

**COMPLEX ADAPTIVE SYSTEMS:
EXPLORING THE KNOWN, THE UNKNOWN
AND THE UNKNOWABLE**

SIMON A. LEVIN

ABSTRACT. The study of complex adaptive systems, from cells to societies, is a study of the interplay among processes operating at diverse scales of space, time and organizational complexity. The key to such a study is an understanding of the interrelationships between microscopic processes and macroscopic patterns, and the evolutionary forces that shape systems. In particular, for ecosystems and socioeconomic systems, much interest is focused on broad scale features such as diversity and resiliency, while evolution operates most powerfully at the level of individual agents. Understanding the evolution and development of complex adaptive systems thus involves understanding how cooperation, coalitions and networks of interaction emerge from individual behaviors and feed back to influence those behaviors. In this paper, some of the mathematical challenges are discussed.

1. INTRODUCTION

The notion of complex adaptive systems has found expression in everything from cells to societies, in general with reference to the self-organization of complex entities, across scales of space, time and organizational complexity. Much of our understanding of complex adaptive systems comes from observations of Nature, or from simulations, and a daunting challenge is to summarize these observations mathematically. In essence, we need a statistical mechanics of heterogeneous populations, in which new types are continuously appearing through a variety of mechanisms, mostly unpredictable in their details.

Complexity comes in a variety of forms, and not all complex, self-organizing systems are adaptive systems in the sense that I will use the term in this paper. Soap bubbles, or the frost-heaving patterns in tundra soils, may arise from self-organization without benefit of any selection or design. Indeed, local variational principles may capture the essential features of such pattern formation mathematically, for example through the minimal surfaces equation, implying some sort of constrained optimization. However, the process by which this optimization takes

Received by the editors December 12, 2000, and, in revised form, February 21, 2002.

2000 *Mathematics Subject Classification.* Primary 92B05, 92D15, 92D40.

Key words and phrases. Complex adaptive systems, self-organization, natural selection, ecosystems.

It is a pleasure to acknowledge the support of the Alfred P. Sloan Foundation, grant award 97-3-5, and of the National Science Foundation, grant award DEB-0083566. Jonathan Dushoff and Helene Muller-Landau provided useful comments.

place is quite different from what happens in biological evolution, in which multiple reproductive lines compete based on realized fitnesses.

Even for systems for which such evolution applies, it is crucial to distinguish the level or levels at which selection takes place. The process of animal development, for example, is one in which macroscopic patterns emerge from largely microscopic interactions as an essentially homogeneous egg becomes differentiated through a self-organizing process guided by local interactions. Natural selection has shaped the local rules of interaction, in accordance with the consequences that different local rules imply for the fitness of the organism as a whole. In contrast, the development of large animal aggregations, ecosystems, economies and societies represent self-organizing processes also shaped by selection, but where the relevant fitness differences are at more local scales, and the emergent patterns are to some extent epiphenomena. In economics, a powerful example is Adam Smith's notion of "the invisible hand," which he argued would serve to assure that "a beneficent social order emerged as the unintended consequences of individual human actions" [57]. Closely related is the notion that efficient (Pareto optimal) utilization of resources will emerge from individual rational economic behavior.

In general, it is not possible to separate complex adaptive systems into neat categories based on whether and where selection is operating. In most systems, selection is manifest on multiple interacting scales. In biological evolution, much theory and research deals for example with intragenomic conflicts, or with the interplay between individual and group selection, or with selection processes that might operate at the level of coevolution between species. In particular, understanding the evolution of altruism and cooperation has been one of the most fascinating and informative challenges in evolutionary theory since Darwin first puzzled about such things. Identical issues arise in the analysis of social and economic systems, and indeed have led to increasing cross-fertilization between evolutionary biology and social science. In economic systems, the self-organizing aspects as reflected in the notion of emergent Pareto optimality must be complemented by the top-down feedbacks that come from central agencies, such as the Federal Reserve in the United States, which modify local rules of interaction in order to direct the system towards desired macroscopic outcomes.

In general, I will define complex adaptive systems by three properties [33]:

- (1) diversity and individuality of components,
- (2) localized interactions among those components, and
- (3) an autonomous process that uses the outcomes of those interactions to select a subset of those components for replication or enhancement.

This is a fairly general and flexible definition, which allows the notion of "local" to be modified as the situation demands. In most cases, the notion of localization of interactions will involve a range of scales.

The prototypical model of a complex adaptive system is the evolving biosphere; hence, this will form the core of my discussion for the rest of this paper. However, all of the essential principles carry over to other complex adaptive systems, and I will point out appropriate parallels along the way.

2. MATHEMATICAL ECOLOGY AND EVOLUTIONARY BIOLOGY:
A SYNTHETIC VIEW WITH IMPLICATIONS FOR THE STUDY
OF COMPLEX ADAPTIVE SYSTEMS MORE GENERALLY

Most mathematicians are aware of the contributions of the great mathematician Vito Volterra to the study of biological populations, and probably of the equations that carry his name and that of the polymath demographer, Alfred Lotka. Volterra, of course, was known first and foremost for his work on integral equations, but was drawn by his son-in-law, the distinguished fisheries biologist Umberto D'Ancona, to consider the mechanisms underlying the observed fluctuations in the Adriatic fisheries. By constructing a set of two ordinary differential equations representing the dynamics of predator and prey, Volterra demonstrated that such systems would have a natural tendency to oscillate. A century of mathematical ecology has built on the Lotka-Volterra equations (Lotka developed similar equations independently) for interacting species to develop not only deep insights into biological dynamics, but also an overbuilt edifice of mathematical complexity. For Volterra, and for this entire tradition, the populations were homogeneous in their composition, made up of identical individuals whose characteristics did not change over time. Evolution was considered to be operating over time scales too long to be relevant to the problems being addressed.

Evolutionary theory, of course, had had a pretty good beginning by the time of Volterra. Darwin's seminal voyage on the *Beagle* took place from 1831-1836, and his masterpiece, *On the Origin of Species*, was published in 1859, the year before Volterra was born; Volterra's main work in mathematical ecology was more than a half-century later. Volterra's work in ecology thus was taking form simultaneously with the mathematical development of evolutionary theory, as pioneered by Sir Ronald Fisher, Sewall Wright and J.B.S. Haldane. Yet the separation between the two was maintained on both sides. For Volterra, evolution was too slow to be important; for the theorists of evolutionary change, interspecific dynamics were also seen to be on a different scale and were largely ignored.

