How Are the Sciences of Complex Systems Possible?

Michael Strevens

Draft of September 2004

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 heart. 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 supposed that complexity can be managed by a two-step process. The process 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 have sufficient simplicity to survive the daunting combinatorics of complexity introduced in the second step.

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.[1] But many complex systems are, it is generally conceded, not decomposable at all, and are therefore unassailable

1. Decomposability is discussed in greater detail in the next section. For more, see Simon (1969), Wimsatt (1986), Bechtel and Richardson (1993), and Auyang (1998), chaps. 4–6. The overview in Bechtel and Richardson (1993), chap. 2 is especially helpful.

by decomposition's divide-and-conquer approach. The aim of this paper is to investigate a way of negotiating such systems' complexity.

2. Non-Decomposable Complex Systems

My inquiry will focus on a particular class of non-decomposable complex systems. The systems I have in mind are made up of many fairly independent, but strongly interacting, parts. By *fairly independent* I mean that, most of the time, any given part does its own thing without being affected significantly by the other parts. By *strongly interacting* I mean that those occasions when the parts do affect one another have potentially radical consequences for the interactors. Typically, due to a complex system's large number of parts, many strong interactions occur over even relatively short periods of time.

Two important examples of complex systems of this sort 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.[2] 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, regardless 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.

From this point on, I refer to the class of systems just described simply as *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

2. Gas molecules do exert a slight force on one another at all times, but it is not comparable to the radical effect of a collision.

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, pretty much any system counts as decomposable, in that 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, rather than changes in the 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, average energy without noticing the energies of individual molecules.

Macrolevel behaviors of complex systems, then, are behaviors such as the following:

1. The way that populations of different kinds of organisms change in an ecosystem;

2. Various social dynamics: suicide rates, the business cycle, the statistical relation between socioeconomic status and success in school;

3. The behavior of the systems studied by statistical physics: diffusion behavior, the ideal gas law, the second law of thermodynamics.

My topic is the scientific endeavor of predicting and explaining these 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 what is known of *microlevel behavior*, by which I mean the individual behavior of a system's components. 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, although what I do say about the foundations of a certain kind of science has implications for these metaphysical questions.

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 provide explanations and predictions of the macrolevel behavior of gases, ecosystems, societies, and other complex systems of this sort. In section 1, I observed that the prima facie greatest obstacle to understanding the macrolevel dynamics of such complex systems is the sheer number of different interactions determining the systems' behavior. The problem posed by the interactions is as follows.

In order to explain or predict the macrolevel dynamics of a complex system using the behaviors of the system's individuals, it is necessary to have some grip on the behaviors of the individuals. To characterize the behavior of an individual, the influences on the 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, it seems that the behaviors of all the other individuals must be characterized already, at least in the aggregate. Such characterization, it seems, cannot proceed one individual at a time. 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.

In decomposable and near-decomposable systems, the leap is not such a challenge after all. This is because the influences between individuals in decomposable systems are 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 the aggregate external influence in the form of a static background condition, such as a "mean field" in physics.[3] 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 (Simon 1969; Bechtel and Richardson 1993).

Observe that, above all else, what makes the decomposition strategy viable 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 *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

---

3. The justification for the 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. As I will show in section 3.4, the same technique of invoking an aggregate, static influence works in some non-decomposable systems for more elaborate, and therefore more interesting, reasons.

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—perhaps including death—will depend almost entirely 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 existence. 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 not to work, because its central assumption about complex systems, indispensable for the derivation of macrolevel from microlevel dynamics, appears to be false. Yet the method does work. To understand why is the main business of this paper.

3.1 The Form of EPA

The method I have in mind has no official 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 so as 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.[4]

I will give a very simple example of EPA. Imagine that you wish to track the population of rabbits in an ecosystem. Suppose that you can assign to any given rabbit 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 what is called *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 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

---

4. For much 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 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.

Some version or other of the technique I have just illustrated—that is, some version of EPA—has been used to understand complexity in statistical physics ever since Maxwell published his seminal work on kinetic theory in the 1860s. Enion probability analysis is also the basis of population genetics, the mathematical engine of modern evolutionary theory, and it is taken for granted in much work on economics and society. Its utility as a route to the macrolevel in complex-system science is, then, enormous.

There is rather more to EPA, both good and bad, than might first appear. Let me begin with the good, in the form of three pleasant surprises.

3.2 The Simplicity of Macrolevel Behavior

The first surprise is that the laws derived using EPA are quite simple. All of a complex system's complexity seems to vanish in the course of the probabilistic derivation. This is not merely a fact about EPA, but a fact about complex systems themselves: at the macrolevel, they tend to behave in simple ways. Some examples:

1. In *ecosystems*, populations are often very stable. Not only do they tend to stay the same; when perturbed by some kind of unusual event, they tend to return to their original levels quickly (see, e.g., Putman and Wratten 1984). This stability is one kind (not the only kind) of simple behavior.

