13 Constitutive Overdetermination

L. A. Paul

Our best philosophical and scientific pictures of the world organize material objects into a hierarchy of levels or layers—microparticles at the bottom, molecules, cells, and persons at higher layers. Are objects at higher layers identical to the sums of objects at lower layers that constitute them? (Note that this question is different from the question of whether *composition*—as opposed to *constitution*—is identity.) As I will define the positions, *reductionists* are monists who claim that constitution is identity and *nonreductionists* are pluralists who deny it.

Paul Oppenheim and Hilary Putnam (1958) present the classic characterization of the layered world model, arguing that the world has multiple layers of objects ordered by the mereological relation of spatiotemporal part to whole.[1] In Oppenheim and Putnam's reductionist model, sums of lower layer objects are identical to objects at higher layers and objects at higher layers are exhaustively (i.e., without remainder) decomposable into parts that are objects at lower layers.

Jaegwon Kim describes how the model of the natural world is stratified into hierarchical layers via mereological compositionality:

> The Cartesian model of a bifurcated world has been replaced by that of a layered world, a hierarchically stratified structure of ''levels'' or ''orders'' of entities and their characteristic properties. It is generally thought that there is a bottom level, one consisting of whatever microphysics is going to tell us are the most basic physical particles out of which all matter is composed (electrons, neutrons, quarks, or whatever). (1993, p. 337; see also Kim 2002, pp. 3–4)

As Kim notes, the layered world was proposed as a substitute for Cartesian dualism, according to which there were mental *substances* that were ontologically independent of physical substances; the layered world model rejects mental substances and organizes the physical domain into layers.

Oppenheim and Putnam categorize and present the objects belonging to different layers this way (Boyd 1991, p. 409):

6. social groups
5. (multicellular) living things
4. cells
3. molecules
2. atoms
1. elementary particles

Any whole that is exhaustively decomposable into parts belonging to layer L is counted as also belonging to L, so objects at each layer include all objects at higher layers. According to Oppenheim and Putnam, different layers correspond to different ways we can mereologically carve objects, but there is no ontological increase or decrease as we move between layers. This suggests that we can model their ontology O of what there *is*, as opposed to how to *carve* what there is, as nonlayered:

O = particles / atoms / molecules / cells / (multicellular) living things / social groups.

Contemporary reductionists are unlikely to accept all the details of the Oppenheim-Putnam model. For example, they might hold that there are additional layers, that the ordering might not always be linear, and that it is unclear where some objects belong in the hierarchy. However, they will embrace its defining characteristics of ontological minimalism and of modeling the world in terms of a hierarchy of objects ordered by exhaustive spatiotemporal mereological composition where sums of lower layer objects *are* the higher layer objects they constitute.

Now, a model of the world as layered can be independent of reductionism: nonreductionists can accept layers even if they deny the reduction of objects at higher layers to sums of objects at lower ones. Why deny the reduction? Because the properties of higher layer (constituted) objects seem to be different from the properties of their lower layer (constituting) sums of objects, and by the Principle of the Indiscernibility of Identicals, things with different properties cannot be identical. The thought is that, for example, statues are valuable, sums of particles of bronze are not; persons are handsome, sums of cells are not (Fine 2003); protein molecules have their carbon atoms accidentally, but sums of atoms that include carbon atoms

have their carbon atoms essentially (e.g., Wiggins 2001; Lowe 1998).[2] Ultimately, nonreductionists mean something ontologically substantial when they use the term 'layer', while reductionists do not.

Thus, nonreductionists make ontological distinctions where reductionists do not. We have the object that is the sum of cells and other matter, and the object that is the person—the nonreductionist holds that these differ (for they differ in their properties) whereas the reductionist identifies them.[3] When the nonreductionist rejects the claim that the sum of the lower layer objects is identical to the higher layer object, she holds that the sum constitutes but is not identical to the higher layer object. Moreover, according to the nonreductionist, in cases when the sum constitutes the higher layer object, it will not share all of its parts with the higher layer object. The sum of cells and other matter constitutes the person, and so the sum and the person share their microparticles and hence some of their parts. But they do not share *all* of their parts—for example, the head of the person is not identical to the head-shaped sum of cells and other material, since their properties differ.[4] The nonreductionist should thus distinguish between constitution and composition: *composition* is the familiar mereological fusion relation of parts to whole and is analogous to (or is) identity, while constitution requires a separate explication and is neither identity nor analogous to it.[5] (This distinction raises several interesting and delicate issues; see my manuscript ''The Ontology of Objects'' for a discussion of the matter.)

