macrolevel is a sea of mathematical calm: somehow the aggregate properties of the parts are stable or predictable even though the parts themselves are not. How can that be?

2.1 The Sources of Convolution

Let me begin by saying something about the sources of microlevel convolution or unpredictability. I see three:

1. *Convolution of the parts:* In some cases, the internals of the individual parts of a complex system are convoluted, and so any given part's behavior is difficult to predict.

2. *Chaos:* The dynamics of the strong interactions between parts is often sensitive to initial conditions (that is, in a loose sense "chaotic"). Small differences in the states of individual parts can make for big differences in the outcomes of their encounters. These big differences then stand to make still bigger differences down the road.

3. *Combinatorial escalation:* There are many parts in a single system. To keep track of these parts is difficult enough under any circumstances; if their individual behaviors are hard to predict (convolution of the parts) or their interactions depend on small details of their states (chaos), then the complexity of the task is multiplied beyond practical feasibility.

The emergence of some kinds of macrolevel simplicity under these conditions is easy to understand. The weight of a gas at a given time is a simple function of the weight of its constituent molecules; the strong interactions between these molecules, and the ensuing chaos and combinatorial escalation, is simply irrelevant to the determination of weight. (This example is a useful reminder that any complex system has many macrolevel properties. Some may behave simply; some more complexly. Of the simple behaviors, some may be easily explained; some not.)

Most macrolevel properties, however, are not like weight. A confined gas's pressure is constituted by the drumming of its many molecules on its container's walls; the force exerted by the molecules depends on their position and velocity, which is determined in turn by the other molecules with which they collide, a process that can be characterized fully only by a stupendous array of equations. The changes in the population of some organism depend on individual births and deaths, events likewise dictated by a microlevel dynamics so complex that it will never be written down.

In these cases, unlike the case of weight, the microlevel dynamics drives most or all of the changes in the macrolevel properties, yet the resulting macrolevel dynamics is as simple as the microlevel dynamics is convoluted. The convolution of the microlevel is suppressed or dissolved without in any way undercutting the microlevel's causal role. The problem is to understand this suppression, this dissolution.

2.2 *Modularity and Near-Decomposability*

Some complex systems have a hierarchical or modular structure, Simon (1996) influentially argues: small groups of parts make up larger assemblies—"modules", if you like—from which the whole is built (or from which higher-level modules are built). The parts of a module interact strongly, as they must if the system is to be considered complex at all, but because of the modular structure these interactions are felt only weakly if at all outside the module. In a completely "decomposable" system, they are not felt at all. The effects of a module part do not pass through the module's boundaries; as a consequence—since a module consists of nothing but its parts—each module is utterly self-contained, a little universe doing its own thing independently of the others. Complete causal independence of this sort is rare, and is in any case hardly a hallmark of complexity. Far more common, Simon observes, is what he calls "near-decomposability": interactions between modules are far weaker in some sense than interactions within modules.

Such near-decomposability can ensure a relatively simple, relatively stable, relatively tractable behavior of the ensemble of modules—of the system as a whole. Simon makes the case, in particular, for systems in which the interaction between modules happens on a much longer time scale than the interaction within modules. The effect of a particular module part, difficult to predict perhaps because of its internal convolutions or because of its strong interactions with the other parts of the same module, will percolate only very slowly from one module to another; consequently, he argued, the effect of one module on another will be dictated not by individual effects but by their long-run average. If this average is stable or has a simple dynamics, then the same should be true for the dynamics of intermodular interactions—in spite of the large number of strongly interacting parts of which the modules are composed. The relatively simple dynamics of intermodular interaction, Simon believed, provided a basis for stability in the system's overall macrolevel behavior—a stability that could be exploited by selection, natural or otherwise, to create sophistication on top of simplicity.[4]

Near-decomposability and the related notion of modularity, though they have remained important in the study of evolvability (section 2.5) and in evolutionary developmental biology (section 3.3), are not enough in themselves to

4. See Strevens (2005) for a more detailed account of the decomposition strategy, and Bechtel and Richardson (1993) for "decomposition" as a research strategy for understanding complex systems.

explain the emergence of macrolevel simplicity from microlevel convolution. First, such emergence occurs even in systems that have no hierarchical organization to confine the influence of the parts to a single locale; examples include gases, many ecosystems, and some social structures. Second, the appeal to near-decomposability explains macrolevel simplicity only if the behavior of individual modules is, in the aggregate or in the long run, simple. But the parts of many modules are sufficiently numerous to produce convolution even within the module; how, then, is the necessary modular simplicity to be explained?

2.3 The Statistical Approach

The first great formal theories of complex systems were those of statistical physics: first, kinetic theory (Maxwell 1860, 1867; Boltzmann 1964), and then the more general apparatus of statistical mechanics (Gibbs 1902; Tolman 1938). For many complexity theorists writing in the wake of the development of statistical physics, it was natural to apply the same statistical methods to ecosystems (Lotka 1925), to "disorganized" complex systems in general (Weaver 1948), and even to an imagined social science capable of predicting the future history of the galaxy (Asimov's *Foundation* trilogy).

Consider again a gas in a box. The statistical approach, in the version represented by kinetic theory, stipulates a physical probability distribution over the position and velocity of each molecule in the gas. In the simplest case, the distribution is the Maxwell-Boltzmann distribution, which specifies that a molecule is equally likely to be found in any part of the box while imposing a Gaussian (normal) distribution over the components of the molecule's velocity that depends only on the temperature of the gas and the mass of the molecule.