2. In *social systems* there is what nineteenth century thinkers regarded as a truly remarkable constancy in the rates of such events as suicides, undeliverable letters, marriages, criminal acts and so on. For a time,

these thinkers envisaged an entire science of iron social laws (Porter 1986; Hacking 1990).

3. In *statistical physics*, you have the dark father of all complex-system laws: the second law of thermodynamics, which says that whatever happens in a system, its entropy must increase.[5] There are many other simple generalizations to be found in statistical physics, such as the ideal gas law, the laws of diffusion, and so on. In each case, systems that are otherwise very diverse in their make-up and dynamics have a common behavior that can be captured by an extremely simple mathematical equation.

It is important to see that the claim that a system's behavior can be characterized by a simple 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. That much is merely the gathering of statistics.

To say that the behavior is characterized by a macrolevel law is, by contrast, to say that 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.[6] This is quite unexpected. Because of the frequent, strong interactions between a complex systems' parts, small differences—differences only apparent in a microlevel description—can make large differences to the fate of any given enion. You might well have thought,

5. Or more accurately, entropy must not decrease—it may remain constant, as it does when already at its maximum value. The second law holds only if the system is energetically closed, and only probabilistically.

6. In the nicest cases, the relation is deterministic, or virtually so. In other cases, the relation is probabilistic, and so can be discerned only by looking at the behavior of a number of systems of the same kind.

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 microlevel realization makes no apparent difference to the way the system behaves.

The relative simplicity of macrolevel dynamics is, it is well worth noting, what makes possible a kind of complex-system science different from that examined in this paper, namely, the empirical modeling of a complex system's macrolevel behavior without regard for its microlevel foundations. It is only because macrolevel behavior is relatively simple—requiring, for its characterization, a few variables rather than variables describing the complete microlevel state of a complex system—that such modeling has a chance of success.

3.3 The Explanation of Macrolevel Simplicity

The second surprise is that EPA seems to provide a simple and accessible, yet extremely general, explanation of *why* complex systems behave in simple ways: EPA's probabilistic treatment allows a perspicuous representation of the way in which complex tangles of strong microlevel interactions cancel out at the macrolevel, leaving simplicity. But to talk of the interactions "cancelling out" is to put things a little 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. Any probabilistic process can be resolved conceptually 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 the first sub-process, and therefore caus-

ing 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 an overlying set of probabilistic fluctuations that cause 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,[7] but due to various perturbations, almost never hits the target exactly.

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 of as producers of individual outcomes, only as producers of a final macrolevel statistic (compare mathematical models in the social sciences that are deterministic save for a probabilistic "error term").[8] Second, the breakdown into two processes is introduced purely for expository

7. Insofar as this is possible given the total number of tosses—obviously, if there are an odd number of outcomes, to achieve precisely one half heads is out of the question.

8. Technically, the process to be decomposed will be a series of random variables, the i^{th} of which represents the statistics for the first i trials (as opposed to the outcome of the i^{th} trial). For example, a Bernoulli process, say a series of coin tosses, will be conceived of as a series of random variables, the i^{th} of which represents the frequency of heads in the first i trials. This process can be resolved mathematically into two component processes, for which the values of the i^{th} terms of each will sum to give the frequency of heads after i trials. The values of the terms of the deterministic process will be the mean of the relevant Bernoulli distribution, hence, in the case of the coin tosses, always equal to one half, and more gener-

purposes, so as to give the reader 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 the values of the relevant probabilities, and second, a set of fluctuations perturbing the outcome of the deterministic process. This *determinism plus fluctuations* way of thinking is especially common in evolutionary biology, where natural selection is conceived of as deterministic, and genetic drift as fluctuation (Sober 1984), but it is possible wherever EPA can be applied.[9]

Now, it is a fact that, the more outcomes are produced by the probabilistic process, the smaller the 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

---

ally, always equal to the single Bernoulli parameter p. The terms of the probabilistic process will represent the deviation from the mean after i trials. The probability distribution over these random variables will be identical to the distribution over the original process, except shifted so as to have zero mean. (If $P_i(x)$ is the distribution over the original process, that is, the probability distribution over the frequency x of heads after i trials, then the distribution over the subprocess is $P_i(x + p)$. It is, in other words, the distribution obtained by transforming the random variable for frequency by subtracting p.) This is, needless to say, a mathematically trivial construction; its value is solely expository. For the mathematical theory of random processes, see Grimmett and Stirzaker (1992).

9. The determinism plus fluctuations characterization of the behavior of complex systems was especially influential in the nineteenth century. The two sub-processes were typically accorded physical reality. See Strevens (2003), §1.14.

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 relevant probabilities; in the case of the complex systems, they are the enion probabilities, such as the probability of a rabbit's surviving for a month. But the enion probabilities depend only, EPA assumes, 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. 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.

