

WEAK EMERGENCE*

Mark A. Bedau
Reed College

An innocent form of emergence—what I call “weak emergence”—is now a commonplace in a thriving interdisciplinary nexus of scientific activity—sometimes called the “sciences of complexity”—that include connectionist modelling, non-linear dynamics (popularly known as “chaos” theory), and artificial life.¹ After defining it, illustrating it in two contexts, and reviewing the available evidence, I conclude that the scientific and philosophical prospects for weak emergence are bright.

Emergence is a tantalizing topic because examples of apparent emergent phenomena abound. Some involve inanimate matter; e.g., a tornado is a self-organizing entity caught up in a global pattern of behavior that seems to be autonomous with respect to the massive aggregation of air and water molecules which constitute it. Another source of examples is the mind; our mental life consist of an autonomous, coherent flow of mental states (beliefs, desires, etc.) that presumably somehow ultimately arise out of the swarm of biochemical activity among our brain’s neurons. Life is a third rich source of apparent emergence. For example, the hierarchy of life embraces ecosystems composed of organisms, which are composed of organs, which are composed of cells, which are composed of molecules, but each level in this hierarchy exhibits behavior that seems autonomous with respect to the behavior found at the level below.

These examples highlight two admittedly vague but nevertheless useful hallmarks of emergent phenomena:

- (1) Emergent phenomena are somehow *constituted by*, and *generated from*, underlying processes.
- (2) Emergent phenomena are somehow *autonomous* from underlying processes.

If we place these hallmarks against a backdrop of abundant apparently emergent phenomena, it is clear why emergence is a perennial philosophical puzzle. At worst, the two hallmarks seem to make emergent phenomena flat-out inconsis-

tent. At best, they still raise the specter of illegitimately getting something from nothing.

So, aside from precisely defining what emergence is, any philosophical defense of emergence should aim to explain—ideally, explain away—its apparently illegitimate metaphysics. Another important goal should be to show that emergence is consistent with reasonable forms of materialism. But perhaps the most important goal should be to show that emergent properties are useful in empirical science, especially in accounts of those phenomena like life and mind that have always seemed to involve emergence. A defense of emergence will be secure only if emergence is more than merely a philosophical curiosity; it must be shown to be a central and constructive player in our understanding of the natural world.

I will argue that *weak emergence* (defined below) meets these three goals: it is metaphysically innocent, consistent with materialism, and scientifically useful, especially in the sciences of complexity that deal with life and mind. But first I will briefly illustrate the scientific irrelevance characteristic of stronger, more traditional conceptions of emergence.

Problems with Strong Emergence.

To glimpse the problems with stronger forms of emergence, consider the conception of emergence defended by Timothy O'Connor (1994). O'Connor's clearly articulated and carefully defended account falls squarely within the broad view of emergence that has dominated philosophy this century. His definition² is as follows: Property *P* is an emergent property of a (mereologically-complex) object *O* iff *P* supervenes on properties of the parts of *O*, *P* is not had by any of the object's parts, *P* is distinct from any structural property of *O*, and *P* has a direct ("downward") determinative influence on the pattern of behavior involving *O*'s parts.

The pivotal feature of this definition, to my mind, is the strong form of downward causation involved. O'Connor (pp. 97f) explains that he wants

to capture a very strong sense in which an emergent's causal influence is irreducible to that of the micro-properties on which it supervenes; it bears its influence in a direct 'downward' fashion, in contrast to the operation of a simple structural macro-property, whose causal influence occurs via the activity of the micro-properties which constitute it.

I call O'Connor's notion "strong" emergence to contrast it with the weaker form of emergence, defended below, that involves downward causation only in the weak form created by the activity of the micro-properties that constitute structural macro-properties.

It is worth noting that strong emergence captures the two hallmarks of emergence. Since emergent phenomena supervene on underlying processes, in this sense the underlying processes constitute and generate the emergent phenomena.

And emergent phenomena are autonomous from the underlying processes since they exert an irreducible form of downward causal influence. Nevertheless, strong emergence has a number of failings, all of which can be traced to strong downward causation.

Although strong emergence is logically possible, it is uncomfortably like magic. How does an irreducible but supervenient downward causal power arise, since by definition it cannot be due to the aggregation of the micro-level potentialities? Such causal powers would be quite unlike anything within our scientific ken. This not only indicates how they will discomfort reasonable forms of materialism. Their mysteriousness will only heighten the traditional worry that emergence entails illegitimately getting something from nothing.

But the most disappointing aspect of strong emergence is its apparent scientific irrelevance. O'Connor finds evidence that strong emergence is useful in the empirical sciences in "the recent proposals of macro-determinative influence on lower-level sub-structure by Polanyi and Sperry with respect to embryonic cells and consciousness, respectively" (p. 99). But these references to Polanyi and Sperry provide little evidence of the empirical viability of strong emergence unless they refer to a flourishing scientific research program. Our doubts about this should be raised when we note that in the recent philosophical literature on emergence (including O'Connor) all citations are to the *same* Polanyi and Sperry papers, which generally date back twenty five years. This is not the trail left by a thriving research program. Strong emergence is perhaps *compatible* with current scientific knowledge. But if Sperry and Polanyi are the best defense of strong emergence's empirical usefulness, then its scientific credentials are very weak. We should avoid proliferating mysteries beyond necessity. To judge from the available evidence, strong emergence is one mystery which we don't need.

Weak emergence contrasts with strong emergence in this respect; science apparently *does* need weak emergence. Fortunately, there are no mysteries like irreducible downward causation in weak emergence, to which we will now turn.

Definition of Weak Emergence.

Weak emergence applies in contexts in which there is a system, call it *S*, composed out of "micro-level" parts; the number and identity of these parts might change over time. *S* has various "macro-level" states (macrostates) and various "micro-level" states (microstates). *S*'s microstates are the intrinsic states of its parts, and its macrostates are structural properties constituted wholly out of its microstates.³ Interesting macrostates typically average over microstates and so compresses microstate information. Further, there is a microdynamic, call it *D*, which governs the time evolution of *S*'s microstates. Usually the microstate of a given part of the system at a given time is a result of the microstates of "nearby" parts of the system at preceding times; in this sense, *D* is "local". Given these assumptions, I define weak emergence as follows:

Macrostate P of S with microdynamic D is *weakly emergent* iff P can be derived from D and S 's external conditions but only by simulation.⁴

Conditions affecting the system's microstates are "external" if they are "outside" the system. If D is deterministic and the system is closed, then there is just one external condition: the system's initial condition. Every subsequent microstate of the system is determined by elements inside the system (the microdynamic D and the system's microstates). If the system is open, then another kind of "external" condition is the contingencies of the flux of parts and states through S . If the microdynamic is nondeterministic, then each accidental effect is an "external" condition. With external conditions understood in this fashion, it is coherent to speak of macrostates being "derivable" from external conditions even in nondeterministic systems.

Although perhaps unfamiliar, the idea of a macrostate being derived "by simulation" is straightforward and natural. Given a system's initial condition and the sequence of all other external conditions, the system's microdynamic completely determines each successive microstate of the system. To simulate the system one iterates its microdynamic, given a contingent stream of external conditions as input. Since the macrostate P is a structural property constituted out of the system's microstates, the external conditions and the microdynamic completely determine whether P materializes at any stage in the simulation. By simulating the system in this way one can derive from the microdynamic plus the external conditions whether P obtains at any given time after the initial condition. What distinguishes a weakly emergent macrostate is that this sort of simulation is required to derive the macrostate's behavior from the system's microdynamic. Crutchfield et al. (1986, p. 49) put the essential point especially clearly: the algorithmic effort for determining the system's behavior is roughly proportional to how far into the future the system's behavior is derived. It is obvious that the algorithmic effort required for a simulation is proportional to how far into the future the simulation goes.

The need for simulations in the study of low-dimensional chaos has been emphasized before (see, e.g., Crutchfield et al. 1986, Stone 1989, Kellert 1993). Weak emergence has a special source in this kind of chaos: exponential divergence of trajectories, also known as sensitive dependence on initial conditions or "the butterfly effect". This particular mechanism does not underlie all forms of weak emergence, though. On the contrary, weak emergence seems to rampant in *all* complex systems, regardless of whether they have the underlying mechanisms that produce chaos. In fact, some include weak emergence as part of the definition of what it is to be a complex adaptive system in general (Holland 1992). Indeed, it is the ubiquity of weak emergence in complex systems that makes weak emergence especially interesting.