The distributions for different molecules are stochastically independent, from which certain conclusions follow immediately by way of the law of large numbers (for the same reason that, from the probability of one-half that a tossed coin lands heads, it follows that a large number of coin tosses will, with very high probability, produce about one-half heads). The distribution over position implies that the gas is at any time almost certainly distributed evenly throughout the box. The distribution over velocity implies that, if the gas is warmed, the pressure it exerts on the container walls will increase proportionately (Gay-Lussac's law): an increase in temperature results in a proportional increase in average molecular velocity, which almost certainly results in a proportional increase in the force exerted on average against the container walls—that is, an increase in pressure. The probability distributions over microlevel properties, then—over the positions and velocities of individual molecules—can be used to explain stabilities or simplicities in a gas's macrolevel properties, such as its temperature and pressure.

A direct generalization of the statistical approach in kinetic theory explains the simple behavior of convoluted systems in three steps:

1. Probability distributions are placed over the relevant behaviors of a complex system's parts: over the positions of gas molecules, over the deaths of organisms, over decisions to vote for a certain electoral candidate, and so on.

2. The distributions are combined, in accordance with the law of large numbers, to derive a probability distribution over the behavior of aggregate properties of the parts: the distribution of a gas; the death rate of a certain kind of organism; the results of an election; and so on.

3. From this distribution, a simple macrolevel dynamics or prediction is derived.

This schematic approach is what Strevens (2003, 2005) calls *enion probability analysis*, or EPA for short.

Let me elaborate on the assumptions and machinery of EPA. First, the parts of a system over whose states or behaviors the probability distributions range are called *enions*; this distinguishes them from other parts of the system (say, the distribution of vegetation in or the topography of an ecosystem) that are excluded from the analysis—their states either being held fixed or having their variation determined exogenously.

Second, the outcomes over which the enion probability distributions range are those that contribute to the macrolevel aggregates that are under investigation. If you are investigating death rates, you need probability distributions over death. If you are investigating the spatial distribution of a gas, you need probability distributions over molecules' positions. By the same token, no probability distribution need be imposed over outcomes that are not relevant to the aggregates.

Third, the enion probability distributions—the probabilities assigned to enion states or behaviors, such as position or death—should depend only on macrolevel properties of the system. The probability of a hare's death, for example, should depend only on the total number of lynxes in the local habitat, and not on the positions of particular lynxes. The reason is this: any variable on which the probabilities depend will tend to find its way into the probability distributions over the macrolevel properties in step (2). If the probability of a certain hare's death depends, for example, on the positions of particular lynxes, then the death rate as a whole will, if mathematically derived from this distribution, in most circumstances depend on the positions of particular lynxes. But then the macrolevel dynamics derived in step (3) will depend on these microlevel variables, and so will not be a simple dynamics.[5]

Fourth, the enion probability distributions should be stochastically independent. One hare's death by predation should, for example, be left unchanged by conditionalizing on another hare's death by predation, just as one coin toss's landing heads makes no difference to the probability that the next toss also lands heads. It is this assumption that allows you to pass, by way of the

5. It is not inevitable that microlevel variables will end up in the macrolevel dynamics: there might be some further mathematical technique by which they can be removed, or perhaps aggregated (so that the death rate depends only on the distribution of lynxes). The same goes for all the strictures in this discussion: any general problem might have a tailor-made solution in some particular case.

law of large numbers, from a probability to a matching long-run frequency—inferring from (say) a 5% probability that any particular hare is eaten in a month that there is a high probability that, over the course of a month, about 5% of a large population of hares will be eaten.

Fifth, as the immediately preceding example suggests, the use of the law of large numbers will tend to give you—provided that each mathematically distinct enion probability distribution is shared by large numbers of enions—probabilities for macrolevel states or behaviors that are close to one. From these probabilities, then, you can derive something that looks much like a definite prediction or a deterministic macrolevel dynamics, with the proviso that there is some chance of deviation. The chance is negligible in the case of a gas with its vast numbers of molecules, but rather more noticeable in the case of hares (though in the ecological case, there are many other sources of deviance; the assumption that outcomes are independent, for example, holds only approximately).

When the suppositions just enumerated hold, macrolevel simplicity emerges from microlevel convolution. The population dynamics of a lynx/hare ecosystem, for example, can be represented by a relatively simple equation; in the best case, a Lotka-Volterra equation containing little more than variables representing the populations of the two species and parameters (derived from the enion probability distributions) representing rates of reproduction, predation, and so on.

Where does all the convolution go? The population of hares depends on the births and deaths of individual hares, which depend in turn on minute details in position and configuration: whether or not a hare is eaten might hinge on just a few degrees in the angle that a particular lynx's head makes to its body at a particular time. Why is the dependence not passed up the chain?

The answer is that, in the aggregate, these dependences manifest themselves in fluctuations that cancel each other out. The outcome of a tossed coin depends sensitively on the speed with which it is spun: a little faster and it would have been tails rather than heads. These dependences push the outcomes of coin tosses this way and that, more or less at random. But precisely because of this randomness, they ultimately make little difference to the frequency of heads. There are as many "nudges" in one direction as in any other; consequently, the nudges more or less balance, leaving the frequency of heads to be determined by fixed, underlying features of the coin's material makeup.