3.4 The Compositionality of EPA

The third surprise—which will, to keep things brief, play no role in the later discussion—is that EPA treats complex systems almost as though they are decomposable. It characterizes the behavior of individual enions first, in the form of a probability distribution. The distributions for all of a system's enions are then aggregated to obtain a formula for the behavior of the system as a whole. This is, on the face of it, just the two-step process characteristic of decomposition (see section 1).

But not quite: the first step of the decomposition approach characterizes the behavior of individual enions *in isolation* or *against a static background*. It does not take into account the dynamic effects of enion interaction; these are introduced only in the second step.

In EPA, by contrast, the probability distributions assigned in the first step already reflect enion interaction. The probability assigned to the event of a rabbit's dying in the course of a month depends in part, for example, on the chance that the rabbit encounters a fox, hence on whatever determines the rate of fox/rabbit encounters—the number of foxes in the ecosystem, the ways in which rabbits avoid foxes, the ways in which foxes find rabbits, the layout of the terrain, and so on. The argument that complex systems are not decomposable (section 2) stands: in dealing with complex systems, interactions between enions must be taken into account right from the very beginning. Enion probability analysis respects this conclusion by building interactions into the enion probability distributions themselves.

The scope and power of EPA nevertheless derives in great part from its ability to mimic the decomposer's two-step. Interactions are taken into account in the first step, but they are only *explicitly* taken into account in the second step, when the probabilities are aggregated. This is possible because the probabilities assigned to individual enion behaviors in the first step can be treated almost as though they are intrinsic properties of individual enions. Once a few parameters are allowed for, enion probabilities travel with the enions wherever they go, so that they can be used to predict and explain the behavior of arbitrary groups of enions.

For example, knowing just the birth rate and the death rate for rabbits, I can predict the macrolevel behavior of any group of rabbits. This is true even if the birth and death rates depend on particular macrolevel facts about particular groups, such as the local fox population: I simply take the parameters that encode these particular facts, drop them into the formulas for the enion probabilities, and aggregate.

Theories of complex systems constructed according to the precepts of EPA are *compositional*, then, in the same way as theories produced by the decomposition strategy, despite the fact that the component parts of an EPA-driven theory—the enion probabilities—are not properties of the component parts of the system, but are properties of the system as a whole. How can this be?

In section 2, I argued that complex systems are not amenable to decomposition because the strong interactions between enions make facts about their dynamics in isolation quite inadequate for understanding their dynamics when embedded in a system. In a decomposable system, by contrast, interactions between parts do not matter in the short term. I was quite right to say that strong interactions make a difference to the fates of individual enions in a complex system, but as you can see from section 3.3, 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. The product of this averaging is a static background that makes a kind of quasi-decomposition possible.

For example, once you take the statistical view of an ecosystem, where you do not attend to the fate of a particular rabbit but only to the fate of rabbits in general, you can consider the behavior of each individual rabbit in the context where it experiences the average number of fox encounters, the average number of rabbit encounters, and so on. This constant background provides the basis for the probability distributions that you impose over rabbit behavior. It is because of the resulting universality—the same probability of death is assigned to each rabbit, for example, once a few important parameters such as fox population are set—that EPA is able to supply a compositional toolbox with which you can construct a useful mathematical model for the macrolevel behavior of any given population of rabbits, provided that the population is sufficiently large.[10]

---

10. This passage is intended to give the reader only a rough sense of the basis of EPA's compositionality. Because I will not, for reasons of space, be able to return explicitly to the topic of compositionality, let me make one forward-looking comment: any enion probabilities that satisfy what I call in section 4 the *probabilistic supercondition* will be compositional in the desired sense.

3.5 Summary

The prima facie obstacle to understanding the macrolevel behavior of complex systems is the sensitivity of the fate of each enion to the many strong interactions between enions, for example, the fact that whether or not a given rabbit dies depends on the positions and dispositions of all the other rabbits and foxes— indeed, of anything that might influence a rabbit's movements, since those movements can make the difference between good luck and bad luck, so between life and death. The obstacle is, in other words, the sensitivity of microlevel outcomes to microlevel details.

Enion probability analysis equates the aspect of the strong interactions that makes prediction and explanation difficult with that part of a probabilistic process that can be conceived of as a series of fluctuations. The fluctuations cancel out and so disappear from the picture at the macrolevel, provided that a system has sufficiently many enions. The cumulative effect of many strong interactions is, then, as *insensitive* to microlevel details as the individual interactions are sensitive. In this way, microlevel complexity coexists with, indeed helps give rise to, macrolevel simplicity. And its doing so can be understood by way of a compositional theory.

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 seems, 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 3.1, 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*.

It is only in virtue of the two parts of the supercondition that EPA's derivation of macrolevel behavior is able to avoid having to take into account the complexities of the microlevel. The reasons for this will become clear later in this section; first, however, let me give you some sense of the supercondition's role in brief by examining its central place in EPA's explanation, presented in section 3.3, of the macrolevel disappearance of complexity.

