

How Are the Sciences of Complex Systems Possible?

Michael Strevens

Forthcoming, *Philosophy of Science*

ABSTRACT

To understand the behavior of a complex system, you must understand the tangled web of interactions between the system's many parts. The web is tractable in systems that are somewhat decomposable, meaning that the interactions influence only weakly the short-term behavior of the parts. But how to handle complexity in systems that are nowhere near decomposable? Science's principal tool for dealing with non-decomposable systems is a variety of probabilistic analysis that I call EPA, exemplified by statistical physics and population genetics.

I show that EPA's power to deal with complex bundles of interactions derives from an assumption that appears to be false of non-decomposable complex systems, in virtue of their very non-decomposability. Yet EPA is extremely successful. I aim to find an interpretation of EPA's assumption that is consistent with, indeed that explains, its success.

1. INTRODUCTION

How is it possible to understand the behavior of extremely complex systems? The *prima facie* obstacle to such understanding is the superabundance of convoluted interactions that collectively determine a complex system's dynamics. A science of complex systems must have some way of comprehending this causal jumble, of making it transparent, so that light can shine through the chaos to illuminate whatever dynamic principles lie at a system's core. The sciences of complex systems are possible, then, only if there are such underlying principles, and only if methods exist to discern them.

Past attempts to provide a foundation for the sciences of complex systems have pursued the possibility of managing complexity by a two step process that resolves a complex system into discrete components, and then

1. Gives a theory of the components' individual behavior, and
2. Gives a theory of the interaction of these behaviors.

The hope is that the theory of individual behavior extracted in the first step will be simple enough to survive the daunting combinatorics of complexity introduced in the second step.¹

Under certain circumstances, this hope is realized. When systems are decomposable or near-decomposable, meaning that the short-term behavior of each of a system's components can be understood largely independently of the behavior of the other components, the effects of the combinatorial explosion are muted or altogether absent. But many complex systems are, it is generally conceded, not decomposable at all, and are therefore unassailable by decomposition's divide-and-conquer approach. The aim of this paper is to investigate a way of negotiating such systems' complexity.

1. On the decomposition approach see, for example, Simon (1969), Wimsatt (1986), Bechtel and Richardson (1993), and Auyang (1998), chaps. 4–6. Chapter two of Bechtel and Richardson is especially helpful.

2. NON-DECOMPOSABLE COMPLEX SYSTEMS

My inquiry will focus on a particular class of non-decomposable complex systems, systems made up of many fairly independent, but strongly interacting, parts. By *strongly interacting* I mean that the parts sometimes interact in ways that have potentially radical consequences for the interactors. By *fairly independent* I mean that for any given part, such interactions are relatively rare—most of the time, a part does its own thing without being affected significantly by the other parts. By *many*, I mean enough parts that in the population as a whole, strong interactions are common over even relatively short periods of time.

Two important kinds of systems that appear to be complex in this sense are the systems studied by statistical physics, such as gases, and ecosystems. In a gas, the many parts are the molecules. At any given instant, two molecules are unlikely to affect one another's direction of travel, speed, mode of vibration, and so on. But when they do interact—that is, when they collide—they stand to change completely all of these properties.

In a typical ecosystem, the many parts are the animals. Like molecules, the animals for the most part chart their own course, largely independently of what is going on around them. But when they do interact, the effect may be life-altering: a meeting often leads to a mating or an eating. (In neither gases nor ecosystems, note, is interaction all or nothing: gas molecules exert a slight force on one another at all times, and animals may monitor one another's movements.)

From this point on, I refer to the class of systems just described as the *complex systems*. Not everything that you would intuitively count as a complex system belongs to this class; I will, therefore, be using *complex system* as a term of art.

The sort of question I am about to ask, whether a given complex system is decomposable, turns not only on the nature of the system itself, but also on the object of the investigation. To give a trivial example: for the purpose

of predicting its total mass, just about any system counts as decomposable, since the mass of each of the system's parts can be determined independently, and these individual masses summed to obtain the total mass.

Let me therefore specify an explanatory goal, namely, to understand the *macrolevel behavior* of a complex system in terms of the behavior of its parts. What is the macrolevel? When you view a complex system's behavior at the macrolevel, you observe changes in *statistics* concerning the system's parts, while ignoring changes in the properties of individual parts. In population ecology, you observe changes in population number without taking note of the fates of individual animals; in sociology you observe suicide rates without noticing who, exactly, is killing themselves; in statistical physics you observe changes in, say, temperature, hence in average heat energy, without noticing the energies of individual molecules.

My topic is the scientific endeavor of predicting and explaining these kinds of behaviors of gases, ecosystems, and societies, and more particularly, the endeavor of understanding them as consequences of the behavior of the molecules, organisms, and people of which they are composed.

What I want to investigate, in other words, is the way in which science predicts and explains macrolevel behavior by deriving it from *microlevel behavior*, by which I mean the individual behavior of a system's parts. This is not the only form that complex-system science can take; often, for example, complex-system scientists model macrolevel behavior without any concern for its microlevel basis. Equally, this is not the only question that can be asked about the relation between a complex system's micro- and macrolevels. I will say nothing here about, for example, questions of supervenience and reduction.

How, then, do scientists parlay their understanding of the dynamics of a complex system's individual components into an understanding of its large-scale dynamics?

Let me begin by showing why the strategy of decomposition cannot pro-

vide explanations and predictions of macrolevel behavior in complex systems. Decomposition requires some grip on the short term behavior of the individuals that are a system's parts. To characterize this behavior, the external influences on an individual must be taken into account. In a complex system, these influences include, potentially, every other individual in the system. To characterize the behavior of one individual, then, the behaviors of all the other individuals must be characterized already, at least in the aggregate. Such characterization cannot, it appears, proceed one individual at a time without circularity. Rather, the cumulative effect of all of a system's interactions must be taken into account in a single theoretical leap. But the number of interactions is so great that the leap may seem an impossible feat.

If a system is decomposable or near-decomposable, the leap is not such a challenge after all. This is because the influences between individuals in decomposable systems are by definition always weak. As a consequence, external influences on an individual have a significant effect only over a relatively long period of time. If you focus on an individual's short-term behavior, then, you can either ignore the influence of other individuals altogether, or you can approximate their aggregate influence in the form of a static background condition, such as a "mean field" in physics, a condition whose state you can calculate with little or no knowledge of the other individuals' dynamics. (The justification for the static approximation is that the influence is weak and that the nuances of the influence—its patterns of change over short time periods, or the directions it comes from—therefore make little difference to the influenced individual's behavior.) This constitutes the first of the two steps in the strategy of decomposition. In the second step, short term individual behaviors are aggregated to yield short term macrolevel behaviors. It is then possible to understand long term macrolevel behavior as a stringing together of the short term macrolevel behaviors derived in this way.