Today, we know that the scales of change are not so easily separable, especially when one is considering the interactions among species with very different generation times. We need look no further than the rapid evolution of antibiotic resistance in bacteria, or the continual emergence of new viral strains, to realize that evolutionary and ecological phenomena can affect one another on comparable time scales. Furthermore, and more subtly, even when we are simply trying to understand purely ecological phenomena, it is valuable to embed them within a broader evolutionary framework so that we can understand how current adaptations arose, and in what sense they are the best that Nature could do.

More generally, the exploration of any complex adaptive systems involves the interplay between two processes—the emergence of pattern in a system of fixed entities, and the continual appearance of new kinds of entities. The first involves the winnowing of variation—the loss of dimensionality and complexity, and the increase of entropy, characteristic of the dynamics of most self-organizing systems. Robust patterns, representing asymptotic outcomes of dynamical systems, typically can be reached from multiplicities of starting conditions. The study of pattern formation in diffusion-reaction and other systems has been a rich area of investigation across the sciences. This, however, is only half of the story. The winnowing of variation

must be balanced against the appearance of new variation; else the systems will run down.

Evolutionary theory teaches that adaptive change due to natural selection depends upon two factors: the selective differential among distinct types, and the degree of variation among those types in the population. Absent either, there can be no selective change. The standard mathematical formalism describes the dynamics of selection given a distribution of types, but is agnostic concerning the generation of new variation. That is left to complementary theories of mutation and recombination, and to the consideration of mutation-selection balance and other frameworks that consider both processes simultaneously. Still, the fundamental problem remains that, though one may be able to estimate the rate of innovation, say through mutation, the exact nature of those innovations cannot be predicted. Thus the dynamics of the full system take place in a high-dimensional space whose axes are unknown.

I will treat these problems in turn—first the problem of pattern formation, and then the evolutionary dynamics of systems subject to innovation. In the latter, especially, there is a plethora of unsolved problems.

3. PATTERN FORMATION

The challenge of explaining pattern formation has attracted the attention of researchers in every branch of science because of its centrality. In general, it deals with the dynamics of collectives and with understanding how patterns manifest at the level of those collectives are mediated at the level of the individual entities. In continuum descriptions, individuals disappear and are replaced by local densities, but the essential issues remain the same.

One common framework for investigating many problems in pattern formation involves the use of systems of diffusion and reaction, in which populations of chemicals or individuals interact locally and are redistributed via random movement (perhaps modified to represent various forms of taxis). Alan Turing, the great mathematician and computer scientist, applied such a framework to understand the fundamental problem in developmental biology—how, without a blueprint, the undifferentiated egg develops into a highly differentiated organism. Turing [56] hypothesized the existence of two chemicals—one an *activator* and the other an *inhibitor*—which interact locally and are redistributed via diffusion. The activator stimulates the production not only of the other type, but also of itself, whereas the inhibitor inhibits the production of both. The relevant system of equations then becomes

$$(3.1) \quad \begin{aligned} \frac{\partial u}{\partial t} &= f(u, v) + D_u \nabla^2 u, \\ \frac{\partial v}{\partial t} &= g(u, v) + D_v \nabla^2 v, \end{aligned}$$

in which u represents the concentration of the activator, and v represents the concentration of the inhibitor. The two diffusion rates, D_u and D_v , are allowed to differ, due to the inherent mobilities of the two molecules.

Turing assumed that, in the absence of diffusion, the underlying system of ordinary differential equations (odes) embedded in (3.1) has an asymptotically stable

equilibrium, $(\underline{u}, \underline{v})$. The linearization (Jacobian) matrix takes the form

$$(3.2) \quad J = \begin{bmatrix} a & b \\ c & d \end{bmatrix},$$

in which a, b, c , and d represent the partial derivatives of f and g with respect to u and v , at the equilibrium. Asymptotic stability means (ignoring the marginal case) that the eigenvalues of the linearization matrix all have negative real part.

For the full system of partial differential equations (pdes), the ode equilibrium becomes translated into a uniform equilibrium valid throughout space. According to standard theory, asymptotic stability is examined with reference to perturbations of the form

$$(3.3) \quad w = \exp(\sigma t)W,$$

in which w is the displacement vector $(u - \underline{u}, v - \underline{v})$, σ is the wave number and W is a time-independent solution of the spatial eigenvalue problem

$$(3.4) \quad \nabla^2 W + \sigma^2 W = 0,$$

subject to the relevant boundary conditions for (3.1) [41].

The corresponding linearization matrix then becomes

$$(3.5) \quad J^* = \begin{bmatrix} a - D_u \sigma^2 & b \\ c & d - D_v \sigma^2 \end{bmatrix}.$$

For appropriate values of a, b, c and d , consistent with stability of the ode system, the pde equilibrium becomes unstable when the ratio D_v/D_u of diffusivities exceeds a threshold amount, greater than unity. That instability is manifest through the growth of a range of perturbations of the form given in (3.3), with maximal growth exhibited for some intermediate “most dangerous mode” (see Figure 1). This is linear stability theory, but it may be shown [52] that the destabilized uniform pattern is replaced asymptotically by a stable non-uniform solution (a “dissipative structure”).

Turing’s interpretation of the significance of this was that, given the existence of two such morphogens, random initial perturbations would invariably translate the initially undifferentiated (homogeneous) distribution into a stable prepattern, reliably reached independent of the nature of the perturbations, which could underlie morphological differentiation through catalytic processes. It was a startling and penetrating insight, and continues to be the preferred model today for many forms of pattern formation. In other situations, where the morphogens cannot be found, or are known not to exist, other mechanisms for endogenous pattern formation have displaced these so-called “diffusive instabilities”. Still, the fundamental ideas remain much the same, with pattern emerging from local breaking of symmetry and long-range stabilization of the new patterns [38].