Derivations that depend on simulations have certain characteristic limitations. First, they are massively contingent, awash with accidental information

about the system's components and external conditions. The derivations can be too detailed and unstructured for anyone to be able to "survey" or understand how they work. The derivations also can obscure simpler macro-level explanations of the same macrostates that apply across systems with different external conditions and different microdynamics. But none of this detracts from the fact that all of the system's macrostates can be derived from its microdynamic and external conditions with a simulation.

The modal terms in this definition are metaphysical, not epistemological. For *P* to be weakly emergent, what matters is that *there is* a derivation of *P* from *D* and *S*'s external conditions and *any* such derivation is a simulation. It does not matter whether anyone has discovered such a derivation or even suspects that it exists. If *P* is a weakly emergent, it is constituted by, and generated from, the system's underlying microdynamic, whether or not we know anything about this. Our need to use a simulation is due neither to the current contingent state of our knowledge nor to some specifically human limitation or frailty. Although a Laplacian supercalculator would have a decisive advantage over us in simulation speed, she would still need to simulate. Underivability without simulation is a purely formal notion concerning the existence and nonexistence of certain kinds of derivations of macrostates from a system's underlying dynamic.

Weak Emergence in the Game of Life.

A good way to grasp the concept of weak emergence is through examples. One of the simplest source of examples is the Game of Life devised more than a generation ago by the Cambridge mathematician John Conway and popularized by Martin Gardner.⁵ This "game" is "played" on a two-dimensional rectangular grid of cells, such as a checker board. Time is discrete. A cell's state at a given time is determined by the states of its eight neighboring cells at the preceding moment, according to the birth-death rule: A dead cell becomes alive iff 3 neighbors were just alive, and a living cell dies iff fewer than 2 or more than 3 neighbors were just alive. (Living cells with fewer than two living neighbors die of "loneliness", those with more than three living neighbors die of "overcrowding", and a dead cell becomes populated by a living cell if it has the three living neighbors needed to "breed" a new living cell.) Although Conway's Game of Life does not represent the state of the art of scientific attempts to understand complex systems, it is a well-known and exquisitely simple illustration of many of the principles of complexity science, including weak emergence, and it illustrates a *class* of systems—so called "cellular automata"—that are one central paradigm for how to understand complexity in general (see, e.g., Wolfram 1994).

One can easily calculate the time evolution of certain simple Life configurations. Some remain unchanging forever (so-called "still lifes"), others oscillate indefinitely (so-called "blinkers"), still others continue to change and grow indefinitely. Figure 1 shows seven time steps in the history of six small initial

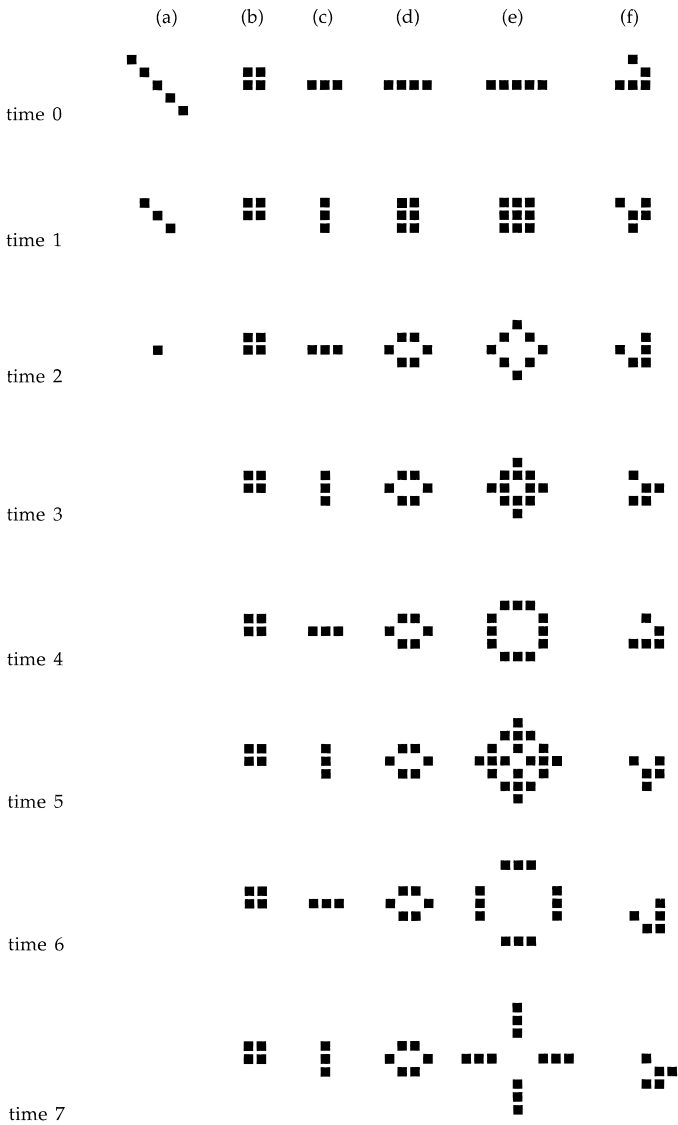


Figure 1. Seven time steps in the evolution of some simple configurations in the Game of Life. Configuration (a) is a “fuse” burning at both ends; after two time steps it is entirely consumed and no life remains. Configuration (b), a still life called the “block”, never changes. Configuration (c), a “traffic light”, is a blinker with period two. Configuration (d) evolves after two time steps into the “beehive,” another still life. Configuration (e) evolves after five time steps into a period two blinker consisting of four traffic lights. Configuration (f) is a glider, a period four pattern that moves diagonally one cell per period.

configurations of living cells; some are still lifes, others are blinkers. Examining the behavior of these initial configurations allows one to derive their exact behavior indefinitely far into the future. More complex patterns can also be produced by the simple birth-death rule governing individual cells. One simple and striking example—dubbed the “glider”, shown as (f) in Figure 1—is a pattern of five living cells that cycles through four phases, in the processes moving one cell diagonally across the Life field every four time steps. Some other notable patterns are “glider guns”—configuration that periodically emit a new glider—and “eaters”—configurations that destroy any gliders that collide with them. Clusters of glider guns and eaters can function in concert just like AND, OR, NOT, and other logic gates, and these gates can be connected into complicated switching circuits. In fact, Conway proved (Berlekamp et al. 1982) that these gates can even be cunningly arranged so that they constitute a universal Turing machine, and hence are able to compute literally every possible algorithm, or, as Poundstone vividly puts it, to “model every precisely definable aspect of the real world” (Poundstone 1985, p. 25).

There is no question that every event and pattern of activity found in Life, no matter how extended in space and time and no matter how complicated, is generated from the system’s microdynamic—the simple birth—death rule. Every event and process that happens at any level in a Life world can be deterministically derived from the world’s initial configuration of states and the birth-death rule. It follows that a structural macrostate in Life will be weakly emergent if deriving its behavior requires simulation. Life contains a vast number of macrostates that fill this bill. Some are not especially interesting; others are fascinating. Here are two examples.

R pentomino growth. The R pentomino is a wildly unstable five-cell edge-connected pattern. Figure 2 shows the first seven time steps in the evolution of the R pentomino; Figure 3 shows the pattern at time step 100 (above) and time step 150 (below). Listen to part of Poundstone’s description (1985, p. 33) of what the R pentomino produces: “One configuration leads to another and another and another, each different from all of its predecessors. On a high-speed computer display, the R pentomino roils furiously. It expands, scattering debris over the Life plane and ejecting gliders.”

