Ontology, Complexity, and Compositionality

Michael Strevens

To appear in *Essays on Metaphysics and the Philosophy of Science*,
M. Slater and Z. Yudell (eds.), Oxford University Press, Oxford.

Abstract

Sciences of complex systems thrive on compositional theories—toolkits that allow the construction of models of a wide range of systems, each consisting of various parts put together in different ways. To be tractable, a compositional theory must make shrewd choices about the parts and properties that constitute its basic ontology. One such choice is to decompose a system into spatiotemporally discrete parts. Compositional theories in the high-level sciences follow this rule of thumb to a certain extent, but they also make essential use of what I call distributed ontologies: divisions of the system into entities or states of affairs each depending on sets of fundamental-level facts that to a large extent overlap. The point is developed using the example of statistical theories in which the probabilities are ontologically distributed.

This paper was originally published with a metaphysical prolegomena, omitted in this version.

1. The Wedding-Cake Ontology

Different sciences have different ontologies—different ways of dissecting the world into individuals, categories, properties. Fundamental physics does particles, chemistry does molecules, biology does cells and organisms and ecosystems, and so on. The list suggests that a certain neat structure is the rule in this grand ontological project: the things at one level are spatiotemporally composed of the things at the next level down. Animals are made of cells, which are made of molecules, which are made of particles…

It follows that the various sciences will cooperate in a straightforward way. Physics, in determining how particles behave, will determine also how spatiotemporal complexes of particles behave—and these rules are the laws of chemistry. Chemistry, in determining how molecules behave, will determine also how complexes of molecules behave—and these rules are some of the laws of geology, cytology, and so on. Each science builds its subject matter from the parts supplied by the science below, and deduces its laws from the laws of those lower-level parts, also supplied by the science below. This wedding-cake picture of the unity of science was made famous by Oppenheim and Putnam (1958); I am here applying the metaphor to its ontological component.

Why should the wedding cake be the rule? One advantage of the cake's structure is that it promises to supply to the scientific enterprise a copious array of *compositional theories*. A compositional theory earns that sobriquet by providing a toolkit for understanding a wide range of systems within a given class (using the term *system* broadly to encompass molecules, organisms, ecosystems, minds). It divides the system into parts and assigns dynamical properties to these components that, when aggregated, predict the behavior of the system as a whole. If you have a compositional theory of a certain kind of system, then, you do not need to theorize anew for each instance of that system. You have a procedure you can use to predict and explain the behavior of every configuration of the system within the allowed range.

Warning: compositionality of theories is not the same thing as composition of objects. The one is a property of a system of representations, in virtue of which the behavior of the whole is derived from, and hence predicted and explained by, the behavior of the parts; the other is a metaphysical relation, in virtue of which the existence of the whole is secured by the existence, properly arranged, of the parts. This paper talks about both compositionality and composition; there is, however, no straightforward relation between the two.

Compositional theories are, quite obviously, greatly desirable; indeed, it is hard for science to make much headway against the world's complexity without them. A compelling attraction of the wedding-cake picture is that it yields a template for building compositional theories. It is spatiotemporal decomposition: divide your system into spatiotemporal parts, then apply the rules that govern the parts, aggregating to deduce the behavior of the whole.

As I will argue in the next section, however, spatiotemporal decomposition frequently leads to a complexity catastrophe: due to the sensitivity and combinatorics of the interactions between parts, it is impossible in practice to aggregate the behavior of the parts to obtain a description of the system that they compose.

The solution is to put aside the wedding cake's spatiotemporal ontology and to find other ways to build compositional theories. In this paper I examine one such approach, a statistical route that might seem at first to help itself to a traditional wedding-cake taxonomy, but that on closer examination slices up the systems to which it applies in a way that is not at all spatial, and that consequently cuts across the wedding-cake division of the world into parts.