The same holds, very broadly, in the ecosystem: an individual hare's life turns on a few small details, but for a system of many hares, these details more or less balance, leaving the rate of hare death to be determined by fixed, underlying features of the ecosystem: hare camouflage, lynx eyesight, ground cover, and of course the overall number of lynxes in the environs. The system's behavior can therefore be characterized by equations representing, explicitly or implicitly, this fixed background and the handful of variable properties.

Enion probability analysis is applicable to a wide variety of complex systems and processes: the systems of statistical physics (and therefore, of physical chemistry); ecosystems (as Lotka hoped) and therefore evolution by both natural selection and genetic drift; various aspects of human societies and economies, such as traffic flow along the highways and through the internet's tubes. Much of the simplicity and stability we see around us can be accounted for in this way. But there is more.

2.4 *Abstract Difference-Making Structures*

The statistical approach to explaining the simplicity of macrolevel behavior—what I call enion probability analysis—can be understood as a demonstration that the incredibly intricate and involved to-ings and fro-ings of a system's microlevel parts make no difference to its macrolevel behavior: they are fluctuations that, because of their statistical profile, cancel out, leaving macrolevel states and changes in state to be determined by relatively stable and measurable features such as molecular mass, lynx physiology, and so on.

Other methods explaining the emergence of macrosimplicity may be viewed in the same way: they show that certain elements of a system, elements that contribute significantly to its convoluted microlevel dynamics, make no difference to the behavior of its high-level properties, though these properties may themselves be aggregates of the very elements that are seen to make no difference. These methods, including EPA, thus identify certain rather abstract properties of the system in question—abstract in the sense that they may be shared by systems that in many other respects differ significantly—and they show that the abstract properties alone are difference-makers for high-level behavior. A system with those properties will exhibit that behavior, regardless of how the properties are realized.[6]

Boltzmann's kinetic theory, for example, uses an implementation of EPA to show that the physical differences between the geometry of different gas molecules make no difference to gases' tendency to move toward equilibrium, and thus to conform to the second law of thermodynamics: all that matters is that, upon colliding, the molecules in a certain sense scatter. This scattering character of collisions is the difference-making property that, however it is realized, secures equilibration.

The simplicity of the equilibration dynamics is, I suggest, closely connected to the abstractness of the difference-making property. The connection is not straightforward or directly proportional, but there is a correlation: where you find simplicity, you tend to find abstract difference-makers.

As a consequence, macrosimplicity also goes along with universality: where one complex system behaves in a simple way, many others, often quite different in their constitution, also tend to behave in that same simple way. Looking beyond EPA, then—though the connection holds there, too—the same techniques that explain macrosimplicity tend also to explain universality,

6. This is not identical, though it is close, to the notion of difference-making that I have developed in my work on scientific explanation (Strevens 2008). This essay does not require, I think, a formal characterization.

in both cases by identifying a highly abstract difference-making structure sufficient, or nearly so, for the phenomenon in question.

Let me give three examples.

Critical Point Phenomena A wide variety of complex systems, when their temperature crosses a certain point, undergo phase transitions in which the macrostate of the system changes qualitatively: liquids freeze; unmagnetized solids become magnetized; disordered rod-like molecules align (Yeomans 1992). For what are called continuous phase transitions, the systems' behavior near the point of transformation—the critical temperature—is strikingly similar: certain physical quantities, such as magnetization, conform to an equation of the form

$$F(T) \propto (T - T_c)^{\alpha}$$

where T is the system's temperature (or other relevant variable), $F(T)$ is the value at temperature T of the physical quantity in question (such as magnetization), T_c is the critical temperature, and α is an exponent that takes the same value for large classes of systems that otherwise differ greatly in their constitution.

Two things to note: first, these are complex systems with enormous numbers of degrees of freedom, yet their macrolevel behavior near the critical temperature is extremely simple; second, there is great universality to this behavior, with many dissimilar systems following the same equation in the critical zone. The explanation of this simplicity and universality consists in a demonstration that almost everything about such systems makes no difference to their behavior in the critical zone; what matters is that they have a certain abstract structure, shared by some simple models used to study critical behavior, such as the Ising model.[7]

7. Wilson (1979) gives an accessible yet satisfying account of critical point universality, written by the physicist who won the Nobel prize for its explanation. For a philosophical treatment, see Batterman (2002).

The Neutral Theory of Biodiversity Certain features of certain complex ecosystems (such as tropical forests) appear to be the same regardless of the species that constitute the system. One such feature is the distribution of species abundance, that is, the relative numbers of the most abundant species, the second most abundant species, and so on down to the rarest species.

The neutral theory of biodiversity explains the universality of these features using models of ecological dynamics that make no distinction among species, and that do not take into account, in particular, the degree of adaptedness of the species to the environment. It is shown that in the models, the observed abundance curve obtains, and it is claimed that real ecosystems conform to the curve for the same reason (Hubbell 2001). If that is correct, then the patterns of abundance in real ecosystems owe nothing to the fitness of the species in the system, but are instead explained as a matter of chance: some species are more plentiful merely because, to put it very simply, they happened to come along at the right time, and so managed to establish a decisive presence in the system.