Indefinite growth (i.e., increase in number of living cells) is a structural macrostate constituted by the cells in Life.⁶ Does the R pentomino (on an infinite Life grid) grow indefinitely? Some Life configurations do grow forever, such as glider guns, which continually spawn five-cell gliders that glide off into the indefinite distance. So, if the R pentomino continually ejects gliders that remain undisturbed as they travel into the infinite distance, for example, then it would grow forever. But does it? There is no simple way to answer this question. As far as anyone knows, all we can do is let Life “play” itself out when given the R pentomino as initial condition, i.e., observe the R pentomino’s behavior. As it happens (Poundstone 1985, p. 35), after 1103 time steps it settles down to a stable

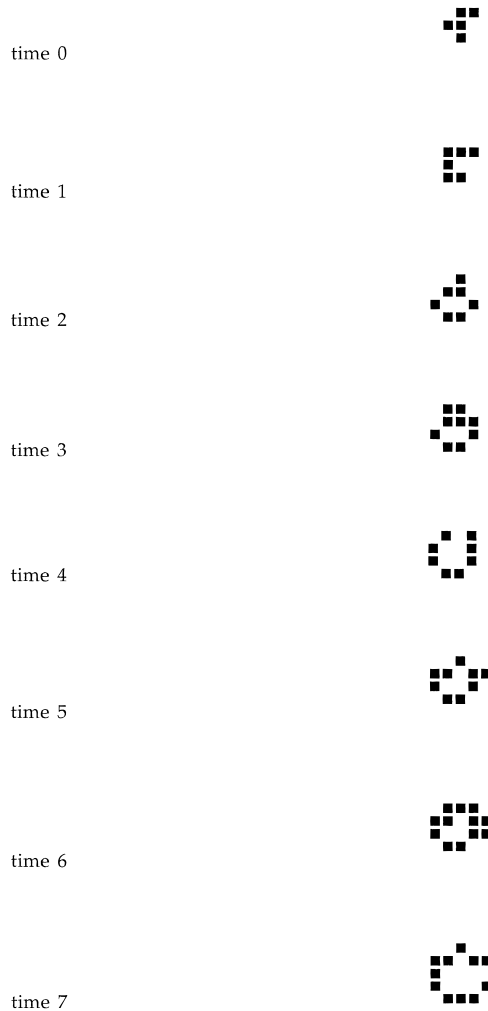


Figure 2. The first seven time steps in the evolution of the R pentomino (the figure at time 0), showing slow and irregular growth.

state that just fits into a 51-by-109 cell region. Thus, the finite bound of the R pentomino is a weak emergent macrostate of the Game of Life.

The R pentomino is one of the simplest Life configurations that is underivable. What makes Life’s underivability so striking is that its microdynamic—the underlying birth-death rule—is so simple.

Glider Spawning. Let *G* be the structural macrostate of quickly spawning a glider. (To make this property precise, we might define *G* as, say, the property of exhibiting a glider that survives for at least a three periods, i.e., twelve time steps,



Figure 3. Above: The R pentomino after 100 timesteps. The configuration contains five blocks, a traffic light, a glider, and some unstable clusters of cells. Below: The R pentomino after 150 timesteps. The configuration now includes three blocks, a traffic light, two gliders, and some unstable clusters of cells. The pattern continues to grow steadily but irregularly.

within one hundred time steps of evolution from the initial condition.) It is easy to derive that certain Life configurations never spawn a glider and so lack property *G*. As illustrations, a little a priori reflection allows one to derive that *G* is absent from each of the five the configurations in Figure 1 (a) - (e), from any configuration consisting of a sparse distribution of those five configurations, from a configuration consisting of all dead cells or all living cells, and from a configuration split straight down the middle into living and dead cells. Similarly, no

simulation is necessary to see that some Life configurations have G ; for example, consider the configuration consisting of one glider, Figure 1 (f). In general, though, it is impossible to tell whether a given initial Life configuration will quickly spawn a glider, short of observing how the initial condition evolves. Thus, G (or non- G) is weakly emergent in most of the Life configurations that possess (or lack) it, as contemplating a couple of examples makes evident. Figures 4 and 5 show two random initial configurations (above) and their subsequent evolution (below). By timestep 115 the configuration in Figure 4 has spawned no gliders, while by timestep 26 a glider has already emerged from the pattern in Figure 5.

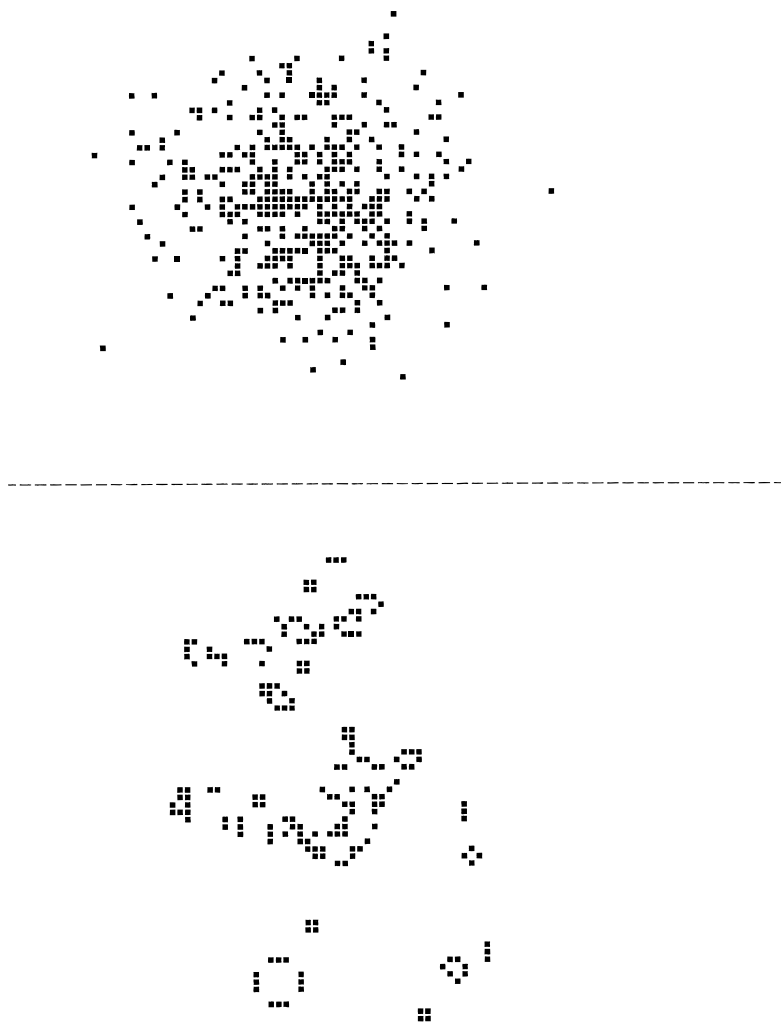


Figure 4. Above: A random distribution of living cells. Below: The distribution after 115 timesteps. No glider has appeared yet.

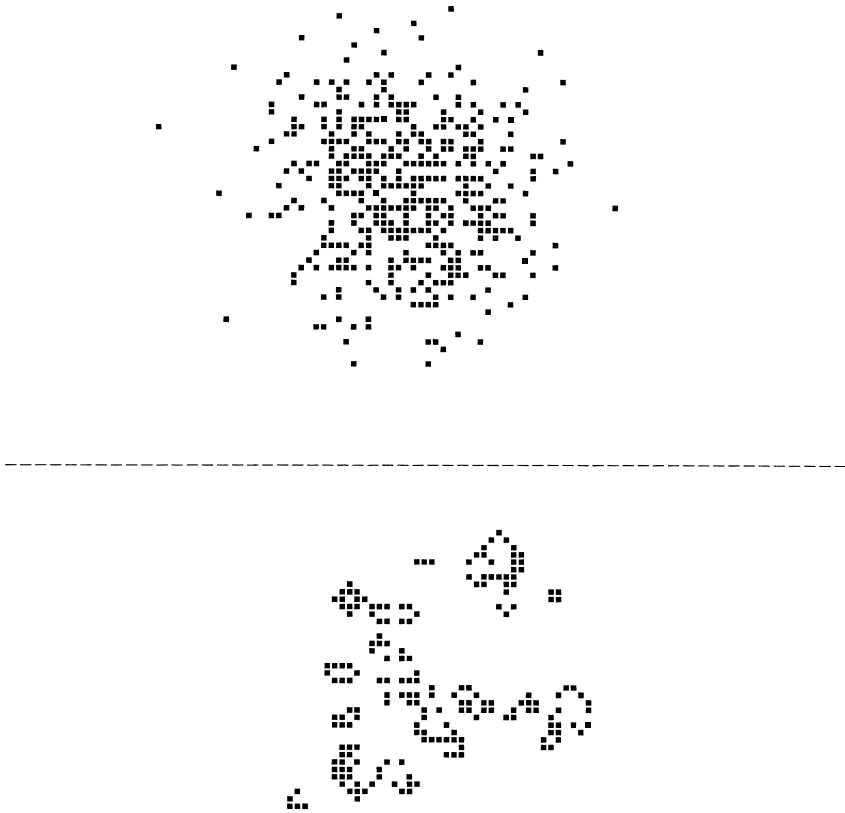


Figure 5. Above: A random distribution of living cells. Below: The distribution after 26 timesteps. A glider is emerging from an unstable cluster of cells at the lower left.

Being weakly emergent does not prevent us from readily discovering various laws involving G . If one observes the frequency of occurrence of gliders in lots of random initial configurations, one discovers that usually gliders are quickly spawned; G is true of most random Life fields. Extensive enough observation allows one to measure the prevalence of G quite accurately, and this information can then be summarized in a little probabilistic law about all random Life fields X , of this form: $\text{prob}(X \text{ is } G) = k$.