Enion probability analysis, recall, 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 set of fluctuations perturbing those statistics. Because a complex system has many enions and so has many instances of the process, 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. Here assumption (1) is invoked to conclude that only macrolevel facts affect the deterministic sub-process, 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.

A few paragraphs, now, spelling out assumptions (1) and (2) in more detail. First, assumption (1). To assume that the probability of rabbit death depends only on macrolevel facts is to assume that it does *not* depend on any 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 very chaotic, in the sense that it is very sensitive to initial conditions,[11] then perhaps these details are very, very small. Enion probability analysis assumes that, despite all of 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.

Second, assumption (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 3.1'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.

---

11. Sensitivity to initial conditions is the core of the notion of chaos; various more technical notions of chaos can be obtained by characterizing the required sensitivity in various ways—as producing exponential divergence of nearby dynamic trajectories, as being produced by a stretch-and-fold process, and so on.

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 sorts of low level information. The first part says that the probability of death is unaffected by details about the state of the system at the beginning of the month; the second part says that the probability of death is unaffected by 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. Add together 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.

Probabilistic reasoning in the style of EPA enables you to ignore the microlevel, then, because your assumption of the probabilistic supercondition builds independence of the microlevel into the probabilities. The probabilistic supercondition 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, this assumption seems false: the 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.

Let me make this point more carefully, once again dividing the supercondition into its two parts. First, consider the assumption that the probabilities for rabbit death depend only on a system's macrolevel state. There is one kind of worry that can be dismissed fairly quickly. You might think that 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.[12]

What the probabilities for death absolutely must not depend on are details such as the initial positions of the rabbits and foxes. But I have already pointed out that the fate of a rabbit does depend on these details. Take two rabbits, alike in every relevant respect: health, strength, age, and the rest. Over the course of a month, one is eaten and one survives. What makes the difference? It can only be those little details of position or whatever. By what right does EPA discount the details? They can make just as much of a difference between survival and death as do health, strength, and age. Enion probability analysis should not ignore them, yet it does, and somehow it all works out.

The supercondition's second part, according to which deaths are stochastically independent, is equally problematic. There is a heuristic that is universally applied to justify assumptions of stochastic independence: two outcomes are stochastically independent if they are causally independent, that is, if the processes that bring about the two outcomes are not causally entangled. Outcomes such as rabbit deaths in a complex system manifestly fail this test: 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.

---

12. In other words, a certain small degree of microlevel dependence in the enion probabilities is tolerable. The reasons for this—not entirely straightforward—are explained in Strevens (2003), §4.6.

Because the causal test gives only a sufficient condition for stochastic independence, failing the test does not imply the absence of independence. But it does leave the independence assumption in complex systems without any known basis.

Let me summarize the situation:

1. Complex systems have chaotic microlevel dynamics, in the sense that small details about the initial conditions of a complex system can radically affect the fate of individual enions.

2. This microlevel chaos prevents the use of the decomposition strategy to predict and explain macrolevel behavior using microlevel behavior. It provides reason, perhaps, to think that no prediction or explanation of any sort is possible, at least insofar as it would require a tractable derivation of macrolevel dynamics from microlevel dynamics.

3. Enion probability analysis purports to show how macrolevel behavior may be derived from microlevel behavior, thanks to microlevel complexity's cancelling out.

4. On closer examination, however, it turns out that EPA simply *assumes* that microlevel complexity falls out of the picture. Nothing is explained after all; chaos menaces the prospect of complex-system science as much as ever.

The founders of EPA were not unperplexed by all of this. Maxwell, who pioneered the application of EPA to gases, wrote that the technique worked well provided that one avoided making any "personal enquiries" of molecules (Garber et al. 1995, 19, 422). Can a policy of social discretion save the day?

The thought is both tempting and bizarre; tempting because avoiding personal enquiries seems to have just the desired effect, and bizarre because it supposes that the irrelevance of microlevel details to macrolevel behavior is due to your ignorance of microlevel detail, as though what you know nothing about, thereby has no power to affect the course of events.

My goal in what follows is to show that the assumption of the probabilistic supercondition is justified, not by any 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.

In attempting to justify the supercondition, I am embarking on a project similar to that of Ornstein and Weiss (1991) and other proponents of modern ergodic theory, who have sought 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. Some 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 further comments 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 13 below)—the supercondition is true, or close enough to true. It is not the definition itself that makes it true, however, so much as the chaotic physical properties of complex systems.

This section discusses the first part of the supercondition, the assumption that enion probabilities are independent of the microlevel details of a system's initial conditions. Let me begin by reminding you why enion probabilities need to be independent of microlevel information, if EPA is to provide a tractable method for deriving macrolevel behavior from, and therefore predicting and explaining macrolevel behavior using, facts about microlevel behavior.