Observe that, above all else, what makes the decomposition strategy vi-

able is the possibility of characterizing the short term behavior of the parts of a system without taking into account the details of their surroundings, and in particular, without taking into account the *short term dynamics* of their surroundings.

The complex systems in which I am interested are not decomposable in this sense: the behavior of any of their parts over even relatively short stretches of time depends greatly on the individual behavior of the other parts. This follows from the definition of complexity alone: a system's complexity, in my proprietary sense, entails frequent interactions between parts with potentially radical effect. The events that occur over a week in the life of a given rabbit, for example—including perhaps death—will depend very sensitively on the rabbit's encounters with foxes and other rabbits, therefore on the behavior of the foxes and other rabbits. Consideration of a rabbit's behavior in isolation or against an unchanging background will tell you very little, then, about the course of a real rabbit's life. For this reason, the decomposition strategy cannot be applied.

It does not follow that there can be no derivation of complex systems' behavior from the behavior of their parts. It does follow that any such derivation will have to take into account, right from the start, the cumulative effect of the strong interactions between the parts. It is a technique for assessing this cumulative effect that will make the sciences of complex systems possible.

3. ENION PROBABILITY ANALYSIS

There is a well-known method for dealing with the many strong interactions of a complex system's parts, so as to achieve a derivation of macrolevel behavior from microlevel behavior, and thus both the prediction and the explanation of macrolevel facts using microlevel facts. The method is introduced in this section; I begin to examine its foundations in the next. The

foundational enquiry unearths a problem: the method ought to fail, because its central assumption about complex systems, indispensable for the derivation of macrolevel from microlevel dynamics, appears to be false. Yet it is rather successful in many different areas, in physics, biology, and elsewhere, both as a predictive and as an explanatory technique. To understand why is the main business of this paper.

The method I have in mind has no name; I will call it *enion probability analysis*, or EPA. *Enion* is my catch-all term for the individual, relatively independent, strongly interacting parts of a complex system: the organisms in an ecosystem, the molecules in a gas, the people in a social system. It is not a technical term, then, but rather a way of referring efficiently to those parts of a system in virtue of which it qualifies as complex in my proprietary sense.

The essence of EPA is to assign a probability distribution over the behaviors of each of a complex system's individual enions. If the distributions are assumed to have certain properties—and this is the central assumption, the apparently *false* central assumption, mentioned a few sentences ago—then they will combine in a straightforward way to yield a macrolevel dynamics for the system. In this way, EPA provides an understanding of macrolevel dynamics in terms of a probabilistically characterized microlevel dynamics.²

Imagine, for example, that you wish to track the population of rabbits in an ecosystem. Suppose that any given rabbit has a 25% probability of being eaten in the course of a month. Suppose further that any rabbit that is not eaten will, if female, produce on average two offspring. Finally, suppose that there are many rabbits, and that the births and deaths of the rabbits are stochastically independent. Then you can conclude from the fact that the probability of death is 25% that, in the course of a month, about a quarter

2. For more on the history, form, and assumptions of EPA, see Strevens (2003), §1.2. Also to be found in Strevens (2003) are defenses, extensions, exemplifications, qualifications, and proofs of many of this paper's key claims.

of the rabbits will die. Of the remainder, about half will be female and will therefore produce about two offspring. Thus, the remaining population will double, for a net increase of 50%. From a few simple probabilistic assumptions about patterns in the lives of individual rabbits, you have derived a mathematically precise law for the dynamics of the rabbit population: the population increases by 50% every month.

This is, of course, a very simple ecosystem, as are the others considered in this paper. In what follows, I will confine my attention to complex systems that are relatively uncluttered, such as ecosystems containing only two species, predator and prey. Despite their lack of complications, these systems raise the same fundamental problem for complex-system science as their real world counterparts: they involve frequent strong interactions between the parts even in the very short term, and so the behavior of the whole cannot be understood by the decompositional two-step described above. I will not be able to prove to you that the same considerations that provide a foundation for the application of EPA to “uncluttered” systems do the same for all complex systems; what I offer is a beginning, not an end, to the study of the foundations of the sciences of non-decomposable systems.

Yet EPA seems to work almost as well in the messy real world as in the sparser environment of idealized models. The range of EPA includes, in addition to population ecology, population genetics—the mathematical engine of modern evolutionary theory—and statistical physics, starting with Maxwell’s seminal work on kinetic theory in the 1860s, as well as much work on economics and society. Its utility as a route to the macrolevel in the science of complex systems, with and without clutter, is, then, enormous.³

How does EPA cut through the complexity of non-decomposable sys-

3. Systems of greater complexity are discussed in any advanced text on, in the case of gases, kinetic theory, and in the case of ecosystems, population/evolutionary ecology. The probabilistic assumptions of EPA are found in advanced as well as elementary treatments. You might begin with Loeb (1934) on gases and Roughgarden (1979) on evolutionary ecology.

tems? To answer this question, consider EPA's end product. What EPA turns out is a purely macrolevel law, that is, a robust generalization in which the macrolevel description of a system at one point in time always bears a determinate mathematical relation to the macrolevel description at an earlier point in time. In the nicest cases—far more common than you would expect—the relation is deterministic, or virtually so. (In other cases, the relation is probabilistic, and so manifests itself only in the behavior of a number of systems of the same kind.)

Applying EPA to ecosystems, for example, normally produces generalizations that relate populations at later times to populations at earlier times. Applying it to gases produces generalizations that relate earlier and later density, temperature, entropy. In each case, information about a system's earlier macrostate alone is sufficient, given the law, to predict its later macrostate.

(It is important to see that the claim that a system's behavior can be characterized by a purely macrolevel law is much stronger and more interesting than the claim that a system's behavior can be described at the macrolevel. The latter claim is trivially true for any system: you can always take a series of macrolevel snapshots of a system's state as it changes over time, but it does not follow that the resulting statistics exhibit any discernible pattern. Note also that the existence of a purely macrolevel dynamics is what makes possible a kind of complex-system science more shallow than EPA-driven science, namely, the empirical modeling of a complex system's macrolevel behavior without regard for its microlevel foundations.)