Although perhaps not especially fascinating or profound, this little law of the Game of Life nicely illustrates how empirical observation of computer simulations can unearth evidence for laws involving the Game of Life's weakly emergent states.

Empirical observation is generally the *only* way to discover these laws. With few exceptions, it is impossible without simulation to derive the behavior of any macrostate in a Life configuration even given complete knowledge of the con-

figuration. In fact, since a universal Turing machine can be embedded in Life, the undecidability of the halting problem proves that in principle there can be no algorithm for determining whether the behavior exhibited in an arbitrary Life world will ever stabilize. Yet all Life phenomena can be derived from the initial conditions and the birth-death rule. Thus, Conway's Game of Life abounds with weakly emergent properties.

The Game of Life is an exceptionally simple system, simpler than many systems studied in the sciences of complexity. For example, recent artificial life work brims with weak emergence. I will present one illustration involving the emergence of evolvability. Although not as simple as the Game of Life, this next illustration will be more typical of current work in the sciences of complexity.

Weak Emergence in a Model of Evolving Life.

Evolving life forms display various macro-level patterns on an evolutionary time scale. For example, advantageous traits that arise through mutations tend, *ceteris paribus*, to persist and spread through the population. Furthermore, organisms' traits tend, within limits and *ceteris paribus*, to adapt to changing environmental contingencies. These sorts of supple dynamics of adaptation result not from any explicit macro-level control (e.g., God does not adjust allele frequencies so that creatures are well adapted to their environment); rather, they emerge statistically from the micro-level contingencies of natural selection.

Norman Packard devised a simple model of evolving sensorimotor agents which demonstrates how these sorts of supple, macro-level evolutionary dynamic can emerge implicitly from an explicit microdynamical model (Packard 1989, Bedau and Packard, 1992; Bedau, Ronneburg, and Zwick, 1992; Bedau and Bahm, 1993 and 1994; Bedau 1994; Bedau and Seymour, 1994; Bedau 1995). What motivates this model is the view that evolving life is typified by a population of agents whose continued existence depends on their sensorimotor functionality, i.e., their success at using local sensory information to direct their actions in such a way that they can find and process the resources they need to survive and flourish. Thus, information processing and resource processing are the two internal processes that dominate the agents' lives, and their primary goal—whether they know this or not—is to enhance their sensorimotor functionality by coordinating these internal processes. Since the requirements of sensorimotor functionality may well alter as the context of evolution changes, continued viability and vitality requires that sensorimotor functionality can adapt in an open-ended, autonomous fashion. Packard's model attempts to capture an especially simple form of this open-ended, autonomous evolutionary adaptation.

The model consists of a finite two-dimensional world with a resource field and a population of agents. An agent's survival and reproduction is determined by the extent to which it finds enough resources to stay alive and reproduce, and an agent's ability to find resources depends on its sensorimotor functionality—that is, the way in which the agent's perception of its contingent local environment

affects its behavior in that environment. An agent's sensorimotor functionality is encoded in a set of genes, and these genes can mutate when an agent reproduces. Thus, on an evolutionary time scale, the process of natural selection implicitly adapts the population's sensorimotor strategies to the environment. Furthermore, the agents' actions change the environment because agents consume resources and collide with each other. This entails that the mixture of sensorimotor strategies in the population at a given moment is a significant component of the environment that affects the subsequent evolution of those strategies. Thus, the "fitness function" in Packard's model—what it takes to survive and reproduce—is constantly buffeted by the contingencies of natural selection and unpredictably changes (Packard 1989).

All macro-level evolutionary dynamics produced by this model ultimately are the result of an explicit micro-level microdynamic acting on external conditions. The model explicitly controls only local micro-level states: resources are locally replenished, an agent's genetically encoded sensorimotor strategy determines its local behavior, an agent's behavior in its local environment determines its internal resource level, an agent's internal resource level determines whether it survives and reproduces, and genes randomly mutate during reproduction. Each agent is autonomous in the sense that its behavior is determined solely by the environmentally-sensitive dictates of its own sensorimotor strategy. On an evolutionary time scale these sensorimotor strategies are continually refashioned by the historical contingencies of natural selection. The aggregate long-term behavior of this microdynamic generates macro-level evolutionary dynamics only as the indirect product of an unpredictably shifting agglomeration of directly controlled micro-level events (individual actions, births, deaths, mutations). Many of these evolutionary dynamics are weakly emergent; although constituted and generated solely by the micro-level dynamic, they can be derived only through simulations. I will illustrate these emergent dynamics with some recent work concerning the evolution of evolvability (Bedau and Seymour 1994).

The ability to adapt successfully depends on the availability of viable evolutionary alternatives. An appropriate quantity of alternatives can make evolution easy; too many or too few can make evolution difficult or even impossible. For example, in Packard's model, the population can evolve better sensorimotor strategies only if it can "test" sufficiently many sufficiently novel strategies; in short, the system needs a capacity for evolutionary "innovation." At the same time, the population's sensorimotor strategies can adapt to a given environment only if strategies that prove beneficial can persist in the gene pool; in short, the system needs a capacity for evolutionary "memory."

Perhaps the simplest mechanism that simultaneously affects both memory and innovation is the mutation rate. The lower the mutation rate, the greater the number of genetic strategies "remembered" from parents. At the same time, the higher the mutation rate, the greater the number of "innovative" genetic strategies introduced with children. Successful adaptability requires that these competing demands for memory and innovation be suitably balanced. Too much mutation

(not enough memory) will continually flood the population with new random strategies; too little mutation (not enough innovation) will tend to freeze the population at arbitrary strategies. Successful evolutionary adaptation requires a mutation rate suitably intermediate between these extremes. Furthermore, a suitably balanced mutation rate might not remain fixed, for the balance point could shift as the context of evolution changes.

One would think, then, that any evolutionary process that could continually support evolving life must have the capacity to adapt automatically to this shifting balance of memory and innovation. So, in the context of Packard's model, it is natural to ask whether the mutation rate that governs *first-order* evolution could adapt appropriately by means of a *second-order* process of evolution. If the mutation rate can adapt in this way, then this model would yield a simple form of the evolution of evolvability and, thus, might illuminate one of life's fundamental prerequisites.

Previous work (Bedau and Bahm 1993, 1994) with *fixed* mutation rates in Packard's model revealed two robust effects. The first effect was that the mutation rate governs a phase transition between genetically "ordered" and genetically "disordered" systems. When the mutation rate is too far below the phase transition, the whole gene pool tends to remain "frozen" at a given strategy; when the mutation rate is significantly above the phase transition, the gene pool tends to be a continually changing plethora of randomly related strategies. The phase transition itself occurs over a critical band in the spectrum of mutation rates, μ , roughly in the range $10^{-3} \leq \mu \leq 10^{-2}$. The second effect was that evolution produces maximal population fitness when mutation rates are around values just below this transition. Apparently, evolutionary adaptation happens best when the gene pool tends to be "ordered" but just on the verge of becoming "disordered."

In the light of our earlier suppositions about balancing the demands for memory and innovation, the two fixed-mutation-rate effects suggest the *balance hypothesis* that the mutation rates around the critical transition between genetic "order" and "disorder" optimally balance the competing evolutionary demands for memory and innovation. We can shed some light on the balance hypothesis by modifying Packard's model so that each agent has an additional gene encoding its personal mutation rate. In this case, two kinds of mutation play a role when an agent reproduces: (i) the child inherits its parent's sensorimotor genes, which mutate at a rate controlled by the parent's personal (genetically encoded) mutation rate; and (ii) the child inherits its parent's mutation rate gene, which mutates at a rate controlled by a population-wide meta-mutation rate. Thus, first-order (sensorimotor) and second-order (mutation rate) evolution happen simultaneously. So, if the balance hypothesis is right and mutation rates at the critical transition produce optimal conditions for sensorimotor evolution because they optimally balance memory and innovation, then we would expect second-order evolution to drive mutation rates into the critical transition. It turns out that this is exactly what happens.