The result is quite compatible with reductionism: it is physicalist and fundamentalist, meaning that everything is made up of physical stuff and everything that happens happens because of the ways that the fundamental laws of nature push around such stuff. The bottommost layer of the cake is, then, firmly in place. From there on up, however, the theoretical patisserie takes on a variety of stranger and more intricate forms. The familiar layers are inarguably there—molecules, cells, organisms, and the rest—but there is much else squeezed into the interstices.

2. Compositionality and Complexity

Choose a class of systems: halide molecules, coronaviruses, predator/prey ecosystems. A compositional theory of such a class divides the systems into parts, assigns the parts properties in virtue of which they behave in certain ways, and then aggregates the behavior of the parts to derive, for predictive or explanatory purposes, the behavior of the whole—for any system in the class.

Newtonian gravitational theory, for example, divides the world into objects, assigns those objects masses in virtue of which they exert and experience gravitational force, and then by way of a principle of aggregation—the rule that individual forces are to be summed as vectors—makes predictions about the movements of a system of such objects.

In principle, Newtonian theory can predict the behavior of any system in its scope; there lies its compositional strength. In practice, great difficulties may emerge in the aggregation. Given only three gravitating bodies, the mathematics of even approximate derivations can, if the bodies are of similar mass, become complex, in part because such systems may be highly sensitive to initial conditions. For greater numbers of bodies, combinatorics piles on difficulty: a small change here can make for moderate changes in many other places, which in turn change things further here—and so on. This is a complexity explosion; it limits the usefulness of even so practically important a compositional theory as Newtonian gravitation.

What you might call the *aggregation problem* arises again and again at many different levels in the sciences. The quantum chemistry of large atoms is difficult enough; that of large molecules is more challenging still. Modeling the complex genetic networks at work in embryological development is fiendishly hard. Predicting many of the significant consequences of interacting human minds—housing bubble collapses, Hollywood megahits, popular revolutions—is quite beyond us. Sometimes, however, compositional theories of complex systems lie within our grasp: we can build successful models in statistical physics, evolutionary biology, and computational psychology. What many of these models have in common is their being non-spatiotemporal in some aspect of their underlying ontology: they decompose systems into parts or properties, some of which are not spatiotemporally discrete units or anything close. The need to solve the problem of aggregation, then, imposes on higher-level sciences ontologies that depart at least in part from the wedding-cake ideal.

Is this because theories with a wedding-cake ontology are especially prone to the aggregation catastrophe? Is there something about decomposing a system into spatiotemporal parts that renders the calculus of aggregation especially intractable? Perhaps not. It might simply be that the great majority of lower-level ontologies are unsuited to tractable aggregation, and that since spatiotemporal cohesion is a strong constraint and so wedding-cake ontologies are only a small proportion of the possible ontologies, we ought not to have expected them to be particularly successful.

Since the spatiotemporal constraint is chosen not at random but because it simplifies the organization of science, that is a somewhat weak explanation of the aggregation problem. I think we can do better: the nature of the wedding-cake ontology to some extent does explain why it, in particular, suffers from the aggregation problem.

A rule of aggregation pulls together the connections between the parts into which a system is decomposed. In a compositional theory with a wedding-cake ontology, then, the aggregation rule will calculate the net effect of the relevant connections—causal connections, let me suppose—between its spatiotemporal parts. The aggregation problem tends to arise for wedding-cake theories because of the combinatorial complexity and sensitivity of these relations.

I call the relations between spatiotemporal parts *sensitive* because their effect on the aggregate is sensitive to small changes in the state of the part.

I do not mean "sensitive" in the chaotic sense, but in a much weaker sense: the effect on the aggregate is not wholly independent of the small changes, or in other words, small changes have some effect on the aggregate. For example, although the gravitational force exerted by an object does not depend chaotically on the object's position, it does depend on exact position: move the object slightly and the force exerted at any point changes slightly, and so the aggregate force exerted on an object at that point changes slightly. The same is true for mass: a slight change in mass means a slight change in force, rather than none whatsoever.