The Topology of the World Wide Web Patterns of linkage in the World Wide Web have, it seems, a certain topology wherever you look: the probability $P(n)$ that a given web site has n incoming links from other sites falls off as n increases, following a power law

$$P(n) \propto n^{-\gamma}$$

where γ is a little greater than 2. As a consequence, the topology of the Web is "scale-free": like many fractals, it looks the same in the large and in the small.

The same structure shows up in many other sizable networks with different values for γ: patterns of scientific citation ($\gamma \approx 3$), electrical power grids ($\gamma \approx 4$), and patterns of collaboration among movie actors ($\gamma \approx 2.3$). Barabási and Albert (1999) cite these phenomena and propose to explain them as a result of a preferential attachment process, in which the probability of a node's gaining a new connection is proportional to the number of the node's existing connections. If this is correct then it seems that the content of a website or a scientific paper makes no difference to its chances of being linked to or cited: all that matters is its current degree of connectivity. That the dynamics of connection depend on so little explains at the same time, then, the simplicity and the (near?) universality of the patterns they produce.

There are other ways to explain probabilistic power laws and the resulting scale freedom, however. Adamic and Huberman (2000), focusing on the topology of the Web in particular, argue that preferential attachment predicts a strong correlation between age and connectedness that does not exist; they propose an alternative explanation, also highly abstract but giving a web site's content a role in attracting links.[8]

Similar critiques have been made of the neutral theory of biodiversity as well as of other models that purport to explain interesting swathes of universality by treating factors that seem obviously relevant as non-difference-makers (such as Bak's (1996) proposal to use a model of an idealized pile of sand to explain patterns in earthquakes, stock market crashes, and extinctions): there are many different ways to account for a given simple behavior. What explains a power law probability distribution in one system may be rather different from what explains it in the next.

2.5 *Evolvability*

Macrosimplicity tends to provide, as I remarked earlier, a relatively stable environment in which sophisticated systems may evolve—an environment in which air is distributed uniformly, the same food plants and animals stay around from year to year, and so on. But stability can be more than just a backdrop against which natural selection and other evolutionary processes (such as learning) operate. It can be the very stuff of which sophisticated

8. In a useful survey of "network science", Mitchell (2009, 253–255) covers these and further sources of skepticism, including assertions that power laws are not nearly so widely observed as is claimed.

systems are made.

Simon (1996) identified modularity and the concomitant stability as a key element of evolvable systems. Thinking principally of gradualist natural selection, in which adaptation is a consequence of a long sequence of small changes in a system's behavior, he reasoned that in a convoluted system such changes will be hard to come by since minor tweaks in implementation will tend—at least sometimes—to effect radical transformations of behavior, in which all accumulated adaptation will be lost. What is needed for evolution, then, is a kind of organization that responds to minor tweaks in a proportional way: the system's behavior changes, but not too much.[9]

This amounts to a kind of measured stability. An evolvable system should have many stable states or equilibria; natural selection can then move among them, finding its way step by step to the fittest. Stuart Kauffman describes such systems—perhaps overly dramatically—as occupying the "edge of chaos": they are on the one hand not totally inert, their behavior too stubbornly fixed to change in response to small tweaks, but on the other hand not too close to chaos, where they would react so violently to small changes that gradual evolutionary progress would be impossible (Kauffman 1993).

The near-decomposability of hierarchical or modular systems described in section 2.2 is, Simon thought, the key to evolvability. When such systems take a microlevel step in the wrong direction, Simon argued, then at worst a single module is disabled; the hard-won functionality of the rest of the system is preserved.[10]

9. Qualitatively the same phenomenon is sought by engineers of control systems, who want their systems to react to input, but not to overreact—a central part of the subject matter of Wiener's (1965) proposed science of sophisticated complexity, which he called *cybernetics*. The additional chapters in the second edition of Wiener's book touch upon evolution by natural selection.

10. At least, that is the upshot of Simon's fable of the watchmakers Tempus and Hora. In real biology, deleterious mutations are often fatal; the evolvability challenge is more to find scope for small improvements than to rule out catastrophic errors. But perhaps it is also a consequence of near-decomposability that a meaningful proportion of tweaks will tend to result in small improvements rather than drastic and almost always disastrous

Kaufmann's best-known work concerns models of genetic regulation (his "*NK* models") that are in no way modular but that exhibit the same modest pliability: they typically respond to small random changes in configuration by changing their behavior, but not too profoundly.

Much recent scientific work on evolvability belongs to the new field of evolutionary developmental biology, in which modularity remains an important theme (see section 3.3).

2.6 Other Approaches

Mathematical derivation has been the preferred tool in the explanations of macrosimplicity and universality surveyed earlier: certain factors are proved not to make a difference to the behavior of certain macrolevel properties (or at least a proof sketch is provided). But complexity theorists have other methods at their disposal.

The first is an empirical search for universal behaviors. Where diverse systems are found exhibiting the same simple macrolevel behaviors, there is at least some reason (perhaps very far from conclusive, as the case of network topology suggests) to think that they share the same abstract difference-making structure.

The second is simulation. Rather than proving that an abstract difference-making structure induces a certain behavior, systems realizing the structure in various ways can be simulated on a computer; if the same behavior appears in each case, there is at least some reason to think that the structure in question is in each case responsible. (The many uses of simulation are treated more generally in a separate entry in this handbook.)