Figure 6 shows four examples of how the distribution of mutation rates in the population change over time under different conditions. As a control, distribu-

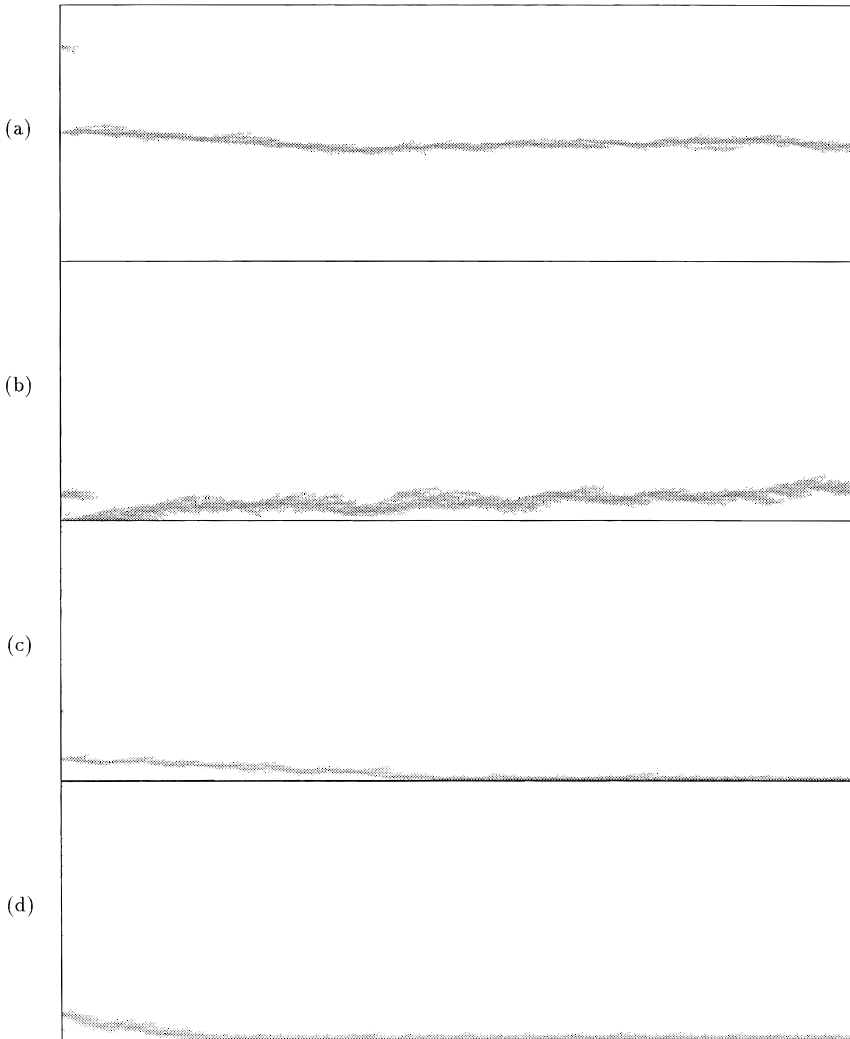


Figure 6. Evolutionary dynamics in mutation rate distributions from four simulations of the model of sensorimotor agents. Time is on the X-axis (100,000 timesteps) and mutation rate is on the Y-axis. The gray-scale at a given point (t, m) in this distribution shows the frequency of the mutation rate m in the population at time t . See text.

tions (a) and (b) show what happens when the mutation rate genes are allowed to drift randomly: the bulk of the distribution wanders aimlessly. By contrast, distributions (c) and (d) illustrate what happens when natural selection affects the mutation rate genes: the mutation rates drop dramatically. The meta-mutation rate is lower in (a) than in (b) and so, as would be expected, distribution (a) is

narrower and changes more slowly. Similarly, the meta-mutation rate is lower in (c) than in (d), which explains why distribution (c) is narrower and drops more slowly.

If we examine lots of simulations and collect suitable macrostate information, we notice the pattern predicted by the balance hypothesis: Second-order evolution tends to drive mutation rates down to the transition from genetic disorder to genetic order, increasing population fitness in the process. This pattern is illustrated in Figure 7, which shows time series data from a typical simulation. The macrostates depicted in Figure 7 are (from top to bottom): (i) the mutation rate distribution, as in Figure 6; (ii) a blow up distinguishing very small mutation rates in the distribution (bins decrease in size by a factor of ten, e.g., the top bin shows mutation rates between 10^{-0} and 10^{-1} , the next bin down shows mutation rates between 10^{-1} and 10^{-2} , etc.); (iii) the mean mutation rate (note the log scale); (iv) the uningested resources in the environment; (v) three aspects of the genetic diversity in the population's sensorimotor strategies; and (vi) the population level.

The composite picture provided by Figure 7 can be crudely divided into three epochs: an initial period of (relatively) high mutation rates, during the time period 0 - 20,000; a transitional period of falling mutation rates, during the time period 20,000 - 40,000; and a final period of relatively low mutation rates, throughout the rest of the simulation. The top three time series are different perspectives on the falling mutation rates, showing that the mutation rates adapt downwards until they cluster around the critical transition region, $10^{-3} \leq \mu \leq 10^{-2}$. Since resources flow into the model at a constant rate and since survival and reproduction consume resources, the uningested resource inversely reflects the population fitness. We see that the population becomes more fit (i.e., more efficiently gathers resources) at the same time as the mutation rates drop. Although this is not the occasion to review the different ways to measure the diversity of the sensorimotor strategies in the population, we can easily recognize that there is a significant qualitative difference between the diversity dynamics in the initial and final epochs. In fact, these qualitative differences are characteristic of precisely the difference between a "disordered" gene pool of randomly related strategies and a gene pool that is at or slightly below the transition between genetic order and disorder (see Bedau and Bahm 1993, 1994, Bedau 1995).

If the balance hypothesis is the correct explanation of this second-order evolution of mutation rates into the critical transition, then we should be able to change the mean mutation rate by dramatically changing where memory and innovation are balanced. And, in fact, the mutation rate *does* rise and fall along with the demands for evolutionary innovation. For example, when we randomize the values of all the sensorimotor genes in the entire population so that every agent immediately "forgets" all the genetically stored information learned by its genetic lineage over its entire evolutionary history, the population must restart its evolutionary learning job from scratch. It has no immediate need for memory (the gene pool contains no information of proven value); instead, the need for inno-

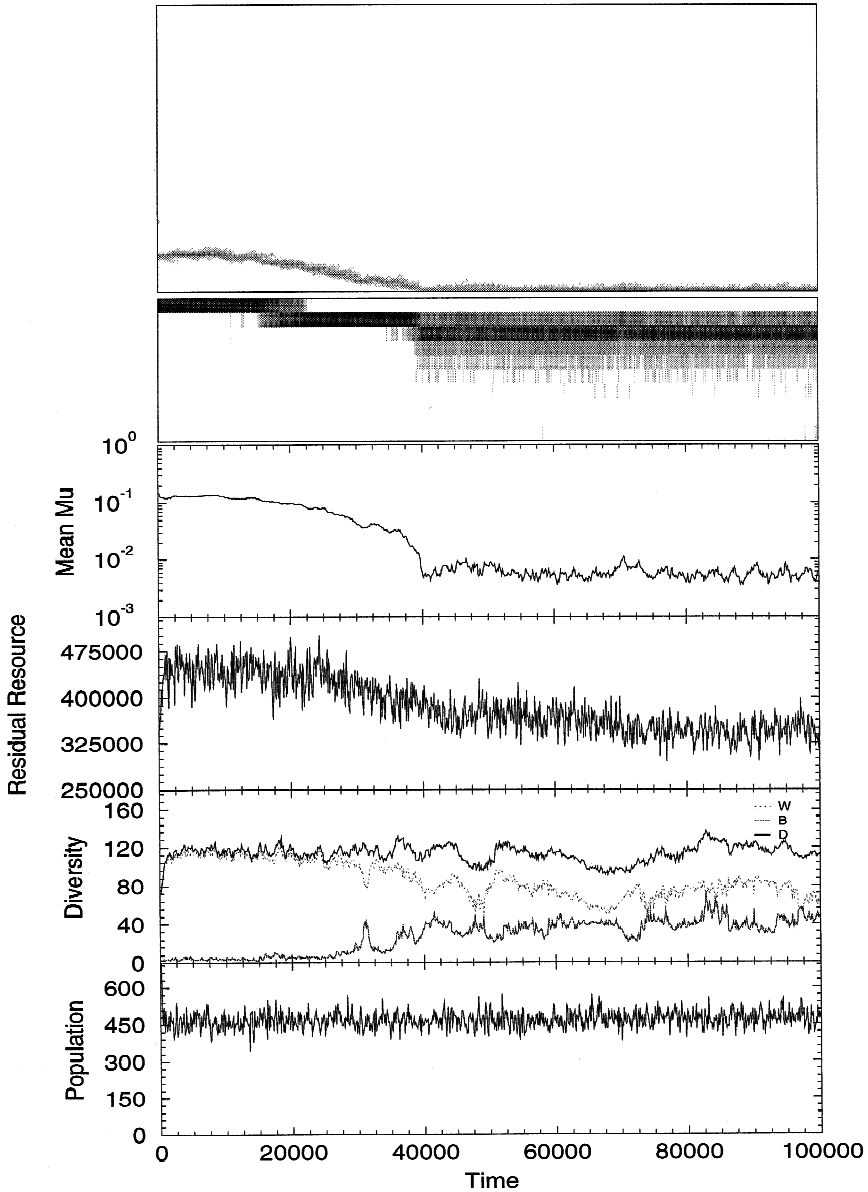


Figure 7. Time series data from a simulation of the model of sensorimotor agents, showing how the population's resource gathering efficiency increases when the mutation rates evolve downward far enough to change the qualitative character of the population's genetic diversity. From top to bottom, the data are: (i) the mutation rate distribution; (ii) a blow up of very small mutation rates; (iii) the mean mutation rate (note the log scale); (iv) the uningested resource in the environment; (v) three aspects of the diversity of the sensorimotor strategies in the population; (vi) the population level. See text.