The relations between spatiotemporal parts are *combinatorially complex* because, as you increase the number of objects, the number of relations to keep track of increases. I do not mean that it increases exponentially, or even non-linearly—just that there is an increase.

These two properties tend to lead, because of increasingly intricate whorls of dependence, to complexity explosions: as in the case of gravitation, small alterations in one place proliferate quickly to many other places, where they further nudge conditions at the original locus of change. The web of mutual influence becomes a hopeless tangle.

For this reason, wedding-cake ontologies typically prove to be unsuitable ground for compositional theories of complex systems. What theories do better?

3. Enion Probability Analysis

A broad class of compositional theories dissolve the aggregation problem by taking a statistical approach: they divide the system into individuals such as molecules or organisms (which I call "enions"), they assign probability distributions to the behavior of the enions, and they derive the behavior of the whole by aggregating the relevant probabilities. I call this strategy *enion probability analysis* or, for short, EPA (Strevens 2003).

Enion probability analysis cannot be used to model just any property of a complex system: it is limited to tracking aggregate properties of populations. You can use it (when conditions are favorable) to model changes in the number of tortoises in an ecosystem, for example, or the number of male tortoises in the system between one and two years of age, but not to follow the day-to-day movements of a particular tortoise. (For the latter purpose, I know of no compositional theory that does not suffer from the aggregation problem.) This is a relatively small subset of the properties of a system, but it is an important one, containing the things we need to predict and explain in order to undertake many of the projects falling within the ambit of statistical physics, evolutionary biology, ecology, economics, and so on. Consequently, EPA has a significant role to play in accounting for the predictive and explanatory success of the compositional sciences of complex systems, and a philosophical understanding of EPA has significant lessons to teach about the ontologies of successful compositional theories—and thus about the ontology of the world.

In the boreal forest of Canada, snowshoe hares eat underbrush and Canadian lynxes eat snowshoe hares. The populations of the two species cycle regularly: every eight to eleven years, the hare population booms (for reasons about which there is still considerable argument) and the lynx population, sustained by regular helpings of *lapin à la mode*, booms soon after. Then the hare population crashes and the lynx population follows on its heels.

Ecologists wishing to understand this famous predator/prey population cycle build mathematical models of the system. Although the models themselves tend to have a deterministic mathematics, their structure can be understood as rooted in stochastic foundations, as it was by Alfred Lotka, who gave his name to these "Lotka-Volterra" models (Lotka 1925). This is the statistical method that I am calling enion probability analysis; after describing its application to lynxes and hares, I will examine its presumptive ontology.

A Lotka-Volterra model tracks populations, attempting to predict the size of future populations—say, the numbers of lynxes and hares in a month's time—from present populations. Some sophistications are possible: weather patterns or the luxuriance of the vegetation can be incorporated into the model, or subpopulations can be tracked (such as sexually mature female hares). But let me put all of this aside to focus on the basics.

Assume that the model has only two variables, representing the number of lynxes and the number of hares. What is wanted is a set of equations that relates the values of the variables at one time to their values at a later time.

It might seem that a complexity explosion is imminent. What happens to an individual hare depends on small wrinkles in local matters of fact. Forage under the old pine this morning and you will be lynx food. Head instead for the spruces, and you will live to graze another day. Tracking the fates of individual hares requires, then, an accurate record of particular movements and the principles that drive them. That, surely, must add up—when hundreds or thousands of hares and many lynxes are involved—into something very complex, certainly more complex than a set of simple equations relating only two variables. Or so it would appear.

Lotka's approach gets around these difficulties using the following recipe. To the members of each population or sub-population in the ecosystem, assign probabilities for the kinds of outcomes that make a difference to populations: birth, death, perhaps migration. Assume that these "enion probabilities" are independent of one another. Use the law of large numbers—assuming here that populations are reasonably large—to derive frequencies for various events equal to the enion probabilities. If, for example, there is a 0.05 probability that any given hare is killed by a lynx over the course of a month, assume that one-twentieth of the hare population is lost to lynx predation every month. You now have a table of per capita rates: the rate of hare reproduction (per hare), the rate of hare predation, the rate of lynx reproduction, the rate of lynx death, and so on.