Anything that a system's enion probabilities depend on must be taken into account in EPA's derivation of the system's macrolevel laws. The more microlevel dependencies there are, the more complex the derivation. If, in particular, the enion probabilities depend on such details as the positions of individual rabbits, then all such positions must contribute to the derivation, resulting in the kind of complexity, and perhaps intractability, that EPA is supposed to avoid.

The demonstration that enion probabilities do not depend on rabbit positions and so on has two parts: a philosophical trick, and a physical investigation of enion probabilities. First, the trick.

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. I remarked above that this seems hopeless, because whether or not the rabbit dies is *determined* in part by its starting position.

I circumvent this initial problem by making enion probability itself a kind of statistical property, employing a definition of enion probability that I will now describe. It is much easier to explain the definition using a simple probabilistic setup than a complex system such as an ecosystem. So, rather than considering the probability of a rabbit's dying over the course of a month, let me consider in its place 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 determine 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. So far, I have been implicitly assuming that the relevant facts about the initial conditions simply are the 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 typical initial condition distribution for some large (perhaps infinitely large) group of coin tosses. On this new assumption, the probability of an outcome is no longer 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 the same as) the frequentist notion (Reichenbach 1949), rather than a single case notion (Giere 1973).[13] I will discuss the implications of this change

---

13. 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 theory must be given of which type probability to use for a given single case. This is what is known (at least 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 other, 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.

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 physics of the coin can be represented by a simple function that maps initial speeds onto outcomes. For any given initial speed, then, the function tells you whether a coin spun with that speed will yield an outcome of heads or tails. I call this function the setup's *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 note 19). The evolution function for the toss determines 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

---

The reference class problem is sometimes taken as a reason to reject type notions of probability. I, of course, 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.)

be tossed with one of the initial speeds that leads to heads. I take it that this will seem eminently sensible.

Let me explain how this definition of probability works in pictures. First, I graph the evolution function for the coin (figure 1). The evolution func-

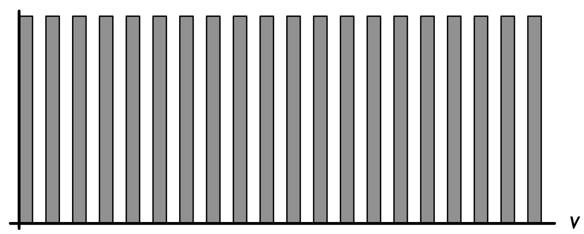

Figure 1: Evolution function for the tossed coin. The horizontal axis gives the coin's initial spin speed v; the vertical axis gives the value of the function, indicating whether or not a particular speed leads to an outcome of heads.

tion 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 2. The density works in

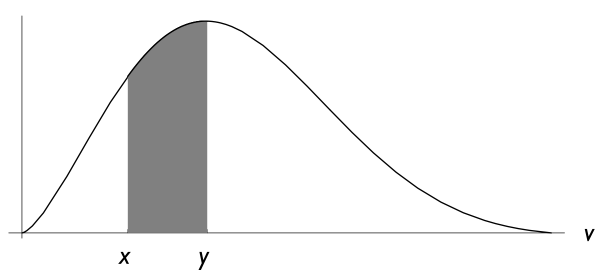

Figure 2: Probability density over the initial speed v of the tossed coin

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 2.

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 3. Then the probability

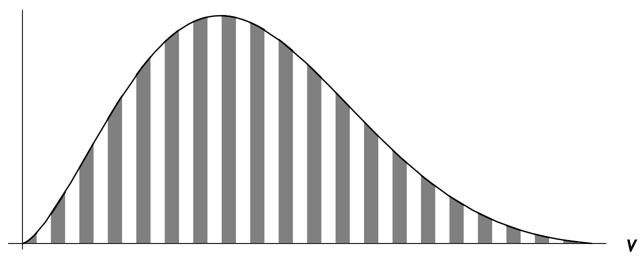

Figure 3: The probability of heads. The vertical axis gives the probability density over initial spin speed v.

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 under the curve in figure 3 that is shaded gray.[14] 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 conditions of that particular toss.

Is this a legitimate move or just a cheap trick? Of course, I am free to define my terms any way I like; all definitions are legitimate in this narrow sense. The real question is whether the property of complex systems captured by the definition is useful for the purposes of complex-system science,

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

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 13), 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 allows population predictions, hence predictions of macrolevel behavior, to be made with near certainty.

I have yet to establish, of course, that the independence assumption needed for the large numbers law holds; that is the topic of section 6. Before independence, however, I must solve a different problem.

5.2 Microconstancy

A statistical definition of enion probability takes care of the problem of enion probabilities' dependence on particular initial conditions, but introduces the possibility 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.)[15]

15. The usual rules for determining the appropriate type probability for a predictive or explanatory task, identified in note 13, do tend towards narrowness of specification.