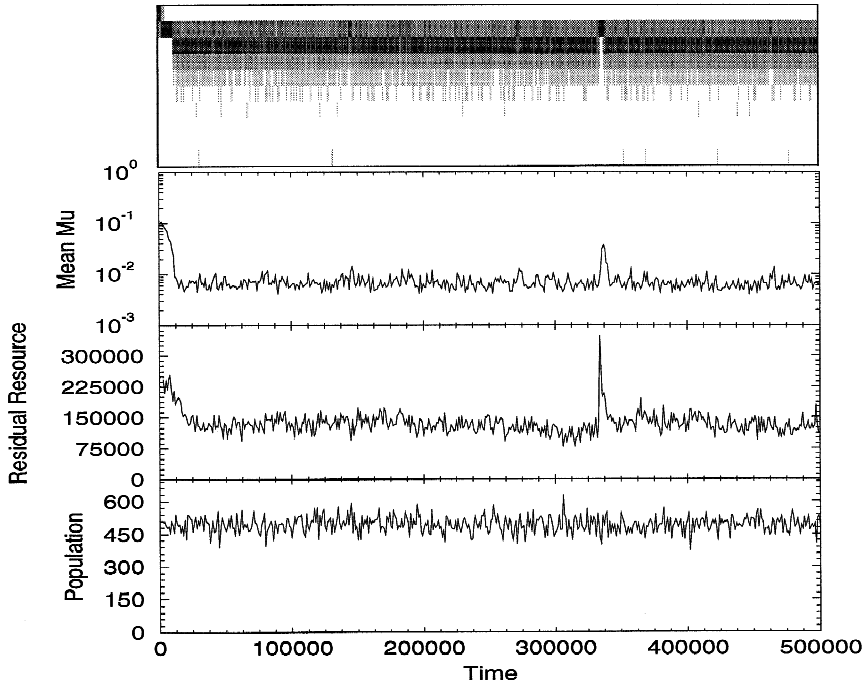


Figure 8. Time series data from a simulation of the model of sensorimotor agents. From top to bottom, the data are: (i) a blow up of very small mutation rates in the mutation rate distribution; (ii) mean mutation rate (note the log scale); (iii) the level of uningested resources in the world; (iv) population level. At timestep 333,333 all sensorimotor genes of all living organisms were randomly scrambled. See text.

vation is paramount. Under these conditions, we regularly observe the striking changes illustrated around timestep 333,333 in Figure 8. The initial segment (time-steps 0 - 100,000) in Figure 8 shows a mutation distribution evolving into the critical mutation region, just as in Figure 7 (but note that the time scale in Figure 8 is compressed by a factor of five). But at timestep 333,333 an “act of God” randomly scrambles all sensorimotor genes of all living organisms. At just this point we can note the following sequence of events: (a) the residual resource in the environment sharply rises, showing that the population has become much less fit; (b) immediately after the fitness drop the mean mutation rate dramatically rises as the mutation rate distribution shifts upwards; (c) by the time that the mean mutation rate has risen to its highest point the population’s fitness has substantially improved; (d) the fitness levels and mutation rates eventually return to their previous equilibrium levels.

All of these simulations show the dynamics of the mutation rate distribution adjusting up and down as the balance hypothesis would predict. Temporarily perturbing the context for evolution can increase the need for rapid exploration of

a wide variety of sensorimotor strategies and thus dramatically shift the balance towards the need for innovation. Then, subsequent sensorimotor evolution can reshape the context for evolution in such a way that the balance shifts back towards the need for memory. This all suggests that, *ceteris paribus*, mutation rates adapt so as to balance appropriately the competing evolutionary demands for memory and innovation, and that, *ceteris paribus*, this balance point is at the genetic transition from order to disorder. An indefinite variety of environmental contingencies can shift the point at which the evolutionary need for memory and innovation are balanced, and the perturbation experiments show how mutation rates can adapt up or down as appropriate.

This sort of supple adaptability in Packard's model can be counted among the hallmarks of life in general (Maynard Smith 1975, Cairns-Smith 1985, Bedau 1996). And, clearly, these evolutionary dynamics are weakly emergent. The model's macro-level dynamic is wholly constituted and generated by its micro-level phenomena, but the micro-level phenomena involve such a kaleidoscopic array of non-additive interactions that the macro-level dynamics cannot be derived from micro-level information except by means of simulations, like those shown above. In a similar fashion, many other characteristic features of living systems can be captured as emergent phenomena in artificial life models; see, e.g., Farmer *et al.* (1986), Langton (1989), Langton *et al.* (1992), Varela and Bourgine (1992), Brooks and Maes (1994), Gaussier and Nicoud (1994), Stonier and Yu (1994), Banzhaf and Eeckman (1995).

Support for Weak Emergence.

Conway's Game of Life and Packard's model of evolving life serve to clarify weak emergence and illustrate its role in the sciences of complexity. But one might still ask whether weak emergence is philosophically interesting and, indeed, whether it deserves the name "emergence" at all. These questions deserve answers, especially since weak emergence differs significantly from traditional twentieth century accounts of emergence.

For example, since weakly emergent properties can be derived (via simulation) from complete knowledge of micro-level information, from that information they can be *predicted*, at least in principle. If we have been observing a simulation of some system *S* and at time *t* we saw that *S* was in state *P*, then we know that there is an appropriate derivation that *S* will be in macrostate *P* at *t*.⁷ So, if we are given a system's microdynamic and all relevant external conditions, then in principle we can derive the system's behavior because we can simulate the system and observe its behavior for as long as necessary. And if we can derive how the system will behave, we can predict its future behavior with complete certainty. Thus, on this key issue weak emergence parts company with at least the *letter* of those traditional conceptions of emergence (e.g., Broad 1925, Pepper 1926, Nagel 1961) that focus on in principle unpredictability of macrostates even given complete microstate information.

At the same time, weak emergence does share much of the *spirit* of those traditional views that emphasize unpredictability. For one thing, in the case of open systems, making the prediction would require prior knowledge of all details of the flux of accidental changes introduced by contact with the external world; and in the case of nondeterministic systems, it would require knowledge of all the non-deterministic events affecting the system's behavior. This sort of knowledge is beyond us, except "in principle;" so weak emergent macrostates of such systems are predictable *only* "in principle." Furthermore, even for closed and deterministic systems, weak emergent macrostates can be "predicted" *only* by observing step-by-step how the system's behavior unfolds. For example, one can "predict" whether an R pentomino will grow forever only by observing in time what happens to the configuration. Some might find this so unlike what should be expected of a prediction that they would agree with Stone (1989) that it is no prediction at all.

One might worry that the concept of weak emergence is fairly useless since we generally have no *proof* that a given macrostate of a given system is underivable without simulation.⁸ For example, I know no proof that the unlimited growth of the R pentomino and glider-spawning probability can be derived only by simulation; for all I know there *is* no such proof. On these grounds some might conclude that weak emergence "suffers in the course of application in practice", to use Klee's words (1984, p. 49). I would strenuously disagree, however, since unproven weak emergence claims can, and often do, still possess substantial *empirical* support. My earlier weak emergence claims about R pentomino growth and random glider spawning, although unproved, still have more than enough empirical support. Similar weak emergence claims have substantial empirical support. A significant part of the activity in artificial life consists of examining empirical evidence about interesting emergent phenomena in living systems; *mutatis mutandis*, the same holds for the rest of the sciences of complexity.