Some of these rates depend on other variables. The rate of hare predation depends, for example, on the population of lynxes: more lynxes means proportionally more hares served up for dinner. These dependences typically (though not invariably) end up in the model. It is crucial, then, that they bring with them into the model's equations only quantities that the model is constructed to represent. In the simple case at hand, the rates should depend only on the total number of hares or the total number of lynxes in the system, or both, or neither. (If the rates as a matter of fact depended on subpopulations, the model would have to track those subpopulations—young female hares, for example—thereby becoming more complex.) The rates will satisfy this independence requirement just in case the enion probabilities satisfy the requirement. The probability of a hare's being killed by lynx over the course of a month should, for example, depend only on the number of lynxes (and perhaps the number of hares).

Let me suppose that the necessary enion probabilities, and so the rates, are fully determined by the physical facts. (Strevens (2003) tries to specify exactly what fundamental states of affairs fix the facts about the probabilities.) Then, with the probabilities' existence secured, it is straightforward to write down two equations representing, respectively, the rate of change of the hare and lynx populations. The equation for the hare population might, for example, set the change in the population equal to the rate of hare reproduction (multiplied by the current hare population) less the rate of predation (multiplied by the current hare and lynx populations) less the rate of hare death from other causes (multiplied by the current hare population). The lynx equation will do something similar. Together, these equations make up a Lotka-Volterra model, which can be simulated or solved and then either tested against observed population changes or used to explain those changes.

That is one way—the EPA way—to model a complex system, deftly avoiding the problems posed by the complexity explosion.

The enion probability analysis of the lynx/hare system appears at first to have a wedding-cake ontology: it proceeds by dividing the system into enions, individual hares and lynxes, that are spatiotemporally discrete. A closer look shows, however, that although the enions are indispensable, the most important elements of EPA—not the enions themselves, but the probabilities that describe the enions' behavior—are individuated rather differently.

4. Where Did the Interactions Go?

An explosion was averted. Every hare and every lynx has manifold interactions with its environment; changes in the hare and lynx populations are nothing but the aggregate consequences of these interactions; and aggregating the interactions looks to be the sort of intractable task that threatens to sink the sciences of complex systems—yet EPA proceeds without a hitch. What happened to the interactions? How were they, in the end, so easily aggregated? By answering this question in part, I will show you that the ontology of EPA does not conform to the model of the spatiotemporally organized layer cake.

Explaining the complexity explosion in wedding-cake models above, I attributed it to two properties of the relations between the parts of those models: sensitivity and combinatorial complexity.

An inter-part relation is sensitive if small changes in the state of a part make for a difference (perhaps slight) in the relation. It gives rise to combinatorial complexity if the number of relations the model must keep track of (when aggregating) increases with the number of parts.

In a Newtonian model the relevant relations are the connections that determine the force exerted by a spatiotemporal part—the forces that must be aggregated, that is, to determine the behavior of the system as a whole. In an EPA model, they are the relations that determine enion probabilities such as the chance of hare death—the probabilities that must be aggregated to determine the behavior of the system as a whole.

Unlike the Newtonian relations, the EPA relations are not sensitive: enion probabilities are not affected by small changes in the state of the relevant enion, such as a shift in position. Indeed, by the independence requirements stated above, they depend on almost nothing about any enion—the probability of a hare's death is not affected by its position or by what happens to any other hare or lynx.

Nor are the EPA relations combinatorially complex: the enion probability analyst must keep track of one set of probabilities per enion type, and that is all. In the lynx/hare system that amounts to two sets, one for lynxes and one for hares, regardless of the population of each. A system with many lynxes is consequently no more difficult to model than a system with a few. Indeed, large populations make things simpler, by making it more likely that actual behavior will correspond to statistically expected behavior.