The derivability of purely macrolevel laws is, perhaps, quite unexpected. Because of the frequent, strong interactions between a complex systems' many parts, small differences—differences apparent only in a microlevel description—can make large differences to the fate of any given enion. You might well have thought, then, that a system's macrolevel state at a given point in time, that is, the collective fate of the system's enions at that time, will depend not only on the previous macrolevel state, but on the way that

the macrolevel state is realized at the microlevel. For some reason, however, the details of the microlevel realization make no apparent difference to the way the system behaves.

That microlevel complications are in a certain sense irrelevant to macrolevel behavior is the key to understanding how EPA makes a predictive and explanatory science of complex systems' macrolevel behavior possible. The prediction and explanation of macrolevel phenomena are within reach when a system's macrolevel dynamics can be derived from aspects of its microlevel dynamics. Given non-decomposability, you might think that such a derivation, though possible in principle, is intractable. But EPA's probabilistic apparatus allows a perspicuous representation of the way in which complex tangles of strong microlevel interactions cancel out, leaving a purely macrolevel dynamics. In so doing, it provides the means for a relatively simple derivation of macrolevel dynamics.

To talk of microlevel interactions "cancelling out" is, however, to put things too crudely. It is not that the system behaves as though the strong interactions were absent, or rabbits would never get eaten, foxes never mate, and so on. Think of it, rather, as follows (further details may be found in section 4). Any probabilistic process can be resolved conceptually and mathematically into two component processes:

1. A deterministic process producing a certain statistic, the value of that statistic being determined by the form of the relevant probability distribution, and
2. A set of fluctuations affecting this process, and therefore causing perturbations in the value of the deterministically produced statistic.

The tossing of a coin, for example, can be thought of as a deterministic process that produces exactly one half heads (one half because, of course, the probability of heads is one half), together with a stochastic process that produces an overlying set of fluctuations, causing the frequency of heads to

deviate to some, usually small, degree from one half. It is as though (but only *as though*) nature takes careful aim at the target of 50% heads, but due to unsteady execution, almost never hits the target dead on.

It is important to note that these “component processes” are not themselves real physical processes—in fact they do not correspond to anything real at all. Nor is their role as producers of individual outcomes well defined: it is not clear how, exactly, the deterministic process produces the statistic on an outcome-by-outcome basis (by producing regular or irregular patterns of outcomes?), and it is therefore equally unclear how the fluctuations affect the production of individual outcomes.

But none of this matters, for two reasons. First, these notional processes are not conceived as producers of individual outcomes, only as producers of a final macrolevel statistic (compare mathematical models in the social sciences that are deterministic except for a probabilistic “error term”). Second, the breakdown into two processes is introduced purely for expository purposes, so as to give you a sense of the way in which microlevel complexity disappears. A more formal treatment of the disappearance need make no use of the “two process” picture—see Strevens (2003).

By representing the microlevel dynamics of a complex system probabilistically, EPA makes it possible to think about the behavior of complex systems in the same way as the behavior of the tossed coin. Enion statistics—that is, macrolevel states—are notionally produced by, first, a deterministic process that generates statistics fixed by what you might call the *enion probabilities*, the probabilities characterizing the behavior of individual enions, and second, a set of fluctuations perturbing the outcome of the deterministic process.

For example, the number of rabbit deaths in a month is notionally produced by a deterministic process that in isolation would kill exactly one quarter of all rabbits every month (since the relevant enion probability, the probability of a rabbit dying the course of a month, is 25%), together with

a stochastic process that makes the actual death rate fluctuate around this mean.

This *determinism plus fluctuations* way of thinking is especially common in evolutionary biology, where natural selection is conceived as deterministic and genetic drift as fluctuation (Sober 1984), but it is apt wherever EPA can be applied. (It was especially influential in the nineteenth century, with the two sub-processes usually accorded physical reality; see Strevens (2003), §1.14.)

Now, it is a fact that, the more outcomes are produced by the probabilistic process, the smaller the likely net effect of the fluctuations, proportionally speaking. The more you toss a coin, for example, the closer the frequency of heads approaches, on average, one half. This is the *law of large numbers*. Apply the law to complex systems, and you find that for systems that have many enions, and that therefore produce many outcomes—that is, enion behaviors—over even relatively short time periods, the fluctuations cancel out; they have a negligible impact on enion statistics.

All that is relevant to understanding macrolevel behavior, then, are the quantities that drive the deterministic process. These are the values of the enion probabilities, the probabilities attached to behaviors of individual enions, such as the probability of rabbit death.

At this stage EPA makes a significant assumption, that the enion probabilities, although they concern microlevel events, have values that depend only on macrolevel quantities, such as the number of rabbits and the number of foxes in an ecosystem. Change at the macrolevel therefore depends only on macrolevel quantities; hence, macrolevel behavior can be characterized by purely macrolevel laws.

While strong interactions make a difference to the fates of individual enions in a complex system, then, those same interactions do not make a difference to a system's macrolevel behavior. The interactions cannot be ignored, but they can be averaged away, because what they contribute above

and beyond their contribution to the average is just a certain fluctuation which, in the long run, disappears. In short, the complexity introduced by strong microlevel interaction simply drops out of the big picture; at the macrolevel, the deterministic production of macrolevel statistics by other macrolevel statistics holds sway. It is by taking advantage of the disappearance of microlevel complexity that EPA makes tractable the derivation of the macrolevel laws, and so makes possible the sciences of complex systems.

The power of EPA is beyond question. Its foundations are, by contrast, rather mysterious. You can make a number of probabilistic posits about a system and show that those posits entail a certain macrolevel dynamics, but you do not thereby understand the dynamics unless you understand the foundation of the posits—unless you understand, that is, why they are true.

Yet they are, it appears, false.

4. EPA SHOULD FAIL

Enion probability analysis makes two crucial, wide-ranging probabilistic assumptions in order to achieve its derivation of macrolevel laws from microlevel behavior. The assumptions, here phrased so as to apply to a rabbit ecosystem, are as follows:

1. The probability assigned to rabbit death depends only on macrolevel facts about the ecosystem, such as the total number of foxes. (In the example in section 3, the same probability of death was assigned in every circumstance, trivially satisfying the assumption.)
2. The deaths of different rabbits are stochastically independent.

Together these requirements make up what I call the *probabilistic supercondition*. You can see the supercondition manifested in the Maxwell-Boltzmann probability distribution, in Boltzmann's assumption of "molecular chaos" in kinetic theory, and in the assumption that the passing down of genes from

one generation to the next can be modeled by a Bernoulli distribution (the same kind of distribution that applies to coin-tossing) in population genetics.