The success of Turing’s insight led to a search for such phenomena in other systems. For example, in consideration of the puzzle of spatial patchiness at virtually all scales in the distribution of oceanic plankton, Levin and Segel [37] and Okubo [48] independently proposed that this patchiness could be understood by blending the insights of Volterra and Turing. In particular, they identified the phytoplankton (small plants) in the upper layers of the oceans with the activator in Turing’s scheme, and the zooplankton (small animals) with the inhibitor. As expected, a spatial pattern emerged, demonstrating patchiness in the distribution of both types

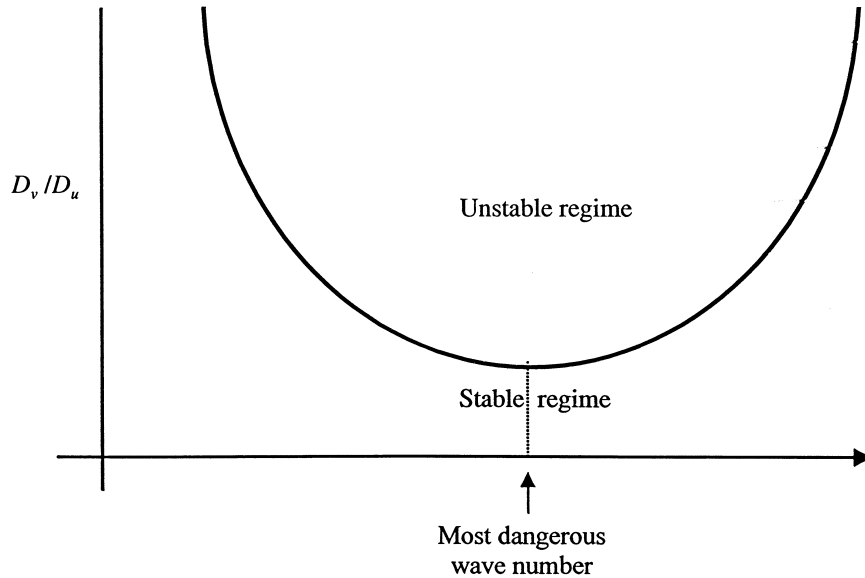


FIGURE 1. Diffusive instabilities: Instability emerges when inhibitor diffusion is high relative to activator diffusion.

of organisms. Because the Turing mechanism relies on high inhibitor diffusion, however, the demonstrated patterns exhibited much more fine-scale patchiness on the part of the phytoplankton than of the zooplankton, in contradiction of observations [22], [34], [60]. In fact, large-scale patchiness in planktonic systems is concordant with patterns exhibited by passive tracers, like temperature or salinity, and hence can probably be adequately explained by reference to the fluid dynamics of the oceans. Small-scale patchiness, on the other hand, primarily of zooplankton, most likely owes its existence to the fact that the animals are active swimmers and preferentially move towards conspecifics, at least on certain scales.

To address this problem, Grünbaum [22], building on the work of others (e.g. [51]), developed an individual-based model of animal aggregation. In essence, he treated individual zooplankton as molecules, and wrote Newtonian laws (force = mass times acceleration) for each. The forces came in a variety of forms, representing the influence of fluid dynamics, the tendency for individuals to follow chemical or resource gradients to find food, and the tendency of individuals to be attracted or repelled by other individuals or their motions. From this Lagrangian description, he derived a statistical mechanics of the ensemble of individuals, and then converted that into an Eulerian description by assuming that local densities followed a Poisson distribution. In this way, the densities in the Eulerian integro-partial-differential equation could be translated into probabilities of observing discrete numbers of other individuals.

Similar work by other investigators built on this theme, and the state of the art is summarized in Flierl et al. [19]. The starting point is the second-order system

$$(3.6) \quad \begin{aligned} \delta \mathbf{x}_i &= \mathbf{v}_i \delta t + \delta \mathbf{X}_i, \\ \delta \mathbf{v}_i &= \mathbf{a}_i \delta t + \delta \mathbf{V}_i, \end{aligned}$$

in which \mathbf{x}_i is the position of an individual at time t , and \mathbf{v}_i is its velocity. The capitalized terms in each equation encapsulate the assumptions made about changes in position (hopping) or changes in velocity (acceleration), and how they depend on extrinsic factors such as the distribution of other individuals. The objective is to proceed from (3.6) to a continuum equation

$$(3.7) \quad \frac{\partial}{\partial t} \rho = -\nabla \cdot \mathbf{J}$$

in which ρ is the local density of organisms (e.g., zooplankton), and \mathbf{J} is the total flux, whose form depends on the assumptions being made.

To proceed from the individual-based models to the continuum description, one writes an equation for the spatial/velocity density $n(\mathbf{x}, \mathbf{v}, t)$. Specifically,

$$(3.8) \quad n(\mathbf{x}, \mathbf{v}, t + \delta t) = \int d\mathbf{x}' d\mathbf{v}' \mathbf{P}_{\delta\mathbf{x}}(\mathbf{x} - \mathbf{x}' - \mathbf{v}'\delta t; \mathbf{x}', \mathbf{v}', t) \mathbf{P}_{\delta\mathbf{v}}(\mathbf{v} - \mathbf{v}' - \mathbf{a}\delta t; \mathbf{x}', \mathbf{v}', t) n(\mathbf{x}', \mathbf{v}', t),$$

where for example $\mathbf{P}_{\delta\mathbf{v}}(\mathbf{z}; \mathbf{x}', \mathbf{v}', t)$ is the probability that a particle at position \mathbf{x}' with velocity \mathbf{v}' at time t has a random acceleration $\delta\mathbf{V} = \mathbf{z}$. The assumption is made that hops and accelerations are uncorrelated. Equation (3.8) is quite complicated, but can be simplified by use of the Fokker-Planck approach [7] of expansion in Taylor series and truncation, provided the probability distributions are sufficiently sharply peaked. A continuum equation can then be obtained. In the absence of random hops, for example, one obtains the Boltzmann equation in differential form,

$$(3.9) \quad \begin{aligned} \frac{\partial}{\partial t} n(\mathbf{x}, \mathbf{v}, t) = & - \frac{\partial}{\partial x_i} [v_i n(\mathbf{x}, \mathbf{v}, t)] \\ & - \frac{\partial}{\partial v_i} [a_i n(\mathbf{x}, \mathbf{v}, t)] \\ & + \frac{1}{2} \frac{\partial^2}{\partial v_i \partial v_j} [\gamma_{ij} n(\mathbf{x}, \mathbf{v}, t)] \end{aligned}$$

summed over v and j .