The third is open-ended computer experimentation. Wolfram (2002) advocates the computer-driven exploration of the behavior of cellular automata in the hope of discovering new abstract difference-making structures and developing new ways to model nature.

---

reconfigurations.

3. From Simplicity to Sophistication

Begin with simplicity—often itself emerging from lower-level convolution, as in the case of the simple, regular behavior of DNA molecules, neurons, even animal populations. Begin, that is, with an inventory of parts whose dynamics conforms to simple mathematical equations. What kinds of surprising behavior can you expect from these parts?

Even in isolation, a system whose dynamics is simple in the mathematical sense can do unexpected things. The "catastrophe theory" of the 1970s showed that systems obeying one of a family of simple, smooth macrolevel laws that induce fixed-point equilibrium behavior can, when subject to small external perturbations, leap to a new equilibrium point that is quite far away. This is a "catastrophe"; the mathematics of catastrophes has been used to explain the buckling of a steel girder, the radically different development of adjacent parts of an embryo, and the collapse of complex civilizations (Casti 1994).[11]

The "chaos theory" of the 1980s showed that systems obeying simple macrolevel laws can exhibit a sensitivity to initial conditions that renders their long-term behavior all but unpredictable. What's more, holding the law fixed but altering one of its parameters can change the system's behavior from a simple fixed-point equilibrium, through periodic behavior—repeatedly visiting the same set of states—to chaos by way of a "period-doubling cascade", in which the set of states visited repeatedly doubles in number as the parameter increases or decreases, in accordance with a pattern that shows remarkable universality captured by the "Feigenbaum constants".[12]

Although interesting and largely unforeseen, these behaviors are not particular to complex systems and are not sophisticated—not goal-directed, not adaptive, not intelligent. For sophistication, it seems, you need to put together

11. For an assessment of catastrophe theory, see Casti (1994, 77–83), which includes a useful annotated bibliography.

12. Stewart (1989) provides an accessible summary of this work and other results in chaos theory.

many simple parts in the right sort of way.

Two related questions, then. First, what is the "right sort of way"? What structures give you not merely convolution, but sophistication? Second, must there be an architect? Or do some of these structures emerge spontaneously, in circumstances that are not vanishingly rare? Are there parts that tend to "self-organize" into the kinds of structures that give rise to sophisticated behavior?

3.1 The Nature of Sophistication

What is sophisticated behavior? The notion, although expositorily useful, is a loose one. In a liberal spirit, let me consider a range of possible marks of sophistication.

The first is adaptation or plasticity. An adapted system is one whose behavior in some sense fits its environment (relative to a presumed or prescribed goal). A plastic system is one whose behavior is capable of changing in reaction to the circumstances in order to realize a presumed or prescribed goal—finding its way around obstacles, anticipating difficulties, abandoning strategies that prove infeasible for promising alternatives. It might be a plasmodium, a person, or a self-driving car.

The second characterization of sophistication is considerably less demanding. Some behavior is so stable as to be deadly boring—as is the macrolevel behavior of the rock sitting in the field. Some behavior is so convoluted as to be bewildering in a way that is almost equally boring—as is the microlevel behavior of the rock, with its multitude of molecules all vibrating this way and that, entirely haphazardly. Between the two, things get interesting: there is order, but there is variation; there is complexity, but it is regimented in ways that somehow call out to our eyes and our minds (Svozil 2008). Systems that have attracted complexity theorists for these reasons include the hexagonal cells of Bénard convection, the wild yet composed oscillations of Belousov–Zhabotinsky chemical reactions, the more intriguing configurations of the Game of Life, and the sets of equations that generate the Mandelbrot set and other complex fractals.[13]

A third criterion for sophistication is as strict as the previous criterion is lax: sophisticated behavior should exhibit intelligence. Intelligence is of course connected to plasticity, but it is a more open-ended and at the same time more demanding kind of sophistication.

My fourth and last criterion attempts to generalize plasticity, not like the previous criterion in the direction of thought, but in the (closely related) direction of animal life. What does that mean? Perhaps having a certain sort of "spontaneity". Perhaps having the characteristics of a typical animal body: appendages, controlled motion with many degrees of freedom, distal senses that control movement, the ability to manipulate objects (Trestman 2013).

These, at least, are some of the behaviors that go beyond simplicity and that complex systems theorists have sought to explain using some variant or other of "complexity theory", understood as a big, broad theory of sophisticated behavior. Two exemplary genres of this sort of complexity theory will be considered here: energetic and adaptive approaches.

3.2 *Energetics*

The universe is always and everywhere winding down—so says the second law of thermodynamics. But this does not mean everywhere increasing decay and disorder. Life on Earth evolved, after all, in accordance with the second law, yet it is a story (if you look at certain branches) of growing order. Might the increasing energetic disorder prescribed by the second law tend to be accompanied, under the right circumstances, by an increasing order of another sort—increasing complexity, increasing sophistication? The fabulous principle specifying this conjectured tendency to complexity is sometimes

13. These phenomena are discussed in many books on complexity and related fields. For Bénard and Belousov–Zhabotinsky, see for example Prigogine (1980); for the Game of Life, Mitchell (2009); for the Mandelbrot set, Stewart (1989).

called, whether affectionately or incredulously, the fourth law of thermodynamics. Were it to come to light, it would provide a foundation for a theory of sophisticated complexity with the widest possible scope.