It is the supercondition that licenses EPA, in its derivation of macro-level behavior, to assume that microlevel complexities “cancel out” in the sense described above. To see this, recall that EPA conceives of the process by which enion behaviors are produced as a probabilistic process, and so as being composed (mathematically, but not physically) of a deterministic sub-process producing statistics that depend on the values of the enion probabilities, and a probabilistic sub-process creating fluctuations in those statistics. Because a complex system has many enions, the law of large numbers can be invoked to conclude that the fluctuations will cancel out. Part (2) of the supercondition plays its role at this point, because application of the law of large numbers requires stochastic independence.

With the fluctuations gone, the dynamics of a system’s enion statistics depends only on what affects the deterministic sub-process, therefore, only on what affects the values of the enion probabilities. Part (1) of the supercondition specifies that only macrolevel facts affect the probabilities, and hence that the dynamics of enion statistics—that is, macrolevel behavior—depends only on macrolevel variables. It follows that the laws of macrolevel behavior will make reference only to macrolevel facts; the microlevel’s causal complexity is as good as gone.

Let me explain what parts (1) and (2) assert, beginning with part (1). To assume that the probability of rabbit death depends only on macrolevel facts is to assume that it is unaffected by low level facts about the ecosystem. Earlier, I said that whether or not a rabbit dies depends on many small details about the initial state of the relevant ecosystem; if the system is chaotic, in the sense that it is very sensitive to initial conditions (the most basic, if not the most subtle, of the various definitions of chaos), then perhaps these details are very, very small. Enion probability analysis assumes that, despite

this, it is reasonable to assign a probability for death that does not depend on any low level details at all. The minutiae of the microlevel are somehow subsumed into the probabilistic haze.

You might think that this cannot possibly be correct, since different probabilities of death ought to be assigned to different rabbits depending on their level of health, their age, and so on. Quite right. There is space within the framework of EPA, however, to make probability assignments that depend on these sorts of microlevel details; they will not, in practice, significantly affect your ability to derive macrolevel behavior from the enion probabilities (Strevens 2003, §4.6). What the probabilities for death absolutely must not depend on are details such as the initial positions of the rabbits and foxes.

On to part (2). To say that rabbit deaths are stochastically independent is to say that the probability that one rabbit dies, given that some other rabbit dies, is just the unconditional probability for death: 25% in section 3's scenario. The probability that your favorite rabbit Hazel dies, for example, will not be impacted, according to the independence assumption, by the fate of any other rabbit, or even the fates of all the other rabbits combined. It is 25% no matter what. Rabbit deaths, then, are assumed to be like tosses of a fair coin: no matter how many heads you may have gotten recently, the probability of heads on the next toss remains one half; likewise, no matter how many other rabbits live or die in the course of a month, the probability that Hazel dies during that same month is one quarter.

You will observe that the two parts of the probabilistic supercondition say the same sort of thing: they say that the probability of rabbit death during a month is unaffected by certain low level details. Part (1) says that the probability of death is unaffected by the details of the state of the system at the beginning of the month; part (2) says that the probability of death is unaffected by the happenings in the system over the course of the month.

This observation reveals how EPA is able to brush off so easily the manifold claims of the microlevel. The probabilistic dynamics that EPA attributes to individual enions in a system—such as the probability of death ascribed to rabbits—is simply assumed to determine the fate of the enions independently of any microlevel details. Aggregate the probabilistic dynamics of a system’s individual enions and you get a dynamics of enion populations—a macrolevel dynamics—that inherits the probabilistic enion dynamics’ independence of the microlevel. The probabilistic supercondition, in other words, constitutes a kind of rewriting of a complex system’s microlevel dynamics without any of the microlevel interdependence described above.

Surely this ought not to be allowed! What was to be proved has simply been assumed, that when taking a statistical view of a system, the details of its microlevel dynamics can be ignored.

Even worse, the supercondition seems not only question-begging, but false. The supercondition’s part (1) supposes that the probability of rabbit death is independent of position. But of two rabbits, alike in every relevant respect—health, strength, age, and the rest—one may die simply because it started out the day in the wrong place. By what right does EPA ignore such factors? They can make just as much of a difference between survival and death as do health, strength, and age.

Part (2) assumes that deaths are stochastically independent, but if causal independence is taken as the foundation of stochastic independence, this too looks dubious: what happens to one rabbit over the course of a month is determined by a series of ecological events, many of which involve, and most of the rest of which indirectly affect, other rabbits.

In short, EPA’s probabilistic rewriting of the microlevel dynamics ignores microlevel interconnections that are clearly present and that can easily make a difference, on their own, to the values of macrolevel variables.

The founders of EPA were not unperplexed by this. Maxwell wrote that his methods worked well provided that one avoided making any “personal

enquiries” of molecules (Garber et al. 1995, 422). Can a policy of social discretion save the day?

My goal in what follows is to show that the assumption of the probabilistic supercondition is justified, not by an appeal to ignorance or unknowability, but because of a certain property of complex systems that is responsible for microlevel chaos. Chaos, then, apparently an impediment to simple, tractable macrolevel behavior, turns out to be its principal cause.

Some fellow travelers: my attempt to justify the supercondition is of a piece with work in modern ergodic theory seeking to show that the independence assumptions and other features of the probability distributions used in statistical mechanics are reasonable, given the special dynamics of statistical mechanical systems (Ornstein and Weiss 1991). Other approaches to the foundations of statistical mechanics reject the independence assumptions but try to show how such assumptions, even if false, can be a part of an empirically adequate theory (Sklar 1993). For more on the relation between my project and the foundations of statistical mechanics, see Strevens (2003).

5. INDEPENDENCE FROM INITIAL CONDITIONS

Is the supercondition true or false? I will show that, provided the notion of probability is judiciously defined—and what is judicious when reasoning about complex systems may make no sense elsewhere (see note 4 below)—the supercondition is true, or close enough to true. It is not the definition itself that makes the supercondition true, however, so much as the chaotic physical properties of complex systems.

The first part of the supercondition, to be discussed in this section, asserts that enion probabilities are independent of the microlevel details of a system’s initial conditions, and so that these details can be safely ignored in a probabilistic derivation of the system’s macrolevel dynamics. The demonstration has two parts: a purely philosophical move, and a physical investi-

gation of enion probabilities. First, the philosophical move.