These negatives—the lack of sensitivity and of combinatorial complexity—go some way toward explaining why EPA models do not suffer from a complexity explosion, but there is much more to be said. The source of enion probabilities' insensitivity is particularly important: the key to understanding the power of EPA is, I think, to understand why there is so little dependence between the statistical behavior of enions and their exact or even approximate states. To put it another way, what should be explained are the independence assumptions upon which the applicability of EPA depends:

1. Enion probabilities depend only on population-level variables of the sort tracked by statistical models.

2. The outcomes to which enion probabilities are attached are stochastically independent.

This is a project I tackle in Strevens (2003); an overview is given in Strevens (2005). The complete story is not something that I will undertake to give here. For the purposes of understanding the implications of compositionality for ontology, it will be enough to answer an easier question: where, in EPA, do the interactions go?

We know that there are many interactions between hares and other hares, lynxes, and their environment. Try to track these many interactions and you will generate—so I have supposed—an immediate combinatorial catastrophe. Enion probability analysis, by representing the behavior of the system as a whole, represents the aggregate effect of these interactions. Yet it somehow, in its formalism, avoids having to represent the interactions explicitly—and so avoids having to aggregate them formally, bypassing the aggregation problem that would result. The interactions are packed away in some place in which they cannot get out of hand. Where?

Consider the probability of a hare's being killed by a lynx over the course of a month—the number that determines (more or less) the rate of hare predation. On what features of the lynx/hare ecosystem does the value of this number depend? What aspects of the system go into determining that the probability of hare death per month is, say, 0.05 rather than 0.1? The relevant factors include the total number of lynxes, the techniques that lynxes use to hunt hares, the techniques that hares use to avoid lynxes, the nature of the vegetation in the habitat, and more. Change any of these things in significant ways, and the magnitude of the probability of hare death will surely change.

The dependence of the hare death probability on the first of the enumerated factors—on the number of lynxes—is represented explicitly in the EPA model. The effect of lynx number on the probability is, in other words, "externalized". What about the rest? They are entirely internal to the probability, which is to say that their net effect is built into the probability—in formal terms, built into the 0.05; in metaphysical terms, built into the physical probability quantified by that number. As a consequence, the model need not explicitly take these interactions into account.

There I will pretty much leave the explanation of the miracle of EPA, taking away two claims about enion probabilities. First, the probabilities are not physically separate and independent entities. They are attached to physically independent entities—to different hares—but they physically overlap, since numerically identical states of affairs contribute to many distinct probabilities. The lynxes' tactics, for example—ultimately a matter of lynx brain configuration, I suppose—help to determine each hare's probability of death, as do many other shared aspects of lynx makeup. Follow the death probabilities for different hares down to the fundamental level, then, and they converge on many of the same fundamental-level facts. In other words, the reduction or supervenience bases for any two hares' death probabilities overlap. Contrast this with a standard wedding-cake theory such as Newtonian gravitation, on which each object contributes to and experiences the net gravitational field in virtue of a wholly spatiotemporally intrinsic property, its mass. The principles for enion probabilities' individuation bear little resemblance to the principles for the construction of the wedding cake.

Second, it is this extrinsic and overlapping quality that makes it possible for EPA to avoid an explosion of complexity, opening the door to a compositional theory that shrugs off the aggregation problem. How so? I venture that a compositional theory must, in order to be useful, individuate a system's parts and properties so that they are in some sense largely independent. Enion probability analysis does not, in its delineation of the determinants of enion behavior, divide the world into factors that are physically independent. It does, however, divide the world into determinants of behavior that are *stochastically* independent, and here lies its power: the rules for aggregating stochastically independent determinants of behavior are far more tractable than, in any interesting system, the rules for aggregating physically independent determinants.

5. Ontologies of the High-Level Sciences

The high-level sciences are ontologically pluripotent. They make great use of physical individuation, that is, individuation according to the spatiotemporal wedding-cake conception: even in EPA, the bearers of probability are typically physically discrete entities such as molecules, animals, or people, and the ultimate aim of models is to track the statistical movements of such things.