What are the right conditions for the thermodynamic generation of sophisticated structure? Consider, naturally, the planet Earth. Energy pours in from the sun, making things interesting. That energy is then mostly radiated into space, in a sufficiently disordered form to comply with the second law. The second law says that systems tend toward equilibrium, as the system containing the sun, Earth, and space very slowly does. But as long as the sun burns, it maintains the surface of the Earth itself in a state very far from thermodynamic equilibrium, a state in which there are large temperature differentials and inhomogeneities in the material structure of things.

Other smaller and humbler systems are also maintained far from equilibrium in the same sense: a retort sitting over a Bunsen burner, or a thin layer of oil heated from below and pouring that heat into the air above. Taken as a whole, the system obeys the second law. But as long as the heat flows, the system in the middle, between the heat source and the heat sink, can take on various intricate and sophisticated structures. In the thin layer of oil, the hexagonal cells of Bénard convection develop, each transporting heat from bottom to top by way of a rolling motion. Such an arrangement is often called a *dissipative structure*, operating as it does to get energy from source to sink as the second law commands.

That the dissipation is accompanied by the law-like emergence of an interesting physical configuration is what captures the attention of the theorist of complexity and energetics—suggesting as it does that the entropic flow of energy goes hand in hand with sophisticated flow-enabling arrangements. Two laws, then: a "second law" to ensure dissipation; a "fourth law" to govern the character of the complex structures that arise to implement the second law. Followed by a bold speculation: that ecosystems and economies are dissipative structures that can be understood with reference to the fourth law (Prigogine

1980).

What, then, does the fourth law say? In one of the earliest and most lucid presentations, Lotka (1945) proposed a fourth law according to which the throughput of energy in a system maintained far from equilibrium tends to increase. The law says, then, that in systems maintained far from equilibrium, dissipative structures will tend to evolve so as to maximize the rate at which energy passes from source to sink.

Lotka was writing principally about biological evolution, and his reasoning was Darwinian: organisms will evolve, he believed, to maximize as far as possible their uptake and use of energy from the environment, and so, as a second-law corollary, to maximize the amount of energy that they dump back into the environment in disordered form.[14] It is selective pressure, then, that powers the fourth law, by choosing, from among many possible physical configurations, those systems that conform to the law.

The fourth law is not, however, confined to biological systems; according to Lotka it holds in any system of "energy transformers" of the right sort, in virtue of something analogous to natural selection. Lotka goes on to characterize the notion of an energy transformer in greater detail, in terms of functional units termed "receptors", "effectors", and "adjusters". He concludes that "a special branch of physics needs to be developed, the *statistical dynamics of systems of energy transformers*" (179).

A tendency for energy flow to increase is not the same thing as a tendency for sophistication to increase. But followers of Lotka, and to a certain extent Lotka himself, have thought there to be a connection: the most effective way to increase the energy flow through a dissipative structure will typically be to make it more sophisticated. So Depew and Weber (1988, 337) write: "It is an essential property… of dissipative structures, when proper kinetic

14. The key underlying assumption is that there is always some way that at least some species in an ecosystem can put additional free energy to use to increase fitness, and that natural processes of variation (such as mutation and sexual reproduction) will sooner or later stumble upon that way.

pathways are available, to self-organize and… to evolve over time toward greater complexity."[15]

There are numerous difficulties in formulating a fourth law, especially one that has something to say about the emergence of sophistication. Most importantly, perhaps, writers after Lotka have tended to ignore the significance of his careful characterization of energy transformers. It is nevertheless worth putting aside these objections to focus on another problem—that a fourth law has meager empirical content—because rather similar complaints can be made about every general theory or principle of sophisticated complexity.

What does the fourth law predict? The law is not intended to prescribe any particular variety of sophistication; its empirical content consists rather in its ruling out a certain kind of event, namely, a system's failure to become more complex. On the face of things, this prediction seems to be a bad one: evolutionary stability, in which nothing much changes, is the statistical norm. More generally, dissipative structures, whether thin layers of oil or tropical rainforests, do not tend to become arbitrarily complex. It is for this reason that Depew and Weber hedge their statement of the fourth law above by specifying that it holds only "when proper kinetic pathways are available".

What, then, makes a kinetic pathway "proper" or "available"? Not bare physical possibility: a system consisting of many agitated molecules has an infinitesimal probability of doing all sorts of astonishing things—currents in gently heated oil might, for example, form the image of a monkey at a typewriter—yet they do not, because of such events' vanishingly low probability. In that case, what is the fourth law saying? That dissipative structures tend to become more complex, provided that they are not unlikely to do so? Without the hedge, the fourth law is false; with the hedge, it seems not much more than a bland truism.

15. To better secure the connection to complexity (in the sense of sophistication), Depew and Weber amend Lotka's fourth law: rather than maximizing energy flow, structures minimize specific entropy—entropy created per unit of energy flow—thus using energy in a certain sense more efficiently.