5.1 *The Definition of Enion Probability*

I want to show that the probability that a particular rabbit dies over the course of a month does not depend on microlevel facts such as the rabbit's position at the beginning of the month. You might think that I have no hope of doing this, since whether or not the rabbit dies is clearly determined in part by its starting position.

I circumvent this initial problem by adopting a definition on which enion probability is a kind of statistical property. For simplicity's sake, instead of the probability of a rabbit's dying over the course of a month, let me consider the more elementary case of the one-half probability of a tossed coin's landing heads. The aim is to state the facts about the world that determine the value of the probability.

There are two kinds of facts that fix the outcome of a coin toss. First, there are the facts about the mechanism of the coin toss, that is, the physics of the toss and the construction of the coin itself. Second, there are facts about the initial conditions of the toss. In the previous section, I assumed implicitly that the relevant facts about the initial conditions are the precise initial conditions of a particular toss. I make enion probability a statistical notion by revising this assumption: the facts relevant to determining an enion probability, I now suppose, concern the distribution of initial conditions for some large—perhaps infinitely large—group of coin tosses. On this approach, the probability of an outcome is not an intrinsic property of the particular process that produces the outcome, but of a whole set of similar processes; for example, the probability of heads is not an intrinsic property of a particular coin toss, but is a property of a set of coin tosses. I intend, then, to define a *type* notion of probability, like (but not identical to) the frequentist notion (Reichenbach 1949), rather than a single case notion (Giere

1973).⁴ I will discuss the implications of this change in view for my project shortly; first, however, I want to go ahead with the definition without debating the consequences.

By the *mechanism* of a probabilistic setup such as a coin toss, I mean the physical process that converts a set of initial conditions into an outcome. For simplicity's sake, suppose that a coin toss has just one initial condition that varies from toss to toss: the speed with which the coin spins. Everything else, in particular, the orientation of the coin at the beginning of the toss and the time for which it is allowed to spin, is fixed. On this assumption, the initial speed of a toss entirely determines the outcome, heads or tails. The relevant facts about the dynamics of the coin can be represented, in a very abstract way, by a simple function that maps initial speeds onto outcomes. For any given initial speed, then, the function tells you whether a coin spun at that speed will yield an outcome of heads or tails. I call this function the setup's

4. The difference between a type notion and a single case notion is that a type notion of probability is attached in the first instance to a type of event or probabilistic setup, a single case notion to a singular event or token setup. To say that probabilities are attached to types does not mean that they cannot be applied to tokens—to predict or explain the occurrence of a singular event, for example—but some rule must be specified to determine which type probability to use for a given single case. This is what is known, in the context of frequentist accounts of probability, as the problem of the reference class.

Although scientists seeking to predict and explain the macrolevel are not interested in the fates of single enions, they are interested in the collective fates that determine macrolevel change. It seems that they will therefore have to solve the reference class problem, deciding which type probabilities are appropriate for predicting or explaining the statistical behavior of the group of enions making up the complex system of interest.

There are two well-known solutions to the problem. The first, endorsed by Reichenbach, is to use the narrowest applicable reference class or setup type for which a probability is known. The second, endorsed by Salmon, is to use the narrowest applicable homogeneous reference class (see Salmon (1984) for the meaning of *homogeneous*). My own treatment of, or rather end-run around, the reference class problem will emerge shortly.

The reference class problem is sometimes taken as a reason to reject type notions of probability. I think that type notions are workable in some circumstances, but nothing that I say here is intended to imply that the correct metaphysical analysis of probability is a type notion. The type notion advanced in this paper is intended solely as a tool for explaining the success of EPA; whether it makes for a good metaphysics of probability is entirely another question. See Strevens (2003), §1.3 for further discussion.

evolution function.

How does the evolution function figure in the definition of the probability of heads? I assume that there is some kind of probability distribution over the initial conditions of the coin toss, that is, over the initial speeds. I will not say anything, yet, about the nature of this distribution, except that it conveys information about a class of coin tosses rather than about any particular toss (for later thoughts, see section 5.2). The facts about the coin toss mechanism represented by the toss's evolution function determine that a certain set of initial speeds will lead to heads. The initial speed distribution determines the probability that the initial speed for any particular toss belongs to this heads-inducing set. That probability is the probability of heads. In other words, the probability of heads is the probability, according to the initial condition distribution, that the coin will be tossed with one of the initial speeds that leads to heads—just as you would expect.

Let me explain this definition of probability in pictures. First, I graph the evolution function for the coin (figure 1a). The function is equal to zero for those values of the spin speed v that lead to tails, and equal to one for those values of v that lead to heads. Thus, the gray parts of the graph correspond to heads-producing speeds, the white parts to tails-producing speeds.

Next I assume that the probability distribution over spin speeds is given by a density function of the sort shown in figure 1b. The density works in the usual way: the probability of obtaining a speed between any two values x and y is equal to the area under the graph between those two values, that is, the shaded area in figure 1b.

The probability of an outcome such as heads can be defined graphically as follows. Superimpose the relevant evolution function over the relevant initial condition density function, as shown in figure 1c. Then the probability of heads is the proportion of the initial condition density that intersects with gray areas in the evolution function, that is, the proportion of the area

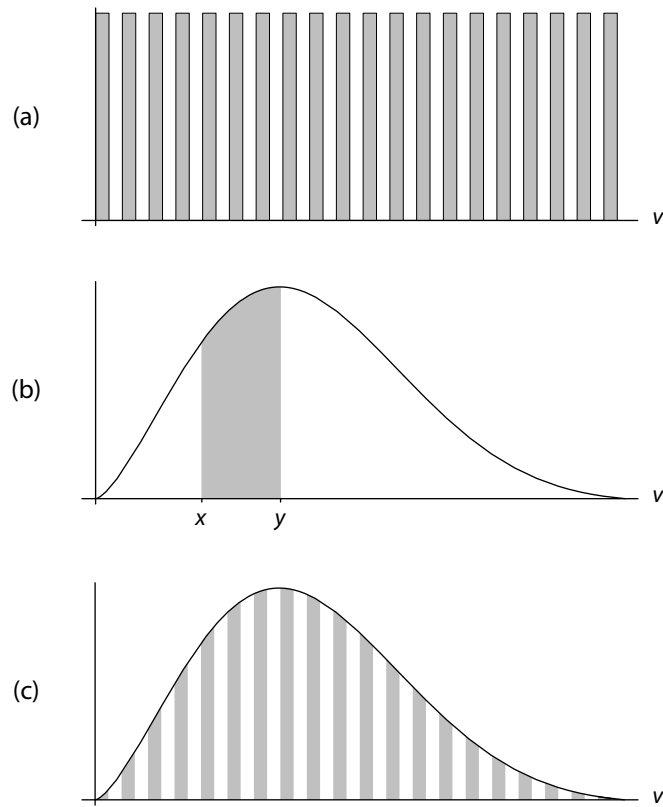


Figure 1: The tossed coin: (a) evolution function, showing values of initial spin speed v that lead to outcome heads; (b) probability density over initial spin speed; (c) the probability of heads

under the curve in figure 1c that is shaded gray.⁵ As you can see, the probability of heads is about one half, as expected.