Yet at least as important are what you might call distributed ontologies, that is, individuations into entities—things and reifications of the tendencies of things, such as causal dispositions and enion probabilities—whose presence is determined by configurations of fundamental physical facts that overlap, so that the same facts contribute to many such entities. Enion probabilities are this paper's paradigms of distributed ontology.

The entities distinguished by the first, spatiotemporal kind of ontological decomposition—the enions themselves, the lynxes and hares—are physically independent but interact with one another in many complex ways. The entities distinguished by the second, distributed kind of ontological decomposition—the enion probabilities, such as the probability of a certain hare's death over the course of a month—are physically overlapping, but stochastically independent and therefore easily aggregated.

These two ontological schemes are not rivals, but rather work together within a single modeling technique in population ecology, serving up a compositional theory that solves the problem of aggregation. The wedding-cake ontologists are right to think that spatiotemporal individuation has been essential to creating compositional theories of the high-level sciences, but wrong in thinking that it has been sufficient. To tackle the sciences of complex systems we need what is, in a mild sense, ontological pluralism.

Can the same be said of other kinds of compositional theories? Let me give you two examples.

The first is spectral analysis in wave theories of various aspects of nature. In the high-level sciences, there are sound waves, ocean waves, waves on the strings of musical instruments, seismic waves, and more—where in each case, the wave is a movement of an underlying medium. In spectral analysis, the medium's movement is decomposed into waves of different frequencies, as when the motion of a vibrating string is decomposed into a tone and various overtones. These waves coexist in the same medium—in the same volume of air, or earth, or water—and indeed in the same movements of that medium; consequently, the fundamental-level matters of fact on which different waves in a spectral decomposition depend are largely identical. The waves form a distributed ontology.

For the predictive and explanatory purposes of many wave models, the only aggregation required is the addition of the effects of these different frequencies, which is accomplished by the straightforward process of linear superposition. The wedding-cake alternative, in which a model keeps track of the movements of different parts of the medium—different segments of a vibrating spring, or different volumes of air or water—is far more difficult to implement. A distributed ontology brings compositional modeling within reach.

(Our theories of the forces that come together to create waves are, however, often sensitive and combinatorially complex. As a result wave theories will, for certain predictive purposes, suffer from a complexity explosion. That is why quantum chemistry is computationally so difficult.)

Another, more speculative, example is belief/desire psychology. Most of us would guess that the facts underlying beliefs and desires—the facts that make it the case that I believe that there is rabbit for dinner or that I desire to wash it down with a glass of wine—are to some extent distributed across the brain in an overlapping way. The propositional attitudes comprise a distributed ontology.

Belief/desire psychology is also a remarkably effective compositional theory of thought (Dennett 1987). The principles of composition are quite familiar to us, but the relation of the whole to the underlying facts remains, for now, opaque. It seems that belief/desire psychology solves an aggregation problem by way of a distributed ontology, then, but we cannot as yet be sure.

Some philosophers have suggested that beliefs and desires not be taken ontologically very seriously at all (Churchland 1981; Dennett 1987), in part, I would guess, because their distributed nature lends them a certain insubstantiality, a lack of proper placement within the great wedding cake of science. Perhaps contemplation of the role of distributed ontologies elsewhere, in wave theories and in EPA, can solidify the attitudes' status, both scientific and metaphysical.

References

Churchland, P. M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy* 78:67–90.

Dennett, D. (1987). *The Intentional Stance*. MIT Press, Cambridge, MA.

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

Oppenheim, P. and H. Putnam. (1958). Unity of science as a working hypothesis. In H. Feigl, M. Scriven, and G. Maxwell (eds.), *Concepts, Theories, and the Mind-Body Problem*, volume 2 of *Minnesota Studies in the Philosophy of Science*. University of Minnesota Press, Minneapolis.

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

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