The terms γ_{ij} in this equation still need to be defined; they represent the second moments of $\delta\mathbf{V}$ averaged over the time step δt , and hence reflect the assumptions made about how accelerations respond to the distribution of individuals. Equation (3.9) still represents the spatial/velocity density, and must be integrated to yield the desired density equation. The systems must then be closed, for example by the method of moments or by the Chapman-Enskog method [19]. In this way, an appropriate statistical mechanics for animal aggregations can be developed. The methods discussed were for a single population, but could in principle be extended to multiple interacting types. Furthermore, the continuum formulation allows the characterization of groups and the examination of the interplay between group properties and individual properties.

The individual-based approach described above assumes no fixed grid, but allows individuals to roam freely over a continuum. A parallel literature, based on spatial stochastic processes in gridded domains [12], has developed from a variety of perspectives. These include highly successful forest growth simulators, such as SORTIE [49], [50], as well as a diversity of efforts to understand the evolution of cooperation [13], [46] and spite [14]. In Durrett and Levin [13] the initial focus is




FIGURE 2. Simulation of an interacting particle system model of competition between selfish individuals (light pattern) and cooperators (dark pattern).

the spatial grid cell rather than the individual, but the rules of interaction still are derived from assumptions about what individuals do (see Figure 2). Within each cell there are a discrete number of individuals of each of several types, which interact with one another according to fixed rules. Again, one can derive appropriate hydrodynamic (continuum) limits by assuming that local densities follow a Poisson distribution [13]. Other approaches to developing reduced-form descriptions in these systems have succeeded by using moment closure methods [4], [36], [50].

The take-home lesson from all of these studies is that it is indeed possible to develop continuum descriptions of relevant phenomena, but that unless one starts at the level of the individual and derives limits formally, the appropriate equations cannot in general be divined. The failure of the Turing equations to capture the behavior of plankton systems is that the relevant diffusion-reaction equations are not those that would emerge from a naive extension of the mean-field equations by appending diffusion terms, because individuals do not move randomly. Similarly, in the work of Durrett and Levin [13], naive extension of mean-field dynamics to a spatial regime fails to explain how cooperative behavior can emerge from local interactions, in this case because of the infinite speed of propagation implicit in the diffusion term. The correct limiting equations still describe diffusive movement, but the interaction terms assume a non-intuitive form, very different from that described by the odes. The appropriate form minimizes the consequences of having small but non-zero densities of non-cooperators.

Similar problems exist in the study of socio-economic systems, and steps towards a statistical mechanics of societies have been taken [11]. In particular, there is fundamental interest in the evolution of social norms [53], [5], [15], or of language [47], and how such group properties emerge from and feed back to influence individual behavior. However, whereas much progress has been made in the analysis of animal

aggregations, and in particular in sorting out the relative importance of endogenous and exogenous factors, the problems in analyzing human social systems are intrinsically more complex than those of oceanic krill and copepods. Still, the potential payoffs are enormous, and the mathematical challenges irresistibly seductive.

4. EVOLUTIONARY DYNAMICS

As discussed earlier, the examination of mechanisms of pattern formation, which has been so successful throughout science, is only half of the problem in the analysis of complex adaptive systems. The methods of statistical mechanics described in the last section assume that the identities of the players are known. What happens if the composition of the assemblage is constantly being changed by the emergence of new types, previously undescribed? For example, a large literature exists about epidemiological dynamics in general, and about the interactions among cocirculating strains of influenza A [1] in particular. Yet such analyses fall short of providing adequate guidance in the design of new vaccines, which must anticipate the emergence of new strains and even new sub-types with novel antigenic properties. Difficult as this may be, remarkable success has been made in addressing this problem [18] by using detailed molecular information to predict what strains will next become widespread. In other situations, even when such prediction is impossible, certain macroscopic statistical properties (such as the rate of appearance of new mutations) may be knowable, allowing some predictive theory to be developed.

Evolutionary theory regularly has to deal with problems of this sort and to embed what is observed in some unknowable higher-dimensional space of what is feasible [26]. The simplest level of evolutionary analysis is content with understanding how things work, and indeed simply in demonstrating that they do work. Thus, one can examine how eyes work, how birds fly or fish swim, or how lizards control heat loss by adjusting their posture in relation to the sun. This kind of analysis, sometimes called *satisficing*, is not technically evolutionary since it looks neither at alternatives nor at history.

The next level of evolutionary analysis views adaptations through the perspective of optimization theory. Thus, one can assign a fitness to each possible phenotype and ask how natural selection will mold the genome in response to this imposed “fitness surface”. Classical theory [32] holds that, under strongly simplifying assumptions, the mean fitness of a population will increase under natural selection, tending towards a local peak. There are many problems with this result, however, including especially the fact that it assumes that fitnesses are constant. Among other things, it is still possible for the fitness landscape to have multiple local peaks; indeed [28], the number of local optima in real situations may be enormous. From a biological point of view, there is nothing wrong with this; evolution is a historically constrained process, shaped in large parts by frozen accidents of times past. This makes retrospective analysis feasible, but prediction impossible, except in the most general terms [25], [33].

To explore the importance and consequences of multiple peaks in the fitness surface, Kauffman and Levin [28] explored how natural selection might have solved the traveling salesman problem. The method, it turned out, was not so different from the combinatorial optimization methods that are commonly used [40]. The traveling salesman problem seeks to find the shortest path connecting a given set of points in the plane. In their approach, Kauffman and Levin [28] represented

a particular path as a haploid genotype, and allowed modifications to the path by various forms of mutation (for example, by interchanging any two cities on a path). In each generation, a given number of mutants were generated, with the mutation rate determining the genetic distances between the mutants and the original genotypes. In the simplest version of the analysis, the most fit mutant was selected if it was an improvement over the existing type. Kauffman and Levin found that the process converged to an equilibrium, but indeed that there were multiple possible equilibria. Furthermore, the rate of improvement depended crucially on the mutation rate. Mutation rates too low kept the system trapped on local peaks, while mutation rates too high led to the loss of local information about the fitness gradient. Thus, optimal improvement occurred for intermediate rates of mutation.