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, in such a way that the standard 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 determined by the distribution remains the same. Then 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. I will show that it is true, however, if the evolution function associated with the enion probability has a certain special 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.[16]

An evolution function for an outcome is microconstant if, over any small set of initial conditions, the ratio of conditions that will bring about the

---

16. The property I call microconstancy was originally used by various luminaries in mechanics—in particular, Poincaré and Hopf—to explain the robust probabilities attached to gambling devices such as roulette wheels (Poincaré 1912; Hopf 1934). In their hands, the approach was called the *method of arbitrary functions*.

outcome to those that will not is about the same. Wherever you look in the evolution function, that is, you find the same ratio. The evolution function for the tossed coin is microconstant, because over any small range of spin speeds, the ratio of heads-producing values to tails-producing values—the ratio of gray to white in figure 1—is 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.[17] In the case of the coin, because over any small area about one half of the initial conditions lead to heads—in graphical terms, one half of the area is shaded gray—any smooth distribution over initial spin speed will determine a probability for heads of about one half.

This result can be proved formally, but it is easy to appreciate in pictures.[18] Figure 4 shows three different initial condition distributions superimposed on the coin's evolution function. It is clear that each of these, and any other smooth distribution, will impose a probability of one half for heads, since about one half of the area under any smooth distribution will be shaded gray by the evolution function.

If enion probabilities are microconstant, you do not have to worry about protean initial condition distributions. On the assumption that the relevant distribution is always smooth, wherever and whenever you calculate an enion probability, you will arrive at the same result, a probability equal

17. By *reasonably smooth*, I mean that the density function should be approximately constant over any small interval.

18. For a more formal approach, see Strevens (2003), chap. 2, especially sections 2.2 and 2.C.

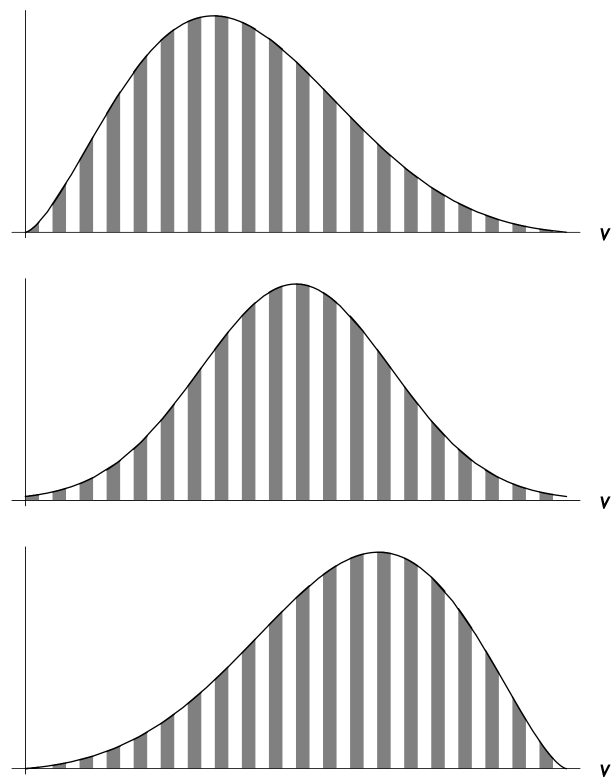

Figure 4: The shape of the initial condition density, if smooth, makes no difference to the value of a microconstant probability. The vertical axis gives the probability density over the initial condition *v*.

to the relevant strike ratio.[19] (Microconstancy over the entire initial condition space is more than EPA requires: since enion probabilities must be the same within any given macrostate, but may vary between macrostates, EPA requires microconstancy within each macrostate, but allows that the strike ratio may be different for different macrostates.)

Then, 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—you have enion probabilities that are not only stable over, but are entirely determined by, the system's macrostate, making for an easy derivation of macrolevel behavior.

It is now possible to see clearly the implications of the move, at the beginning of this section, 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.[20] If you know that you are dealing with microconstant probabilities, however, everything changes. Provided that you are right in assuming that the initial conditions of the system in question are smoothly distributed, you

---

19. This is why I did not worry much about which initial condition distribution to use for determining an enion probability. Provided that the candidate distributions are smooth, you do not need to know which is correct. Perhaps it is indeterminate which is *the* correct distribution function; this would be fine, provided that all the good choices are smooth.

20. The more so because you need to know a joint probability distribution over all the different configurations of initial conditions that matter.

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.[21]

The philosophical trick, 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 the reader to Strevens (2003), where they are answered in depth.[22]

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 groundwork for the answer to this question has already been laid in the last section.

21. 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.

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

The great impediment to the assumption of stochastic independence, you will recall, 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. If I am to assume stochastic independence, I need to establish that, in some circumstances, causal interaction does not destroy stochastic independence.

In general, causal interaction *does* destroy stochastic independence. A major exception to this truth concerns probabilities with microconstant evolution functions. Microconstancy, I will show, works to preserve stochastic independence in the face of causal influence.