3.3 *Adaptation*

The literature on general theories of complexity in the 1980s and 1990s contained numerous references to "complex adaptive systems". Is it their complexity that allows such systems to adapt? Often, what is meant is something different: they have become complex through adaptation (perhaps after some other process of emergence has provided "evolvable" parts, as discussed in section 2.5). In Lotka's theory, organisms evolve to maximize their uptake and use of energy from the environment, and so, as a second-law corollary, they maximize the amount of energy that they dump back into the environment in disordered form.[16] It is selective pressure, then, that secures the truth of the fourth law, by choosing, from among many possible physical configurations, those systems that conform to the law.

Putting thermodynamics entirely to one side, it has of course been appreciated ever since Darwin that natural selection is capable, given the character of life on Earth, of building ever more sophisticated systems (though as remarked in the previous section, complexification is only one evolutionary modus operandi among many). So what of novel interest does putting "adaptive" in the middle of "complex systems" accomplish? To what extent is there a theory of the development of sophisticated complexity by way of natural selection that goes beyond Darwin and modern evolutionary biology?

The answer is that attempts at a general theory of "complex adaptive systems" hope to do one or both of two things:

1. Apply broadly Darwinian thinking outside its usual biological scope, using new methods for discovering and understanding the consequences of selection.

2. Find patterns or tendencies common to a large class of systems that

16. In some places, Lotka implies that it is an ecosystem as a whole that maximizes the rate of energy throughput; how exactly selection explains this systemic maximization is left unspecified.

> includes both those in which adaptation is driven by natural selection and those in which it is effected by other means, such as learning. Gell-Mann (1994a), for instance, provides a very liberal definition of adaptivity, and writes that "an excellent example of a [complex adaptive system] is the human scientific enterprise".

The investigation of "artificial life" is an example of the first sort of project. Under this heading, researchers aim to abstract away from the implementation of reproduction and inheritance on Earth—from DNA and RNA and the cellular mechanisms that coordinate their replication and variation—and model, using computer programs or other fabricated constructs, systems in which these things happen by alternative means. In Tom Ray's influential Tierra model, the stuff of life is code itself: the "organisms" are small computer programs that consume CPU time and use it to replicate themselves, sometimes recombining and sometimes mutating (Ray 1992). What happens in Tierra is not the simulation of evolution by natural selection in some other system, then, but the real thing: the scarce resource is CPU time and survival and reproduction is not represented but rather instantiated in the persistence and replication of units of code. The behavior of these and similar experiments is quite engaging: researchers see the development both of sophisticated complexity (up to a point) and of parasitism. Yet it is as yet unclear whether artificial life has any general lessons to teach about complexity, above and beyond what is already known from mainstream evolutionary biology (Bedau et al. 2000).

When the Santa Fe Institute was founded in 1984 to promote the interdisciplinary study of complex systems, the ideal of a general theory of sophisticated adaptive systems was an important element of its credo. A series of popular books celebrated the promise of this program of research (Lewin 1992; Waldrop 1992; Gell-Mann 1994b). But by the late 1990s, many figures associated with the theory of sophisticated complexity had begun to worry, like Simon (1996), that "complexity… is too general a subject to have much content" (181). Surveying the ghosts of general theories of complexity that have paraded by in the past one hundred years—cybernetics (Wiener 1965), general system theory (von Bertalanffy 1968), autopoiesis (Varela et al. 1974), synergetics (Haken 1983), self-organized criticality (Bak 1996)—the inductively inclined non-specialist might well, at least in the case of sophisticated complexity, concur.

That does not mean that an emphasis on complexity in the abstract is not theoretically fruitful. One of the most exciting fields in science in the last few decades has been evolutionary developmental biology, which uses discoveries about the genetic modulation of development to better understand how complex body plans evolved (Raff 1996; Gerhart and Kirshner 1997). An important part of the explanation seems to be, as Simon envisaged, modularity: some important evolutionary steps were taken not by drastically re-engineering modules' internal workings but by tweaking the time and place of their operation.

The feel of the science is typical of the study of many complex systems. On the one hand, the empirical findings, though hard won through great amounts of tedious and painstaking research on particular systems, cry out to be expressed at the highest level, in abstract vocabulary such as "complexity", "modularity", "switching". On the other hand, attempts to generalize, leaving the particular topic of the development and evolution of life on Earth behind to say something broader about the connection between modularity and complexity, seem to produce merely truisms or falsehoods. Thinking about sophisticated complexity in the abstract remains as enticing, as tantalizing as ever—but the best theories of sophisticated complexity turn out to have a specific subject matter.

4. Conclusion

Two intriguing features of complex systems have been discussed in this essay: simple behavior at the high level emerging from convoluted underpinnings, and sophisticated behavior at the high level emerging from simple underpinnings. Complexity theory has sometimes concerned itself with the one sort of emergence, sometimes with the other, and sometimes it seems to aim for both at the same time, seeking to explain behaviors that are both surprisingly stable and surprisingly sophisticated.

The default approach to complex systems, registered in the segregation of the university departments, is to tackle one kind of stuff at a time. The term "complexity theory" implies a more interdisciplinary enterprise, an attempt to identify commonalities among complex systems with radically different substrates: to find connections between ecosystems and gases, between power grids and the World Wide Web, between ant colonies and the human mind, between collapsing cultures and species headed for extinction.

With respect to the emergence of simplicity from convolution, I have been (as a partisan) optimistic about the possibility of substantive general theories, although a complete understanding of simplicity's basis in complexity will surely require many distinct ideas at varying levels of abstraction.