For Kauffman [27] the lessons learned from this example suggested broad principles, namely that systems function best when they are at the “edge of chaos”, poised between too much rigidity in the face of change and too much change in the face of achieved progress. The term “chaos” here is, of course, not employed in the technical mathematical sense, but the term stuck, especially since the word “chaos” does have meanings outside mathematics. Kauffman’s is an intriguing notion, and one can imagine natural selection shaping our tendency to take risks in accordance with the rate that new threats and opportunities present themselves over evolutionary time. But what of complex adaptive systems, in which the relevant level of selection is well below that of the whole system? Kauffman [27] argues that such systems also will seek out the “edge of chaos”. Is there any reason to expect that?

The most intriguing metaphor for Kauffman is Per Bak’s notion of self-organized criticality, illustrated best perhaps by Bak’s toy model of sandpile development [2]. Bak argues, in essence, that self-organizing systems will show increasing degrees of connectivity until the correlations are pervasive enough that random perturbations will cause cascading patterns of collapse, satisfying a power law in terms of the statistical distribution of disturbances. In his examination of a variety of phenomena (such as mass extinctions, or stock market fluctuations), he claims support for this theory in a variety of power laws, as would be predicted from the general theory. Intriguing as this notion is (and I believe there is some truth in it), there are a number of problems. First of all, power laws can arise from a variety of mechanisms; so, whereas their existence is consistent with the notion of self-organized criticality, it hardly is proof of it. As concerns the evidence for mass extinctions in the fossil record, we know well that most of these were not endogenous in origin, and alternative models explain the data equally well or better [44].

Secondly, ecological systems are not sandpiles. They exhibit a high degree of heterogeneity, as do all complex systems, and it is this heterogeneity that is the stuff of evolution. The fitnesses of individual types are, furthermore, not constants, but are contingent on the composition of the population; this is what is known in the evolutionary literature as frequency dependence. Finally, through evolution and local interaction, patterns of modularity develop that limit the spread of disturbances. It is these three features—heterogeneity (and its maintenance), frequency dependence and modularity—that complicate the picture and that will occupy most of my attention for the rest of this paper.

The importance of frequency dependence in considering evolutionary phenomena has long been understood and has long stimulated efforts to apply the theory of

games [39], [54]. The evolution of cooperation, in particular, cannot succeed without some reciprocity, but the context-dependent valuation of adaptations applies to everything from the peacock's ostentatious plumage to the simplest devices for seed dispersal [23]. Maynard Smith [55] initiated a revolution in evolutionary ecology, however, with his modification of simple formal notions from game theory to characterize evolutionary behaviors. In particular, he introduced the notion of an *evolutionarily stable strategy (ESS)* as one that, once established in the population, could not be invaded by a mutant type. In particular, if $E(u, v)$ is the payoff to strategy u when playing against strategy v , an ESS u is defined by the properties that, for all alternative strategies v ,

$$(4.1) \quad \begin{array}{l} \text{either} \quad E(u, u) > E(v, u) \\ \text{or} \quad E(u, u) = E(v, u) \text{ and } E(u, v) > E(v, v). \end{array}$$

It is easy to confirm that, if individuals play other types at random, these are exactly the conditions needed to guarantee that the average payoff to u cannot be exceeded by the average payoff to any other type, provided that the proportion of type u in the population is arbitrarily close to 1. (Indeed, as long as u is not fixed in the population (i.e. at 100%), its average payoff must exceed any other type, provided interactions are random.) u may represent a pure strategy, or it may represent a mixed strategy, in which each of several pure strategies is played with a given probability. In the latter case, the ESS probabilities of playing each strategy correspond to a Nash equilibrium [42], [43], and the population equilibrium could also contain the corresponding pure strategists in proportions according to the same probabilities.

In concept, the notion of the ESS is powerful; in practice, it often can be difficult to define. Individual fitnesses are not defined by games played once during a lifetime, but by repeated games carried out in an explicit spatial context and distributed over the lifetime of an individual. Throughout that lifetime, the individual will be encountering a diversity of other individuals and strategies, and under changing environmental conditions. The fitness may need to be measured as an integral over time, for example, and the game theoretic problem becomes a problem in the calculus of variations [24]. Furthermore, the notion of an ESS is a local one in frequency space, so it is not unusual for systems to possess multiple ESSs. Biologically, this is quite comforting, since many different strategies do exist in Nature; mathematically, it enriches the problem, but makes solution more difficult.

Other complications exist. The notion of an ESS describes a static situation. If the payoff in an evolutionary game translates into fitness, the question still remains as to whether an ESS is dynamically achievable. Furthermore, as in the case of influenza, there may be constraints on what novel types may arise, but there is still an infinity of possible mutations to contend with. We must therefore introduce a dynamic, in a space permitting an infinity of possibilities, in order to address these challenges.

Suppose that there is a continuum of possible strategies, and that the dynamics of evolution assumes small mutations from existing strategies. Although there is a continuum of possible strategies, populations are made up of discrete individuals; thus, in reality, only a finite number of strategies will be represented at any one time. Let the extant strategies be indexed by i or j , and let for example p_i denote

the fraction of the population composed of type i individuals (strategists). Then the average payoff to a type i individual, under the assumption of random interactions, is

$$(4.2) \quad W_i = \sum p_j E(u_i, u_j)$$

and we define the *relative fitness* Φ_i of u_i to be

$$(4.3) \quad \Phi_i = \frac{W_i}{W},$$

where W is the average payoff in the population, namely

$$(4.4) \quad W = \sum p_j W_j.$$

This allows us to embed the dynamic in the standard replicator scheme in population genetics. In particular, if we track changes in discrete time and denote by p'_i the frequency of u_i in the successor generation, then

$$(4.5) \quad p'_i = p_i \Phi_i = p_i \frac{W_i}{W}.$$

All of this could be extended naturally to a continuum of types by use of density functions, but the discrete formulation is clearer for exposition.