I now return to the problem of showing that a complex system's enion probabilities are independent of microlevel information. I have solved a part of the problem. To put it in terms of the coin toss, I have defined the probability of heads on a given toss so that it does not depend on the initial con-

5. A more formal and detailed presentation of this definition of probability can be found in Strevens (2003), chap. 2.

ditions of that particular toss.

Is this a legitimate move or just a cheap trick? In the context of the present paper, this amounts to the question whether the property of complex systems captured by the definition is useful for the purposes of complex-system science, that is, whether, given its principled neglect of particular initial conditions, it can be used to draw sufficiently strong conclusions about the macrolevel behavior of systems.

Leaving particular initial conditions out of the definition handicaps enion probability as a predictor of the fate of individual enions: although enion probabilities can be used to forecast individual fates (see note 4), you could make much better predictions, at least in principle, by taking individual initial conditions into account. Disregarding particular details does not, however, affect significantly the power of enion probability to predict the fate of populations of enions, because the law of large numbers will allow population predictions, hence predictions of macrolevel behavior, to be made with near certainty.

5.2 *Microconstancy*

A statistical definition of enion probability takes care of the problem of enion probabilities' dependence on particular initial conditions, but introduces a parallel problem, that enion probabilities may depend on very particular *distributions* of initial conditions: it may be that the appropriate distribution of initial conditions to use for assessing an enion probability changes from enion to enion and from time to time, and that the value of the probability changes as the distribution changes. (The probabilities would still be type probabilities, based on sets of initial conditions and attached to types of probabilistic setup, but the types would be very finely specified.)

This would leave you not much better off than you were before: microlevel details about a complex system, reflected in the changing initial condition distributions, would make a difference to your assignments of enion

probabilities and so would have to be taken into account in deriving the system's macrolevel behavior from those enion probabilities.

There are two ways you might deal with this problem. First, you might formulate some standard of what is an appropriate initial condition distribution to use for a given enion probability, a standard that holds the distribution constant across different enions and times. You would then have to argue that a definition of enion probability that incorporates such a standard will be predictively and explanatorily useful.

Mine is the second kind of solution: I will explore the possibility that, even if the appropriate initial condition distribution changes from instance to instance, the probability induced by the distribution remains the same. It will follow that the probability, though perhaps not the distribution that determines the probability, is independent of microlevel details.

Clearly, it is not in general true that an enion probability will remain constant across shifts in the initial condition distribution. It is true, however, if the evolution function associated with the enion probability has a certain property, which I call *microconstancy*. I go on to claim (though this claim will not be defended here) that enion evolution functions tend to be microconstant. In the remainder of this section, then, I will characterize microconstancy and I will show how it creates probabilities that are robust in the face of changing distributions of initial conditions.⁶

A (single variable) evolution function for an outcome is microconstant if its range of possible initial conditions can be divided into small intervals, over each of which the ratio of conditions that will bring about the outcome to those that will not is about the same. Wherever you look in a microconstant evolution function, then, you will find roughly the same ratio. The evolution function shown in figure 1a is microconstant, you will observe,

6. The property I call microconstancy was originally used by Poincaré, Hopf, and others to explain the robust probabilities attached to gambling devices such as roulette wheels (Poincaré 1912; Hopf 1934). They called their approach the *method of arbitrary functions*.

because it can be divided into small, contiguous pieces each having a ratio of heads-producing values to tails-producing values—a ratio of gray to white—of one-to-one, or $1/2$. I call this constant ratio the evolution function's *strike ratio*. The strike ratio does not have to be $1/2$, of course; in a tossed die, for example, the strike ratio for the outcome of obtaining a six will be $1/6$.

Now observe that any reasonably smooth distribution of initial conditions will, in conjunction with a microconstant evolution function for an outcome, determine approximately the same probability for that outcome, equal to the outcome's strike ratio—where by *reasonably smooth*, I mean a density that is approximately constant over any small interval. In the case of the coin, because a smooth distribution over initial spin speed is constant over each of the small intervals having a $1/2$ strike ratio for heads, the total probability for heads is also approximately one half.⁷

This result is easily appreciated in pictures. You should be able to see that the shape of the initial condition density used in figure 1c does not matter very much at all. Any smooth density shaded in the same way would be about half gray and half white, resulting in a one half probability for heads. To induce a probability for heads other than one half, a density would have to look something like the density in figure 2; such a function, being sharply sloped over even quite small ranges of initial conditions, is not smooth in the relevant sense.

If your enion probabilities have microconstant mechanisms, then, you do not have to worry about protean initial condition distributions, as long as the distributions, however they change, remain smooth, since wherever and whenever you calculate a microconstant enion probability, you will arrive at the same result, a value equal to the relevant strike ratio. (This is why I am unconcerned as to which initial condition distribution to use for determining an enion probability. Provided that the plausible candidates are all

7. For formalization and proofs, see Strevens (2003), chapter 2.




Figure 2: A density that yields a probability for heads not equal to one half cannot be smooth

smooth, you do not need to know which is correct—if indeed there is a fact of the matter at all.)

Note that EPA does not require microconstancy over the entire initial condition space: since enion probabilities must be the same within any given macrostate, but may vary between macrostates, so EPA requires microconstancy within each macrostate, but allows that the strike ratio may be different for different macrostates.

Enion probabilities are, if what I have assumed so far is correct, equal to their corresponding strike ratios. Thus, provided that the values of the strike ratios are determined by fairly large-scale facts about the system—an ecosystem’s number of foxes, or a gas’s temperature—enion probabilities are not only stable over, but are entirely determined by, a system’s macrostate, making for an easy derivation of macrolevel behavior. The first part of the supercondition is vindicated.