With respect to the emergence of sophistication from simplicity, there is less progress on general theories to report, and considerable if reluctant skepticism even in friendly quarters—but plenty of interesting work to do on individual systems, with results that will surely continue to fascinate us all.

References

Adamic, L. A. and B. A. Huberman. (2000). Power-law distribution of the world wide web. *Science* 287:2115a.

Bak, P. (1996). *How Nature Works: The Science of Self-Organized Criticality*. Copernicus, New York.

Barabási, A.-L. and R. Albert. (1999). Emergence of scaling in random networks. *Science* 286:509–512.

Batterman, R. W. (2002). *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford University Press, Oxford.

Bechtel, W. and R. C. Richardson. (1993). *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton University Press, Princeton, NJ.

Bedau, M. A., J. S. McCaskill, N. H. Packard, S. Rasmussen, C. Adami, D. G. G. T. Ikegami, K. Kaneko, and T. S. Ray. (2000). Open problems in artificial life. *Artificial Life* 6:363–376.

von Bertalanffy, L. (1968). *General System Theory: Foundations, Development, Applications*. George Braziller, New York.

Boltzmann, L. (1964). *Lectures on Gas Theory*. Translated by S. G. Brush. University of California Press, Berkeley, CA.

Casti, J. (1994). *Complexification*. HarperCollins, New York.

Depew, D. and B. Weber. (1988). Consequences of non-equilibrium thermodynamics for the Darwinian tradition. In Weber et al. (1988).

Gell-Mann, M. (1994a). Complex adaptive systems. In G. Cowan, D. Pines, and D. Meltzer (eds.), *Complexity: Metaphors, Models, and Reality.* Addison-Wesley, Reading, MA.

———. (1994b). *The Quark and the Jaguar: Adventures in the Simple and the Complex*. Henry Holt, New York.

Gerhart, J. and M. Kirshner. (1997). *Cells, Embryos, and Evolution*. Blackwell, Oxford.

Gibbs, J. W. (1902). *Elementary Principles in Statistical Mechanics*. Scribner's, New York.

Haken, H. (1983). *Synergetics, an Introduction: Nonequilibrium Phase Transitions and Self-Organization in Physics, Chemistry, and Biology*. Third edition. Springer-Verlag, Heidelberg.

Hubbell, S. P. (2001). *The Unified Neutral Theory of Biodiversity and Biogeography*. Princeton University Press, Princeton, NJ.

Kauffman, S. (1993). *The Origins of Order*. Oxford University Press, Oxford.

Lewin, R. (1992). *Complexity: Life at the Edge of Chaos*. Macmillan, New York.

Lotka, A. J. (1925). *Elements of Physical Biology*. Williams and Wilkins, Baltimore, MD.

———. (1945). The law of evolution as a maximal principle. *Human Biology* 17:167–194.

Maxwell, J. C. (1860). Illustrations of the dynamical theory of gases. *Philosophical Magazine* 19 and 20:19–32 and 21–37.

———. (1867). On the dynamical theory of gases. *Philosophical Transactions of the Royal Society of London* 157:49–88.

Mitchell, M. (2009). *Complexity: A Guided Tour*. Oxford University Press, Oxford.

Prigogine, I. (1980). *From Being to Becoming*. W. H. Freeman, New York.

Raff, R. A. (1996). *The Shape of Life: Genes, Development, and the Evolution of Animal Form*. University of Chicago Press, Chicago.

Ray, T. S. (1992). An approach to the synthesis of life. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (eds.), *Artifical Life II*. Addison-Wesley, Redwood City, CA.

Simon, H. A. (1996). *The Sciences of the Artificial*. Third edition. MIT Press, Cambridge, MA.

Stewart, I. (1989). *Does God Play Dice? The Mathematics of Chaos*. Blackwell, Oxford.

Strevens, M. (2003). *Bigger than Chaos: Understanding Complexity through Probability*. Harvard University Press, Cambridge, MA.

———. (2005). How are the sciences of complex systems possible? *Philosophy of Science* 72:531–556.

———. (2008). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, MA.

Svozil, K. (2008). Aesthetic complexity. http://arxiv.org/abs/physics/0505088.

Tolman, R. C. (1938). *The Principles of Statistical Mechanics*. Oxford University Press, Oxford.

Trestman, M. (2013). The Cambrian Explosion and the origins of embodied cognition. *Biological Theory* 8:80–92.

Varela, F. J., H. R. Maturana, and R. Uribe. (1974). Autopoiesis: The organization of living systems, its characterization and a model. *Biosystems* 5:187–196.

Waldrop, M. M. (1992). *Complexity: The Emerging Science at the Edge of Order and Chaos*. Simon & Schuster, New York.

Weaver, W. (1948). Science and complexity. *American Scientist* 36:536–544.

Weber, B., D. Depew, and J. Smith (eds.). (1988). *Entropy, Information, and Evolution*. MIT Press, Cambridge, MA.

Wiener, N. (1965). *Cybernetics, or Control and Communication in the Animal and the Machine*. Second edition. MIT Press, Cambridge, MA.

Wilson, K. (1979). Problems in physics with many scales of length. *Scientific American* 241:158–179.

Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media, Champaign, IL.

Yeomans, J. M. (1992). *Statistical Mechanics of Phase Transitions*. Oxford University Press, Oxford.