Rejecting the identification of higher layer objects with their constituting sums of lower layer objects implies that there are robust ontological differences between objects and their constituting sums. This means that, according to the nonreductionist, *in addition* to having a person, we have a sum of cells, a sum of atoms, a sum of particles, and so on. All of these entities (somehow) occupy the same region of spacetime and involve the same matter. When this nonreductionist stance is combined with a picture of the world as layered, we end up with objects in higher order layers that are irreducible to objects in lower order layers. This suggests that nonreductionists model their *ontology* as layered. For the nonreductionist, particles and molecules and so on, belong to lower layers, but *contra* the Oppenheim-Putnam interpretation, there are ontological differences as we move up layers (there is an increase in the number of objects in the world as each layer is added).

Why would nonreductionists accept a version—albeit a modified version—of the reductionists' layered model? First, perhaps, because having layers of some sort in the world seems to be a fairly natural world view: early versions of a (nonmereologically) layered model were proposed by emergentists, whose views bear an affinity to the nonreductionists' views.[6] Second, having a layered world model allows us to represent relationships between the smaller and larger objects referred to by theories of natural and social science in a way that is consistent with contemporary philosophical work on supervenience and related topics. Third, and perhaps most important, if the model is rejected, it is unclear what to put in its place. Recall the context in which the layered world model was adopted—it was seen as the antidote to substance dualism. If the layered world model is rejected, it seems we must revert to a mysterious Cartesian-like picture—with the unwelcome twist that there are now many material substances (one for each former layer). It is surely incorrect to hold that nonreductionism about objects related by constitution implies some sort of substance pluralism. (Compare the thought that nonreductionism in mind implies substance dualism.)

But for the nonreductionist who accepts the layered world, trouble comes quickly. Assume that the nonreductionist can adequately motivate her view and can explain enough about the constitution relation for us to understand how sums are numerically distinct from what they constitute and how sums and higher layer objects share their matter and spatiotemporal region without sharing all their parts.[7] Even if all this is accomplished, the nonreductionist still faces the problem of symmetric causal overdetermination.[8] (The substance pluralist would also face the problem of symmetric causal overdetermination, so rejecting the layered world model in favor of substance pluralism won't help.)

Those familiar with the debate over mental causation will recognize that the problem is the sort of problem that arises for nonreductionists about the mind (except that here the focus is on objects rather than properties). When philosophers of mind hold that mental properties are not identical to or reducible to physical properties, they need to explain the causal efficacy of mental properties given the sufficiency of physical properties for our actions, thoughts, beliefs, and so on; similarly, nonreductionists about layers need to explain the causal efficacy of objects that are constituted by sums of smaller objects. (Philosophers of mind sometimes argue that the

overdetermination problem for nonreductionists about the mind arises in the context of the *exclusion thesis*: if an event *c*, together with associated background conditions and laws, is entirely sufficient for an event *e*, then no additional event c^* is a cause of *e*. Although the exclusion thesis is trying to capture a compelling intuition about the way causation works, the thesis itself is not particularly plausible, since some sorts of overdetermination—such as the case where the three golf balls shatter a window, or perhaps cases involving overdetermination between parts and a whole—are surely possible.)[9]

Imagine that I hit a tennis ball and it bounces off my racquet at a speed of 100 miles per hour. What causes this effect? Well, *I* do. (I am eliding the difference between object causation and event causation here.)[10] But even though I cause the bouncing of the tennis ball, I am constituted by a sum of cells and other matter, and this sum causes the bouncing as well. In fact, according to the nonreductionist, there are *many* different objects causing the bouncing. Just for starters, I, the sum of elementary particles that constitutes the sum of atoms, the sum of atoms that constitutes the sum of molecules, the sum of molecules that constitutes my cells and other matter, and the sum of cells and other matter that constitutes me, all cause the bouncing. How can we explain this? The ball didn't bounce off my racquet at 500 miles per hour, and yet the action of each of the five objects causes a bouncing of 100 miles per hour. (Strictly speaking, multiple effects are overdetermined, since the sum of molecules that constitutes the tennis ball moves at 100 miles per hour, as does the sum of particles that constitutes the tennis ball, and so on; but I ignore this complication for simplicity's sake.) The same worry arises when I write a book: what causes the book to be written? Do I, the sum of my cells and other matter, the sum of my fundamental particles, and so on, all cause the writing of the book? The point here, of course, is that it seems as though nonreductionism generates massive amounts of symmetric causal overdetermination, which I shall call "constitutive overdetermination."