One might object that weak emergence is *too* weak to be called "emergent", either because it applies so widely or arbitrarily that it does not demark an interesting class of phenomena, or because it applies to certain phenomena that are not emergent. For example, indefinitely many arbitrary, ad hoc Life macrostates are (for all we know) underivability without simulation. Or, to switch to a real world example, even though the potentiality of a certain knife to slice a loaf of bread is "not the sort [of macrostate] emergence theorists typically have in mind" (O'Conner 1994, p. 96), the knife's potentiality might well be weakly emergent with respect to its underlying molecular microdynamic. But this breadth of instances, including those that are arbitrary or uninteresting to "emergence theorists", is not a problem or flaw in weak emergence. Weak emergence explicates an everyday notion in complexity science. It is not a special, intrinsically interesting property; rather, it is widespread, the rule rather than the exception. So not all emergent macrostates are interesting; far from it. A central challenge in complexity science is to identify and study those exceptional, especially interesting weak emergent macrostates that reflect fundamental aspects of complex systems and are amenable to empirical investigation. Simulation gives us a new capacity to

identify and study important macrostates that would otherwise be beyond the reach of more traditional mathematical or empirical methods.

The micro-level derivability of weak emergent phenomena might be thought to deprive them of the right to be called “emergent”; they might not seem “emergent” enough. The impetus behind this worry might come partly from the history of emergence concepts being ineliminably and unacceptably mysterious—as if no acceptable and non-mysterious concept could deserve to be called “emergence.” By contrast, part of my defense of weak emergence is precisely that it avoids the traditional puzzles about emergence.

In any event, there are good reasons for using the word “emergence” in this context. For one thing, complexity scientists themselves routinely use this language and weak emergence is an explication of their language.⁹ Another compelling reason for allowing the “emergence” language is that weak emergence has the two hallmarks of emergent properties. It is quite straightforward how weak emergent phenomena are constituted by, and generated from, underlying processes. The system’s macrostates are constituted by its microstates, and the macrostates are entirely generated solely from the system’s microstates and microdynamic. At the same time, there is a clear sense in which the behavior of weak emergent phenomena are autonomous with respect to the underlying processes. The sciences of complexity are discovering simple, general macro-level patterns and laws involving weak emergent phenomena. There is no evident hope of side-stepping a simulation and deriving these patterns and laws of weak emergent phenomena from the underlying microdynamic (and external conditions) alone. In fact, as I emphasized earlier, the micro-level “explanations” of weak emergence are typically so swamped with accidental micro-details that they obscure the macro-level patterns. In general, we can formulate and investigate the basic principles of weak emergent phenomena only by empirically observing them at the macro-level. In this sense, then, weakly emergent phenomena have an autonomous life at the macro-level. Now, there is nothing inconsistent or metaphysically illegitimate about underlying processes constituting and generating phenomena that can be derived only by simulation. In this way, weak emergence explains away the appearance of metaphysical illegitimacy.

It is also clear why weak emergence is consistent with reasonable forms of materialism. By definition, a weak emergent property can be derived from its microdynamic and external conditions. Any emergent phenomenon that a materialist would want to embrace would have materialistic micro-level components with materialist micro-properties governed by a materialistic microdynamic. Thus, the weak emergent phenomena of interest to the materialists would have a completely materialistic explanation.

Conclusion.

Weak emergence is no universal metaphysical solvent. For example, if (hypothetically, and perhaps *per impossible*) we were to acquire good evidence that

human consciousness is weakly emergent, this would not immediately dissolve all of the philosophical puzzles about consciousness. Still, we *would* learn the answers to some questions: First, a precise notion of emergence *is* involved in consciousness; second, this notion of emergence is metaphysically benign. Thus, free from special distractions from emergence, we could focus on the remaining puzzles just about consciousness itself.

As Conway's Game of Life and Packard's model of evolving sensorimotor agents illustrate, weak emergence is ubiquitous in the burgeoning, interdisciplinary nexus of scientific research about complex systems. The central place of weak emergence in this thriving scientific activity is what provides the most substantial evidence that weak emergence is philosophically and scientifically important. It is striking that weak emergence is so prominent in scientific accounts of exactly those especially puzzling phenomena in the natural world—such as those involving life and mind—that perennially generate sympathy for emergence. Can this be an accident?

Notes

- * Many thanks for comments on earlier drafts and discussion on these topics to Hugo Bedau, John Collier, Seth Crook, Dick Grandy, Mark Hinchliff, Terry Horgan, Stephen Leeds, Norman Packard, Teresa Robertson, Bill Uzgalis, and to audiences at the Third Workshop on Artificial Life (in Santa Fe during summer 1993), at Oregon State University (spring 1994), and at the Northwest Philosophy Conference (fall 1994). This research was supported with the help of an external fellowship from the Oregon Center for the Humanities and a research grant from the Oregon Council for the Humanities.
1. Accessible introductions to the study of chaos, with references to more technical treatments, include Crutchfield et al. (1986), Gleick (1987), and Kellert (1993). The bible of connectionism is Rumelhart and McClelland (1986); discussions for philosophers, and references to the technical literature, can be found in Bechtel and Abrahamsen (1990), Horgan and Tienson (1991), and Ramsey, Stich, and Rumelhart (1991). The locus of much recent activity in the "sciences of complexity" is the Santa Fe Institute, a private, independent multidisciplinary research center. Semi-popular introductions to some of the research centered at the Santa Fe Institute include Levy (1992), Lewin (1992), and Waldrop (1992). A representative range of technical work can be found in the series Santa Fe Institute Studies in the Sciences of Complexity, published by Addison-Wesley; e.g., Langton (1989) and Langton et al. (1992).
 2. O'Connor adapts Kim's notion of "strong supervenience" (Kim 1990) and Armstrong's definition of structural property (Armstrong 1978).
 3. The macrostate *P* might fall into a variety of categories. It might be a property of *S*, possibly one involving various other macrostates of *S*; it might be some phenomenon concerning *S*, possibly involving a variety of *S*'s other macrostates; it might be a pattern of *S*'s behavior, possibly including other macrostates of *S*. There are also more complicated cases, in which the macrostate is "supple" or "fluid", and the structural definition of the macrostate might be infinitely long. This latter issue is developed in Bedau (forthcoming).

4. This definition is explicitly restricted to a given macrostate of a given system with a given microdynamic. This is the *core* or *focal* notion in a family of related notions of weak emergence, all others of which would be defined by reference to the core notion and would crucially invoke underivability without simulation. For example, one can speak of a weak emergent *law* when it is a law that a given macrostate of a given system with a given microdynamic is weakly emergent from a range of initial conditions; this law is underivable without simulations across many initial conditions. Similarly, one can speak of a weak emergent *pattern* involving a range of suitably related macrostates, microdynamics, or systems, but I will not attempt here to define weak emergence in this whole family of contexts. The guiding strategy behind these definitional extensions is reasonably clear. The range of new contexts for weak emergence is limited only by our imagination.

It is worth at least mentioning that the notion of underivability without simulation provides another dimension along which notions of weak emergence can vary. There is a range of more or less stringent conditions. For example, consider a macrostate that in principle is derivable without simulation, yet the derivation uses vastly more resources (e.g., “steps”) than any human could grasp; or consider a macrostate that is derivable (only) by simulation but the simulation is infinitely long. I will not elaborate on this dimension here. The paradigm of weak emergence involves underivability except by finite simulation.

5. See Berlekamp et al. (1982) and Gardner (1983). An excellent introduction to the intellectual delights of Conway’s Game of Life is Poundstone (1985).
6. Specifically, indefinite growth is the macrostate defined as the (infinite) disjunction of all those (infinite) sequences s of life states such that, for each positive integer n , there is a time t when s contains more than n living cells.
7. This can be spelled out as follows: Let C_i be the set of microstates of all the parts of S at time i . Apply D (possibly with nondeterministic steps) to the S ’s initial condition C_0 (and possibly include a property synchronized sequence of external conditions) through successor conditions C_i until D yields C_t . From C_t and the structural definition of P , determine whether P obtains at t .
8. It is a mathematical fact whether a given macrostate of a given system is underivable from the system’s microdynamics and external conditions. So, unless it’s undecidable, it’s provable. Nevertheless, being provable does not entail that it is easy, or even humanly possible, to find and evaluate the proof.
9. Even if we adopt the quite simplistic expedient of restricting our attention to the *titles* of research reports, we can easily generate a rich range of examples of this language. E.g., rummaging for a few minutes in a handful of books within easy reach produced the following list, all of which speak of emergence in the weak sense defined here in their titles: “Emergent Colonization in an Artificial Ecology” (Assad and Packard 1992), “Concept Formation as Emergent Phenomena” (Patel and Schnepf 1992), “A Behavioral Simulation Model for the Study of Emergent Social Structures” (Drogoul et al. 1992), “Dynamics of Artificial Markets: Speculative Markets and Emerging ‘Common Sense’ Knowledge” (Nottola, Leroy, and Davalo 1992), *Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks* (Forrest 1989), “Emergent Frame Recognition and its Use in Artificial Creatures” (Steels 1991), “The Coreworld: Emergence and Evolution of Cooperative Structures in a Computational Chemistry”

(Rasmussen, Knudsen, and Feldberg 1991), "Spontaneous Emergence of a Metabolism" (Bagley and Farmer 1992).