In the dynamic setting, an ESS is a type that, once established in the population, cannot be displaced by an invader. Let $\phi(v, u)$ denote the (relative) fitness of strategy v in a population dominated by strategy u . Thus, according to (4.3),

$$(4.6) \quad \phi(v, u) = \frac{E(v, u)}{E(u, u)}.$$

We may now redefine the concept of an ESS in terms of ϕ . Note first that

$$(4.7) \quad \phi(u, u) = 1 \text{ for all } u.$$

Furthermore, an ESS w is defined by the condition that

$$(4.8) \quad \phi(v, w) \leq \phi(w, w) \text{ for all } v.$$

We broaden this a bit to introduce the notion of a *local ESS* as one for which (4.8) holds for all v close to w . Thus, a local ESS w is one for which $\phi(v, w)$ is maximized as a function of v at $v = w$.

Next, consider what would happen if

$$(4.9) \quad E(w, w) > E(v, w)$$

for all other strategies v , or at least for all other nearby strategies in strategy space, so that w is an ESS, but that, in addition,

$$(4.10) \quad E(w, u) < E(u, u)$$

for all nearby strategies. The consequence of (4.10) would be that w , once established, cannot be invaded, but that, in turn, it cannot invade any nearby strategies. Dynamically, it is not clear from these facts alone whether the ESS w will survive or be eliminated from the population, or indeed how it might ever arise in a population. This is not a pathological case, but one that can arise quite robustly in realistic problems; see for example [8], [16], and [35]. To address this anomaly, Eshel [16] introduced the notion of a continuously stable strategy (CSS) as one

that would be achieved dynamically. The notion of an ESS requires that $\phi(v, u)$ be maximized as a function of v at $v = u = w$, so that

$$(4.11) \quad \frac{\partial^2 \phi}{\partial^2 v} \leq 0$$

there. Continuous stability requires that, in addition,

$$(4.12) \quad \frac{\partial^2 \phi}{\partial^2 v} + \frac{\partial^2 \phi}{\partial u \partial v} \leq 0,$$

which means that types closer to a CSS can invade types farther away.

Alternatively, by differentiating (4.7) twice, one finds that

$$(4.13) \quad \frac{\partial^2 \phi}{\partial^2 u} + \frac{2\partial^2 \phi}{\partial u \partial v} + \frac{\partial^2 \phi}{\partial^2 v} = 0,$$

so that (4.12) is equivalent to

$$(4.14) \quad \frac{\partial^2 \phi}{\partial^2 v} \leq \frac{\partial^2 \phi}{\partial^2 u}.$$

Note that, in the anomalous example I introduced earlier by conditions (4.9) and (4.10), it follows that $\phi(v, u)$ is maximized as a function of u at $v = u = w$. Hence $\partial^2 \phi / \partial^2 u \leq 0$, so that (4.14) is automatically a more stringent condition than (4.11). This need not always be the case, and realistic situations do indeed arise in which (4.14) is satisfied, but not (4.11). In this case, it may be shown that the dynamic will drive the system towards w , but that w is not an ESS and hence cannot be sustained in the population. This has been termed an evolutionary branch point [20].

These examples demonstrate both the complexity of outcomes that can arise when the full adaptive dynamics of a community are considered, and also the rapidly developing tool chest of methods that are available [20], [35]. We are, to some extent, where population genetics was a half-century ago. From Mendel's initial experimental work, there developed a reasonably complete theory of single-locus population genetics [10]. As understanding of the inter-locus interactions, especially linkage and epistasis, grew, the single-locus theory became extended to two-locus theory [10]. When many loci are involved, however, as is the case for the predominance of traits of ecological importance, a fundamentally different approach is needed, cognizant of the particulate basis of inheritance, but tailored to deal with a virtual continuum of phenotypic outcomes. This quantitative theory [17] has led to powerful new insights at the interface between genetics and ecology. Similarly, in ecology, the literature as regards single species and pairs of interacting species has a long and rich tradition. To deal with the evolution of whole communities, however, and ultimately with the biosphere, we need a quantitative theory that is cognizant of the roles of individuals and individual species, but is tailored to deal with a virtual continuum of behaviors. The lessons and methods of quantitative genetics promise to be a powerful aid in this quest to understand the adaptive dynamics of whole communities.

What are the challenges? In ecological and socioeconomic systems alike, patterns emerge at higher levels as a result of adaptive dynamics at lower levels of integration. The nature of these systems as highly nonlinear complex adaptive systems means that prediction is limited, and that we must develop statistical mechanical methods to extract the knowable from the unknown. Ecological systems show

regularities in the distribution of key nutrients and trace elements, characteristic abundance distributions for species, and patterns of nutrient cycling that are robust over time and sustain life as we know it. Yet it is not enough to suggest, as some extreme versions of GAIA theory do (see [30] and [33]), that biotic evolution has shaped the evolution of the physico-chemical environment to exactly those conditions that humans need for survival. One must examine the coevolution of life and environment by means of a theory that begins from the levels at which evolution operates—organisms and their genomes. Daisyworld theorists [31] recognize this and have tried to demonstrate how biospheric properties may have evolved. But those models are obviously cartoons at this point, caricatures designed to illustrate how system properties might evolve from individual-based selection. More realistic models are needed, tailored to the dynamics of systems where data are available to test hypotheses. For socioeconomic systems, similar efforts are needed to demonstrate how macroscopic properties of resource utilization and population distribution can be related to the behaviors of individual agents. Much research is ongoing based on simulations, for example of market interactions, but an analytical theory is lacking.

One of the fundamental problems in ecological theory today is developing an understanding of what maintains biodiversity and of the consequences of the loss of biodiversity for system functioning. Here, the parallels to similar problems in population genetics are strong; indeed, as discussed earlier in this paper, the fundamental problems in the study of any complex system are understanding what maintains diversity, and how the existence of diversity affects system dynamics. Thus it is not surprising that the methods of population genetics can be transported from one realm to another and brought to bear upon these fundamental issues. In particular, Norberg et al. [45] use a variant of (4.5) to describe the dynamics of a community of interacting species in the face of a changing environment. In essence, the payoff matrices underlying (4.5) are not constant, and immigration events continuously introduce new types. In an effort to extract the knowable from the unknown, Norberg et al. [45] focus attention not on the details of which species are present, but on macroscopic descriptors of system function, for example the means and variance of performance measures as distributed over the population of all species. This serves to reduce dimensionality, but raises the problem of closing the system of descriptors. Borrowing heavily from the techniques of quantitative genetics [3], they use moment closure methods to develop a description that tracks simulated dynamics remarkably well.