You can now see the implications of the move, in section 5.1, to a statistical characterization of enion probabilities. Because I am concerned only with macrolevel behavior, what matters are not individual enions but enion statistics. Provided that I use the right statistics in my derivation of macrolevel behavior, my macrolevel predictions will be accurate. What statistics count as *right* in this sense? I must use statistics that represent accurately the initial conditions of the actual system whose behavior is to be predicted or

understood.

In the entirely general case, the move to statistical thinking would not constitute much of a gain, then, due to the difficulty of knowing the precise distribution of initial conditions for a particular system at a particular time (the more so because you need to know a joint probability distribution over all the different configurations of initial conditions that matter). If you know that you are dealing with microconstant probabilities, however, you are in luck. Provided that you are right in assuming that the actual initial conditions of the system in question are smoothly distributed, you can conduct your statistical analysis using any smooth distribution of initial conditions whatsoever, or equivalently, you can set all probabilities equal to the corresponding strike ratios. (Observe that the more enions there are in a system, the more likely it is that the initial conditions of the system are smoothly distributed, assuming that a smooth distribution is the norm. This is the law of large numbers at work creating the almost paradoxical state of affairs that the more complex a system—the greater the number of its enions—the more surely the complexity will fall out at the macrolevel.)

The philosophical move, then—the move to statistical thinking—is very little use in itself. It is a property of the dynamics of complex systems, microconstancy, that makes EPA, and thus the sciences of complex systems, possible.

The discussion in this section leaves open, I hardly need remark, two questions:

1. Why expect initial condition distributions to be smooth? And considerably more pressing,
2. Why think that the evolution functions of enion probabilities are microconstant?

To deal with these questions in a short paper is not possible. I refer you to

Strevens (2003), where they are answered in depth.⁸

6. INDEPENDENCE OF OUTCOMES

Why it is reasonable to regard the outcomes to which enion probabilities are attached, such as rabbit death, as stochastically independent? Much of the foundation for the answer to this question was laid in the previous section.

The great impediment to the assumption of stochastic independence is the evident causal entanglement of the processes determining the fates of different enions, for example, the processes determining the death and survival of different rabbits. To justify the supercondition's assumption of stochastic independence, I must establish that, in certain circumstances, causal interaction does not destroy stochastic independence.

In general, causal interaction does destroy stochastic independence—but not, it turns out, when the probabilities involved have microconstant evolution functions. Microconstancy, I will show, works to preserve stochastic independence in the face of causal influence.

Suppose that two coins are tossed at the same time (but with uncorrelated initial spin speeds). They collide in mid-air. The processes producing the outcomes of the two tosses are not, then, causally independent. Normally, however, the outcomes are stochastically independent: the probability of heads on one toss, conditional on the outcome of the other, is equal to the unconditional probability of one half. (Intuitively, knowing how one coin landed gives you no information about the outcome of the other toss.)

My goal is to understand why the interaction between the two tosses does not destroy their independence, and to show (as far as is possible in a short paper) how the result generalizes to all microconstant processes. To this end, I will introduce the notion of a composite evolution function. The

8. Question (1) is answered in section 2.5 of the book, question (2) in chapter four.

evolution function of the previous section expressed the dependence of the outcome of a single coin toss on a single initial spin speed. The composite function does the same for the two colliding coins: it yields the outcomes of both tosses, given their initial spin speeds.

Let me begin with the evolution function for the simplest case of a double toss, that in which there is no collision. I will then examine the effects various kinds of collision on the dynamics.

To keep things simple, consider just two possible outcomes from a toss of two coins: either both coins land the same side up, that is, both heads or both tails, or they do not. If they land the same side up, call the outcome a *double*. The evolution function for a double can be drawn as in figure 3. The




Figure 3: Composite evolution function for two non-colliding tossed coins with initial spin speeds v_1 and v_2

diagram shows the outcome of any toss as a function of the initial speeds v_1 and v_2 of the two coins. For pairs of initial speeds shaded gray, both tosses yield the same outcome, so you have a double; otherwise they yield different outcomes, so no double. You should satisfy yourself that figure 3 is the correct evolution function for the case where the coins do not interact.

The next thing to establish is that, in this simple no-collision case, the

outcomes of the tosses are stochastically independent. In principle, I should go through and check the probabilities of every possible outcome of two tossed coins. In fact, I will focus on the probability of a double; filling out the story is a simple exercise.

For the tosses to be independent, the probability of a double needs to be one half. This is because stochastic independence implies that, if one coin lands one way—say, heads—then the probability that the other coin also lands that way, yielding a double, is just the usual, unconditional probability, namely, one half. You will see from the microconstancy and one half strike ratio of the evolution function in figure 3 that the probability is indeed one half, provided that the initial condition distribution over spin speeds is reasonably smooth.

This result is quite general. Whenever you look at the outcomes of several causally independent probabilistic processes with microconstant evolution functions, you will find a composite evolution function that is microconstant with the strike ratio needed for stochastic independence. This is just as well, since it is a commonplace that where there is causal independence, there is stochastic independence (Strevens 2003, §§3.3, 3.4).

I will continue to assume that the probability distribution over initial spin speeds is smooth, while introducing some degree of interaction between the two coins. First, suppose that there is a linear interaction, meaning that the effect of the collision between the two coins is to alter the spin speeds so that each is a linear function of their pre-collision values.⁹ The composite evolution function for two linearly interacting tosses will look something like figure 4a. You can see that the interaction has a real effect on the evolution function. But it does not alter the function's microconstancy, nor the strike ratio for a double, which is still one half. Thus it does not destroy

9. Writing the spin speed of the first coin after the collision v'_1 , then, it must be the case that $v'_1 = av_1 + bv_2$ for some a and b , where v_1 and v_2 are the coins' spin speeds before the collision.




Figure 4: Composite evolution function with (a) linear interaction and (b) nonlinear interaction

stochastic independence. This is in general true of linear interactions: they preserve microconstancy and strike ratios in composite evolution functions, thus they do not affect independence.

In geometrical terms, a linear interaction has the effect of changing the point from which the “chessboard” structure of the evolution function is apparently being viewed, but nothing else. Different linear interactions between the tossed coins will, therefore, produce evolution functions that look like chessboards viewed from different angles and distances. The viewing angle will not change the microconstant aspect or the strike ratio of the function; this is why stochastic independence is preserved.¹⁰

10. You may notice that there are two ways of looking at a chessboard that destroy microconstancy. First, you might look at it side on. Second, you might look at it very close up, so that the squares are so large that the evolution function is no longer *micro*-constant: you have the pattern but it is too big. These cases are exceptions to the rule. It turns out that the exceptions will not occur if the dynamics of the collision is *inflationary*, meaning roughly that small differences in pre-collision variables make for larger differences in post-collision variables. Consequently, the exceptions will not occur when interactions between enions are sensitive to initial conditions, or chaotic, as in a complex system.