References

- Armstrong, D. 1978. *Universals and Scientific Realism—Vol. II: A Theory of Universals*. Cambridge: Cambridge University Press.
- Assad, A. M., and N. H. Packard. 1992. "Emergent Colonization in an Artificial Ecology." In Varela and Bourgine (1992), pp. 143-152.
- Bagley, R. J., and J. D. Farmer. 1992. "Spontaneous Emergence of a Metabolism." In C. Langton *et al.* (1992), pp. 93-140.
- Banzhaf, W., and Eeckman, F. eds. 1995. *Evolution and Biocomputation—Computational Models of Evolution*. Berlin: Springer.
- Bechtel, W. and A. Abrahamsen. 1990. *Connectionism and the Mind : An Introduction to Parallel Processing in Networks*. Cambridge, Mass.: B. Blackwell, 1990.
- Bedau, M. A. 1994. "The Evolution of Sensorimotor Functionality." In Gaussier and Nicoud (1994), pp. 134-145.
- Bedau, M. A. 1995. "Three Illustrations of Artificial Life's Working Hypothesis." In Banzhaf and Eeckman (1995), pp. 53-68.
- Bedau, M. A. 1996. "The Nature of Life." In M. Boden (1996), pp. 332-357.
- Bedau, M. A. forthcoming. "Emergent Models of Supple Dynamics in Life and Mind." *Brain & Cognition*.
- Bedau, M. A., and Bahm, A. 1993. "Order and Chaos in the Evolution of Diversity." In *Proceedings of the Second European Conference on Artificial Life (ECAL'93), Brussels, Belgium*, pp. 51-62.
- Bedau, M. A., and Bahm, A. 1994. "Bifurcation Structure in Diversity Dynamics." In Brooks and Maes (1994), pp. 258-268.
- Bedau, M. A., Bahm, A., and Zwick, M. 1995. "Variance and Uncertainty Measures of Population Diversity Dynamics." *Advances in Systems Science and Applications*, 1, 7-12.
- Bedau, M. A., Giger, M., and Zwick, M. 1995. "Evolving Diversity of Population and Environment in Static Resource Models." *Advances in Systems Science and Applications*, 1, 1-6.
- Bedau, M. A., and Packard, N. 1992. "Measurement of Evolutionary Activity, Teleology, and Life." In C. Langton *et al.* (1992), pp. 431-461.
- Bedau, M. A., Ronneburg, F., and Zwick, M. 1992. "Dynamics of Diversity in a Simple Model of Evolution." *Parallel Problem Solving from Nature 2*: 94-104.
- Bedau, M. A. and Seymour, R. 1994. "Adaptation of Mutation Rates in a Simple Model of Evolution." In Stonier and Yu (1994), pp. 37-44.
- Berlekamp, E.R., J.H. Conway, and R.K. Guy. 1982. *Winning Ways*. Vol. 2. New York: Academic Press.
- Boden, M. 1996. *The Philosophy of Artificial Life*. New York: Oxford University Press.
- Broad, C. D., 1925. *The Mind and Its Place in Nature*. London: Routledge and Kegan Paul.
- Brooks, R. and P. Maes, eds. 1994. *Artificial Life IV*. Cambridge, Mass.: MIT Press.
- Cairns-Smith, A. G. 1985. *Seven Clues to the Origin of Life*. Cambridge: Cambridge University Press.
- Crutchfield, J.P., J.D. Farmer, N.H. Packard, and R. S. Shaw. 1986. "Chaos." *Scientific American* 255 (December): 46-57.
- Drogoul, A., J. Ferber, B. Corbara, and D. Fresneau. 1992. "A Behavioral Simulation Model for the Study of Emergent Social Structures." In Varela and Bourgine (1992), pp. 161-170.
- Farmer, J. D., Lapedes, A., Packard, N., and Wendroff, B., eds. 1986. *Evolution, Games, and Learning: Models for Adaptation for Machines and Nature*. Amsterdam: North Holland.
- Forrest, S., ed. 1989. *Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks*. Amsterdam: North-Holland.
- Gardner, M. 1983. *Wheels, Life, and Other Mathematical Amusements*. New York: Freeman.

- Gaussier, P., and Nicoud, J. -D., eds. 1994. *From Perception to Action*. Los Alamitos, Calif.: IEEE Computer Society Press.
- Gleick, James. 1987. *Chaos: Making a New Science*. New York: Viking.
- Holland, J. H. 1992. *Adaptation in Natural and Artificial Systems*, 2nd edition. Cambridge, Mass: MIT Press.
- Horgan, T. and J. Tienson, eds. 1991. *Connectionism and the Philosophy of Mind*. Dordrecht: Kluwer Academic.
- Kellert, S. H. 1993. *In the Wake of Chaos: Unpredictable Order in Dynamical Systems*. Chicago: The University of Chicago Press.
- Kim, J. 1984. "Concepts of Supervenience." *Philosophy and Phenomenological Research* 45: 153-176.
- Klee, R. 1984. "Micro-Determinism and Concepts of Emergence." *Philosophy of Science* 51: 44-63.
- Langton, C., ed. 1989. *Artificial Life*. SFI Studies in the Sciences of Complexity, Vol. VI. Redwood City, Calif.: Addison-Wesley.
- Langton, C., C. E. Taylor, J. D. Farmer, S. Rasmussen, eds. 1992. *Artificial Life II*. SFI Studies in the Sciences of Complexity, Vol. X. Reading, Calif.: Addison-Wesley.
- Levy, S. 1992. *Artificial Life: The Quest for a New Creation*. New York: Pantheon Books.
- Lewin, R. 1992. *Complexity: Life at the Edge of Chaos*. New York: Macmillan.
- Maynard Smith, J. 1975. *The Theory of Evolution*, 3rd edition. New York: Penguin.
- Nagel, E., 1961. *The Structure of Science*. New York: Harcourt, Brace & World.
- Nottola, C., F. Leroy, and F. Davalo. 1992. "Dynamics of Artificial Markets: Speculative Markets and Emerging 'Common Sense' Knowledge." In Varela and Bourguine (1992), pp. 185-194.
- O'Conner, T. 1994. "Emergent Properties." *American Philosophical Quarterly* 31: 91-104.
- Packard, N. 1989. "Intrinsic Adaptation in a Simple Model for Evolution." In Langton (1989), pp. 141-155.
- Patel, M., and U. Schnepf. 1992. "Concept Formation as Emergent Phenomena." In Varela and Bourguine (1992), pp. 11-20.
- Pepper, S., (1926), "Emergence." *Journal of Philosophy* 23: 241-245.
- Poundstone, W. 1985. *The Recursive Universe*. Chicago: Contemporary Books.
- Ramsey, W., S. P. Stich, and D. E. Rumelhart, eds. 1991. *Philosophy and Connectionist Theory*. Hillsdale, N.J.: L. Erlbaum Associates.
- Rasmussen, S., C. Knudsen, and R. Feldberg. 1991. "The Coreworld: Emergence and Evolution of Cooperative Structures in a Computational Chemistry." *Physica D* 42: 121-30.
- Rumelhart, D. E., J. L. McClelland, and the PDP Research Group. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Cambridge, Mass.: MIT Press.
- Steels, L. 1991. "Emergent Frame Recognition and its Use in Artificial Creatures." In *Proceedings of the 10th IJCAI*. San Mateo, Calif.: Morgan Kaufmann.
- Stone, M. A. 1989. "Chaos, Prediction, and Laplacean Determinism." *American Philosophical Quarterly* 26: 123-131.
- Stonier, R. and X. H. Yu, eds. 1994. *Complex Systems—Mechanisms of Adaptation*. Amsterdam: IOS Press.
- Varela, F., and P. Bourguine. 1992. *Towards a Practice of Autonomous Systems*. Cambridge, Mass.: MIT Press.
- Waldrop, M. M. 1992. *Complexity: The Emerging Science at the Edge of Order and Chaos*. New York: Simon & Schuster.
- Wolfram, S. 1994. *Cellular Automata and Complexity*. Reading, Mass.: Addison-Wesley.