The efforts so far serve mainly to demonstrate that progress is possible, and that the problems are rich both mathematically and biologically. The competitive communities considered by Norberg et al. [45] (see also [29] and [6]) are the simplest ecological communities imaginable. How does cooperation arise and become sustained? If ecological and socioeconomic communities really have unique topological structures and patterns of modularity and connectivity [9], [58], [59], how do these emerge from the local dynamics of self-organization? And ultimately, perhaps, how do social norms arise in human societies and become sustained against external influences? In particular, can this understanding help us to sustain patterns of behavior that serve the common good, and to change antiquated frameworks that are resistant to progress simply because they have become frozen accidents of a

cultural evolutionary process? These are the issues that make the study of complex adaptive systems so seductive.

5. THE KNOWN, THE UNKNOWN AND THE UNKNOWABLE

The study of complex adaptive systems is the study of systems limited in their predictability. Because complex adaptive systems are systems in which microscopic interactions and evolutionary processes give rise to macroscopic phenomena through nonlinear interactions, these systems are subject to path dependence, with implications for the likelihood of multiple stable states, chaotic dynamics and frozen accidents.

However, although it is impossible to predict every detail of system evolution, broad features are knowable. In the metaphor of Ralph Gomory [21], the central problem is to develop an appropriate statistical mechanics that allows one to separate the knowable unknown from the truly unknowable. Such a mechanics will have to deal with heterogeneous ensembles of interacting agents and with the continual refreshment of that ensemble by novel and unpredictable types. The challenges are immense, but sufficient progress already is being made to suggest that the payoffs will be great.

REFERENCES

- [1] V. Andreasen, J. Lin, S. A. Levin, *The dynamics of cocirculating influenza strains conferring partial cross-immunity*, J. Math. Biol. **35** (1997), 825–842. MR **99a**:92013
- [2] P. Bak, *How Nature Works*, Springer-Verlag, New York, 1997.
- [3] N. H. Barton, M. Turelli, 1994. *Genetic and statistical analyses of strong selection on polygenic traits: What, me normal?* Genetics **138** (1994), 913–941.
- [4] B. M. Bolker, S. W. Pacala, S. A. Levin, *Moment methods for stochastic processes in continuous space and time*, The Geometry of Ecological Interactions: Simplifying Spatial Complexity, Cambridge University Press, Cambridge, 2000, pp. 388–411.
- [5] S. Bowles, *Cultural evolution and human social structure*, unpublished manuscript.
- [6] L. A. Buttel, R. Durrett, S. A. Levin, *Competition and species packing in patchy environments*, Theoret. Population Biol. **61** (2002), 265–276.
- [7] S. Chandrasekhar, *Stochastic problems in physics and astronomy*, Rev. Modern Phys. **15** (1943), 1–89. MR **4**:248i
- [8] D. Cohen, S. A. Levin, *The interaction between dispersal and dormancy strategies in varying and heterogeneous environments*, Mathematical Topics in Population Biology, Morphogenesis and Neurosciences, Proc. Kyoto 1985, Springer-Verlag, Heidelberg, 1987, pp. 110–122. MR **88h**:92002
- [9] J. E. Cohen, *Food webs and community structure*, Perspectives in theoretical ecology, Princeton University Press, Princeton, NJ, 1989, pp. 181–202. MR **92a**:92007
- [10] J. F. Crow, M. Kimura, *An Introduction to Population Genetics Theory*, Harper & Row, New York, 1970. MR **42**:8944
- [11] S. Durlauf, *Statistical mechanics approaches to socioeconomic behavior*, The Economy as a Complex Evolving System II, Addison-Wesley, Redwood City, CA, 1997.
- [12] R. Durrett, S. A. Levin, *Stochastic spatial models: A user's guide to ecological applications*, Philos. Trans. Roy. Soc. London Ser. B **343** (1994), 329–350.
- [13] R. Durrett, S. A. Levin. *The importance of being discrete and (spatial)*, Theoret. Population Biol. **46** (1994), 363–394.
- [14] R. Durrett, S. A. Levin. *Allelopathy in spatially distributed populations*, J. Theoret. Biol. **185** (1997), 165–171.
- [15] P. Ehrlich, S. A. Levin, *In preparation*.
- [16] I. Eshel, *Evolutionary and continuous stability*, J. Theoret. Biol. **103** (1983), 99–111. MR **85a**:92024
- [17] D. S. Falconer, *Introduction to Quantitative Genetics*, The Ronald Press Co., New York, 1960.