I will use the tossed coin to illustrate this point. 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 the unconditional probability of one half. To put it another way, 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 evolution function of the last 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. That is, 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 effect of introducing various kinds of collision into 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 5. The

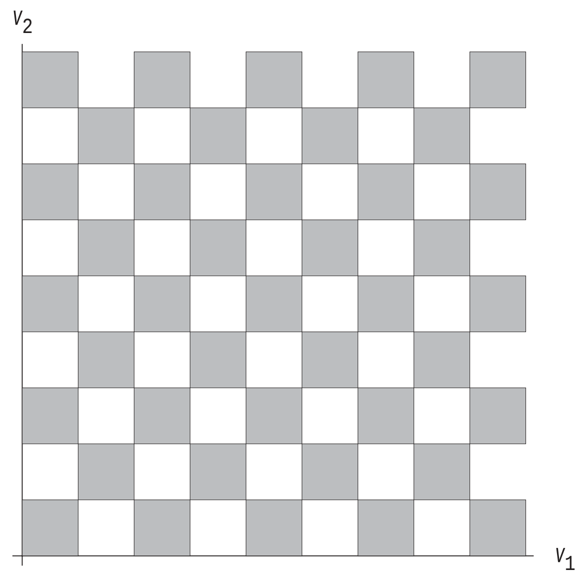

Figure 5: Composite evolution function. Shaded areas indicate pairs of initial spin speeds v_1 and v_2 for two tossed coins that yield a double, that is, an outcome in which both coins land same side up.

diagram is to be interpreted as follows: to find the outcome of any toss, find the point corresponding to the initial speeds v_1 and v_2 of the two coins. If it is gray, both tosses yield the same outcome, so you have a double; if white, they yield different outcomes, so no double.

You should satisfy yourself that figure 5 is the correct evolution function for the case where the coins do not interact. In effect, it is created by taking two copies of the evolution function for a single tossed coin shown in figure 1, turning one on its side, superimposing them, and coloring those sections which are either both gray or both white. The process is shown in figure 6.

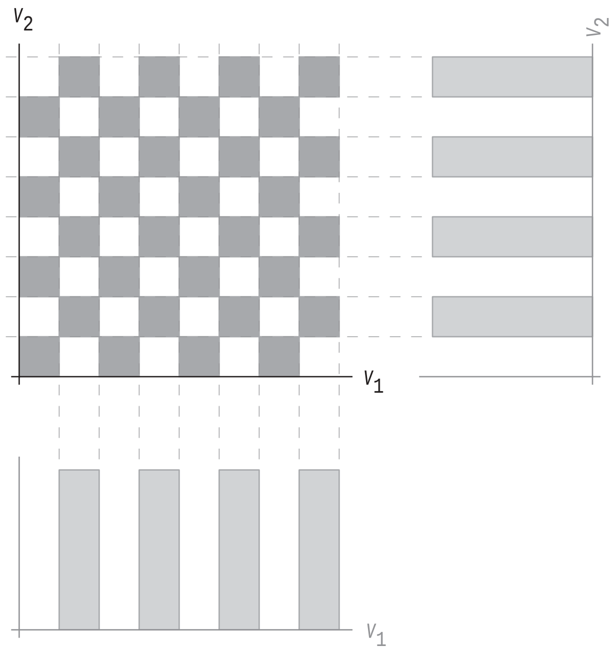

Figure 6: Construction of a composite evolution function. See main text for the interpretation.

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, while promising you that everything else works out just as it should.

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 that, because the composite evolution function for a double in the no-collision case is microconstant with strike ratio one half, then provided that the initial condition distribution over spin speeds is reasonably smooth, the probability of a double will be one half. So I have what I want for stochastic independence.

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.[23]

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.[24] The compos-

---

23. A more complete picture of the connection between causal independence and stochastic independence is presented in Strevens (2003), sections 3.3 and 3.4.

24. 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

ite evolution function for two linearly interacting tosses will look something like figure 7. 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 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 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.[25]

Most interactions are not linear. I had better, then, examine the effect of a nonlinear interaction on the coin tosses. A typical nonlinear interaction produces the composite evolution function shown in figure 8. 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.

---

collision.

25. 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 big 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. As a result, the exceptions will not occur when interactions between enions are sensitive to initial conditions, or chaotic, as in a complex system.