Most interactions are not linear. What, then, is the effect of a nonlinear interaction on the coin tosses? A typical nonlinear interaction produces the composite evolution function shown in figure 4b. The microconstancy of the composite evolution function is now visibly strained. But it holds up: the new evolution function is approximately microconstant with a strike ratio of one half. The nonlinear interaction has left stochastic independence intact.

Why? The nonlinear interaction I used to construct figure 4b has an effect that is very close to linear over small areas. Because microconstancy is a property that holds, in effect, in virtue of the layout of small areas, and linearity preserves the properties required for microconstancy, microconstancy survives a nonlinear interaction that is linear in the small. Provided that nonlinear interactions are linear over very small scales, then, stochastic independence will be preserved, or rather, approximately preserved (in the sense that all probabilities will be very close to the probabilities obtained when independence is strictly preserved). This property of a function's being approximately linear over very small intervals or regions of its domain is typical of functions that represent physically real processes, including the nonlinear processes studied by chaos theory (and is, in its limiting version, the property that underlies the applicability of calculus to physical processes).

It follows that the fates of enions in a complex system are stochastically independent provided that:¹¹

1. Interactions between enions are either linear, or if nonlinear, are approximately linear over small regions of initial conditions, and
2. The evolution functions of enion probabilities are microconstant.

11. The issue of independence in complex systems is more involved than I have made out here, because to show that enion evolution functions are microconstant, I need further independence results not presented in this paper. For the details, see Strevens (2003), chap. 4. Once these additional complications are taken into account, however, it turns out that the claim about to be made in the main text is true all the same, if you add the requirement that enion behavior is chaotic, that is, sensitive to initial conditions in a certain way. For an insight into the role of chaos in promoting independence, see note 10.

Condition (1) is, as I have just remarked, typically true of real world interactions, while condition (2) is, as earlier noted, shown to be true for enion dynamics in Strevens (2003). The explanation above of the stochastic independence of two colliding coins looks to generalize, then, to interacting enions in most or all complex systems.

But, I hardly need add, there is much work left to do. Even in the case of molecular collisions, the complicated quantum mechanics of intermolecular force makes the demonstration of (1) difficult (though it is straightforward for idealized models of collision, as shown in Strevens (2003), §4.8), and that (1) applies to predator/prey interactions is to a great extent simply a matter of faith. I take it, however, that (1) is plausible enough in these and other cases to offer a foundation for complex-system sciences that is of real interest.

A final observation: to say that one rabbit's death is stochastically independent of another is to say that taking into account one death does not affect the *probability* of the other, in the population-level sense of probability defined in section 5.1. Stochastic independence is quite compatible with, and indeed in your typical ecosystem coexists with, both the causal and the epistemic relevance of one death to another (though to take advantage of the epistemic relevance, your knowledge of the system would have to be vast). What the independence results show is that neither form of relevance makes a difference to macrolevel behavior.

7. CONCLUSION

The probabilistic supercondition is true, or at least true enough, if probability is defined in the right way, thanks to certain special properties of complex systems. The most important of these properties is the microconstancy of enion probabilities. Microconstancy is, as it were, what neutralizes microlevel chaos and allows the emergence of a macrolevel dynamics that is

independent of low level facts. In so doing, microconstancy ensures that macrolevel behavior is derivable, in a tractable way, from microlevel behavior, so provides a foundation for EPA, and thus makes the sciences of complex systems possible.

If you look closely, however, you will see that microconstancy is itself a form of chaos. When a mechanism has a microconstant evolution function, small changes in initial conditions will cause a shift from white to gray or vice-versa, and so a change in outcome.

This observation goes just as deep as you might hope. In order to show that enion probabilities in complex systems have microconstant evolution functions—something that I have not done here—I take advantage of the microlevel chaos inherent in every complex system (see notes 10 and 11). The same chaos that appears to threaten the tractability of complex-system science is revealed as its enabler and guarantor, its friend.

REFERENCES

- Auyang, Sunny Y. (1998), *Foundations of Complex-System Theories in Economics, Evolutionary Biology, and Statistical Physics*. Cambridge: Cambridge University Press.
- Bechtel, William and Robert C. Richardson (1993), *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton, NJ: Princeton University Press.
- Garber, Elizabeth, Stephen G Brush, and C W F Everitt (eds.) (1995), *Maxwell on Heat and Statistical Mechanics*. Bethlehem, PA: Lehigh University Press.
- Giere, Ronald (1973), "Objective Single Case Probabilities and the Foundation of Statistics", in P Suppes, L. Henkin, G. C. Moisil, and A. Joja, (eds.), *Logic, Methodology and Philosophy of Science IV: Proceedings of the Fourth International Congress for Logic, Methodology and Philosophy of Science, Bucharest, 1971*. Amsterdam: North Holland.
- Hopf, Eberhard (1934), "On Causality, Statistics and Probability", *Journal of Mathematics and Physics* 13:51–102.
- Loeb, Leonard B (1934), *The Kinetic Theory of Gases*. 2nd edn., New York: McGraw Hill.
- Ornstein, Donald S and Benjamin Weiss (1991), "Statistical Properties of Chaotic Systems", *Bulletin of the American Mathematical Society* 24:11–116.
- Poincaré, Henri (1912), *Calcul des Probabilités*. 2nd edn., Paris: Gauthier-Villars.
- Reichenbach, Hans (1949), *The Theory of Probability*. Berkeley, CA: University of California Press.

- Roughgarden, Jonathan (1979), *Theory of Population Genetics and Evolutionary Ecology*. New York: Macmillan.
- Salmon, Wesley (1984), *Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.
- Simon, Herbert A (1969), *The Sciences of the Artificial*. 1st edn., Cambridge, MA: MIT Press.
- Sklar, Lawrence (1993), *Physics and Chance*. Cambridge: Cambridge University Press.
- Sober, Elliott (1984), *The Nature of Selection*. Cambridge, MA: MIT Press.
- Strevens, Michael (2003), *Bigger than Chaos: Understanding Complexity through Probability*. Cambridge, MA: Harvard University Press.
- Wimsatt, William C (1986), "Forms of Aggregativity", in A Donagan, A N Perovich, and M V Wedin, (eds.), *Human Nature and Natural Knowledge*. Dordrecht: D. Reidel.