- [18] W. M. Fitch, R. M. Bush, C. A. Bender, N. J. Cox, *Long term trends in the evolution of H(3) HA1 human influenza type A*, Proc. Nat. Acad. Sci. U.S.A. **9** (1997), 7712–7718.
- [19] G. Flierl, D. Grünbaum, S. Levin, D. Olson, *From individuals to aggregations: the interplay between behavior and physics*, J. Theoret. Biol. **196** (1999), 397–454.
- [20] S. A. H. Geritz, J. A. J. Metz, E. Kisdi, G. Meszéna, *Dynamics of adaptation and evolutionary branching*, Phys. Rev. Lett. **78** (1997), 2024–2027.
- [21] R. E. Gomory, *The known, the unknown and the unknowable*, Sci. Amer. **272** (1995), 120.
- [22] D. Grünbaum, *Local processes and global patterns: Biomathematical models of bryozoan feeding currents and density dependent aggregations in Antarctic krill*, Ph. D. thesis, Cornell University, Ithaca, NY, 1992.
- [23] W. D. Hamilton, R. M. May, *Dispersal in stable habitats*, Nature **269** (1977), 578–581.
- [24] Y. Iwasa, S. A. Levin, *The timing of life history events*, J. Theoret. Biol. **172** (1995), 33–42.
- [25] F. Jacob, *Evolution and tinkering*, Science **196**, (1977), 1161–1166.
- [26] F. Jacob, *The Possible and the Actual*, Pantheon, New York, 1982.
- [27] S. Kauffman, *The Origins of Order*, Oxford University Press, New York, 1993.
- [28] S. Kauffman, S. A. Levin, *Towards a general theory of adaptive walks on rugged landscapes*, J. Theoret. Biol. **128** (1987), 11–45. MR **88h**:92031
- [29] A. P. Kinzig, S. A. Levin, J. Dushoff, S. Pacala, *Limiting similarity and species packing in competition-colonization models*, Amer. Natur. **153** 1999, 371–383.
- [30] J. Kirschner, *Presentation for the epistemology panel discussion during the Chapman Conference on the Gaia Hypothesis*, March 7-11, 1988, San Diego, CA.
- [31] T. M. Lenton, J. E. Lovelock, *Daisyworld is Darwinian: Constraints on adaptation are important for planetary self-regulation*, J. Theoret. Biol. **206** (2000), 109–114.
- [32] S. A. Levin, *On the evolution of ecological parameters*, Ecological Genetics: The Interface Sciences, Springer-Verlag, New York, 1978, pp. 2–26.
- [33] S. A. Levin, *Fragile Dominion: Complexity and the Commons*, Perseus Books, Reading, MA, 1999.
- [34] S. A. Levin, A. Morin, T. H. Powell, *Patterns and processes in the distribution and dynamics of Antarctic krill*, Scientific Committee for the Conservation of Antarctic Marine Living Resources Selected Scientific Papers Part 1, SC-CAMLR-SSP/5, CCAMLR. Hobart, Tasmania, Australia, 1989, pp. 281–299.
- [35] S. A. Levin, H. C. Muller-Landau, *The evolution of dispersal and seed size in plant communities*, Evol. Ecology Res. **2** (2000), 409–435.
- [36] S. A. Levin, S. W. Pacala, *Theories of simplification and scaling of spatially distributed processes*, Spatial Ecology: The Role of Space in Population Dynamics and Interspecific Interactions, Princeton University Press, Princeton, NJ, 1997, pp. 271–296.
- [37] S. A. Levin, L. A. Segel, *An hypothesis for the origin of planktonic patchiness*, Nature **259**, no. 5545, (1976), 659.
- [38] S. A. Levin, L. A. Segel, *Models of the influence of predation on aspect diversity in prey populations*, J. Math. Biol. **14** (1982), 253–284.
- [39] R. C. Lewontin, *Evolution and the theory of games*, J. Theoret. Biol. **1** (1961), 382–403.
- [40] S. Lin, B. W. Kernighan, *An effective heuristic algorithm for the traveling-salesman problem*, Operations Res. **21** (1973), 498–516. MR **50**:12194
- [41] J. D. Murray, *Mathematical Biology*, Springer-Verlag, Heidelberg, Germany, 1990. MR **90g**:92001
- [42] J. F. Nash, *Equilibrium points in n-person games*, Proc. Nat. Acad. Sci. U.S.A. **36** (1950), 48–49. MR **11**:192c
- [43] J. F. Nash, *Non-cooperative games*, Ann. of Math. **54** (1951), 286–295. MR **13**:261g
- [44] M. E. J. Newman, G. J. Eble, *Decline in extinction rates and scale invariance in the fossil record*, Paleobiology **25** (1999), 434–439.
- [45] J. Norberg, D. P. Swaney, J. Dushoff, J. Lin, R. Casagrandi, S. A. Levin, *Phenotypic diversity and ecosystem functioning in changing environments: A theoretical framework*, Proc. Nat. Acad. Sci. U.S.A. **98**, no. 20, (2001), 11376–11381.
- [46] M. A. Nowak, S. Bonhoeffer, R. M. May, *Spatial games and the maintenance of cooperation*, Proc. Nat. Acad. Sci. U.S.A. **91** (1994), 4877–4881.
- [47] M. Nowak, D. Krakauer, *The evolution of language*, Proc. Nat. Acad. Sci. U.S.A. **96** (1999), 8028–8033.

- [48] A. Okubo, *Diffusion-induced instability in model ecosystems*, Chesapeake Bay Institute Technical Report 86, Johns Hopkins University, Baltimore, MD, 1974, 17 pp.
- [49] S. W. Pacala, C. D. Canham, J. Saponara, J. A. Silander, R. K. Kobe, E. Ribbens, *Forest models defined by field measurements: II. Estimation, error analysis and dynamics*, Ecol. Mono. **66** (1996), 1–43.
- [50] S. W. Pacala, S. A. Levin, *Biologically generated spatial pattern and the coexistence of competing species*, Spatial Ecology: The Role of Space in Population Dynamics and Interspecific Interactions. Princeton University Press, Princeton, NJ, 1997, pp. 204–232.
- [51] S. Sakai, *A model for group structure and its behavior*, Biophysics (Japan), **13** (1973), 82–90.
- [52] L. A. Segel, S. A. Levin, *Applications of nonlinear stability theory to the study of the effects of diffusion on predator prey interactions*, Topics in Statistical Mechanics and Biophysics: A Memorial to Julius L. Jackson. American Institute of Physics, New York, 1976, pp. 123–152.
- [53] B. Skyrms, *Evolution of the Social Contract*, Cambridge University Press, New York, 1996, 146 pp.
- [54] L. B. Slobodkin, *The strategy of evolution*, Amer. Sci. **52** (1964), 342–357.
- [55] J. Maynard Smith, *Evolution and the Theory of Games*, Cambridge University Press, Cambridge, 1982.
- [56] A. M. Turing, *The chemical basis of morphogenesis*, Philo. Trans. Roy. Soc. London Ser. B **237** (1952), 37–72.
- [57] K. J. Vaughan, *Invisible hand*, The New Palgrave: A Dictionary of Economics, Stockton Press, New York, 1987, pp. 997–999.
- [58] D. J. Watts, *Small Worlds: The Dynamics of Networks Between Order and Randomness*, Princeton University Press, Princeton, NJ, 1999. MR **2001a**:91064
- [59] D. J. Watts, S. H. Strogatz, *Collective dynamics of ‘small-world’ networks*, Nature **393** (1998), 440–442.
- [60] L. H. Weber, S. Z. El-Sayed, I. Hampton, *The variance spectra of phytoplankton, krill and water temperature in the Antarctic Ocean south of Africa*, Deep-Sea Res. **33** (1986), 1327–1343.

DEPARTMENT OF ECOLOGY AND EVOLUTIONARY BIOLOGY, PRINCETON UNIVERSITY, PRINCETON, NEW JERSEY 08544-1003

E-mail address: slevin@princeton.edu