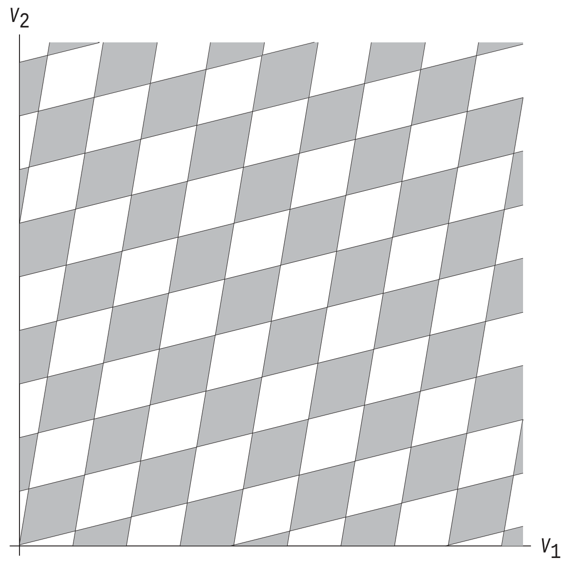

Figure 7: Composite evolution function with linear interaction

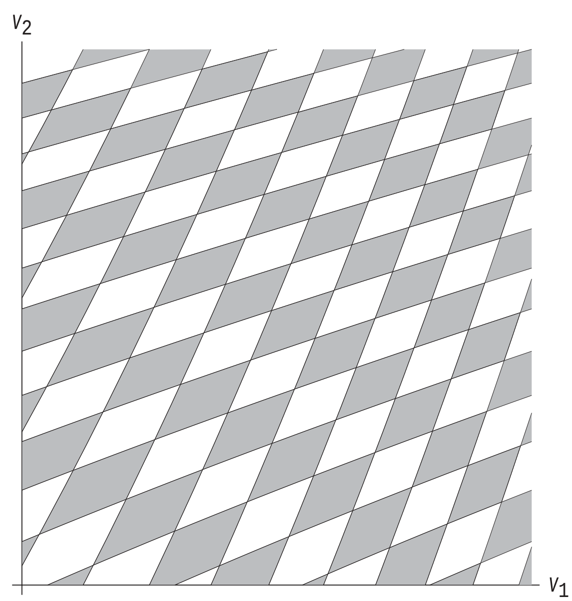

Figure 8: Composite evolution function with nonlinear interaction

Why? The nonlinear interaction I used to construct figure 8 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 is not destroyed by 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 approximately so.[26] 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 non-linear processes studied by chaos theory and so on (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:[27]

1. Interactions the system are either linear, or if non-linear, are approximately linear over small regions of initial conditions, and

2. The evolution functions of enion probabilities are microconstant.

Condition (1) is, as I have just pointed out, 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

---

26. Approximately preserved in the sense that all probabilities will be very close to the probabilities obtained when independence is strictly preserved.

27. The issue of independence in complex systems is more involved than I have made out so far, because to show that enion evolution functions are microconstant, I need further independence results not presented here. 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 25.

enions in most or all complex systems. But I hardly need add, there is much work left to do.

7. Conclusion

The probabilistic supercondition turns out to be true, or at least true enough, if probability is defined in the right kind of 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 macrolevel simplicity. In this way, 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 25 and 27). The same chaos that appears to threaten macrolevel simplicity and tractability is revealed as their enabler and guarantor, their friend.

References

Auyang, S. Y. (1998). *Foundations of Complex-System Theories in Economics, Evolutionary Biology, and Statistical Physics.* Cambridge University Press, Cambridge.

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

Garber, E., S. G. Brush, and C. W. F. Everitt (eds.). (1995). *Maxwell on Heat and Statistical Mechanics.* Lehigh University Press, Bethlehem, PA.

Giere, R. (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.* North Holland, Amsterdam.

Grimmett, G. and D. Stirzaker. (1992). *Probability and Random Processes.* Second edition. Oxford University Press, Oxford.

Hacking, I. (1990). *The Taming of Chance.* Cambridge University Press, Cambridge.

Hopf, E. (1934). On causality, statistics and probability. *Journal of Mathematics and Physics* 13:51–102.

Ornstein, D. S. and B. Weiss. (1991). Statistical properties of chaotic systems. *Bulletin of the American Mathematical Society* 24:11–116.

Poincaré, H. (1912). *Calcul des Probabilités.* Second edition. Gauthier-Villars, Paris.

Porter, T. M. (1986). *The Rise of Statistical Thinking 1820–1900*. Princeton University Press, Princeton, NJ.

Putman, R. J. and S. D. Wratten. (1984). *Principles of Ecology*. University of California Press, Berkeley, CA.

Reichenbach, H. (1949). *The Theory of Probability*. University of California Press, Berkeley, CA.

Salmon, W. (1984). *Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, NJ.

Simon, H. A. (1969). *The Sciences of the Artificial*. First edition. MIT Press, Cambridge, MA.

Sklar, L. (1993). *Physics and Chance*. Cambridge University Press, Cambridge.

Sober, E. (1984). *The Nature of Selection*. MIT Press, Cambridge, MA.

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

Wimsatt, W. C. (1986). Forms of aggregativity. In A. Donagan, A. N. Perovich, and M. V. Wedin (eds.), *Human Nature and Natural Knowledge*. D. Reidel, Dordrecht.