In contemporary discussions of causation, standard cases of symmetric causal overdetermination are defined (roughly) as cases involving multiple distinct causes of an effect where the causation is neither joint, additive, nor preemptive (and it is assumed the overdetermining causes do not cause each other). Common examples involve cases where three bullets simultaneously enter a victim's heart, or three golf balls simultaneously shatter a

window. Each cause makes exactly the same causal contribution as the other causes to the effect (so the causal overdetermination is *symmetric*); each cause without the others is sufficient for the effect; and for each cause the causal process from cause to effect is not interrupted.

Many of those who contribute to the literature on philosophical theories of causation think that symmetric causal overdetermination is problematic but peripheral.[11] There are several reasons for this. First, such cases are supposed to be rare or nonexistent in the actual world. Second, it is supposed that even if cases of overdetermination do exist, most of them are artifacts of the coarse or robust individuation of effects. If so, effects can be individuated more finely so as to remove the overdetermination: for example, in the case where three golf balls shatter a window, one can hold that the shattering is a slightly different effect when there are three golf balls as opposed to one or two. This makes the overdetermination in the golf ball case disappear, turning the case into a case of joint causation that was misidentified under a too-robust individuation of the shattering. Third, since our commonsense intuitions are vague as to how to understand symmetric overdetermination in the context of a reductive analysis of causation, some hold that true cases of symmetric overdetermination should be treated as spoils to the victor; that is, if better understood, more central cases are solved by an analysis, whatever verdict given by this analysis on symmetric overdeterminers should simply be accepted.[12] Finally, there is a special sort of overdetermination that may result (if composition is not identity but merely analogous to identity) such that the collection of objects that compose *c* and the whole that is *c* overdetermine *e*. The thought here is that it is unclear how concerned we should be about this sort of overdetermination or whether we should even balk at accepting it. The reason for a laid-back approach is that compositionality—the special mereological relationship of ''almost identity'' between things and the whole they compose—is so intimate that this sort of overdetermination is not metaphysically troubling.[13]

But none of this will help the nonreductionist. First, the sort of overdetermination implied by her view would not be rare—it would be the norm. Practically every instance of the sort of garden-variety causation involved in our commonsense, scientific, or philosophical claims would be overdetermined! Second, merely individuating effects more finely will not solve the problem, since the effects in question will not be any different if there is no overdetermination. The bouncing of the tennis ball is not

changed by the fact that both I and my constitutive sum of cells hit the tennis ball; the bouncing would be unchanged if, for example, I cause the bouncing but my sum of cells somehow does not. Third, the ''spoils to the victor'' solution is only plausible (if it is plausible at all, which I doubt) in a context where we are performing a conceptual analysis and the troublesome cases are not the central sort of case handled by the analysis. As we have seen, the nonreductionist's view makes symmetric overdetermination *the most* central and common sort of causation around. Finally, the nonreductionist denies that constitution is merely composition, that is, she denies that when the sum of lower layer objects constitutes the higher layer object that this is just the relation of the sum composing the higher layer object. Indeed, the nonreductionist denies that the sum and the object it constitutes share their larger spatiotemporal parts, so the constitution relation cannot be the composition relation. Since constitution is not composition, the claim that overdetermination is not problematic in the special case of compositionality will not help the nonreductionist. Result: the nonreductionist cannot ignore the problem of constitutive overdetermination.

Of course, the nonreductionist could respond by giving up her views and granting that objects are identical to their constituting sums. If objects are identical to their constituting sums then the problem with overdetermination that the nonreductionist faces is avoided. But there are other options that the nonreductionist might prefer to consider.

One option is to be an eliminativist about objects at all higher layers. (The nonreductionist is unlikely to find this option attractive, but may feel pushed towards it given her view that should objects like statues and persons exist, they must differ from their constituting sums.) Outright eliminativism is clearly more radical than reductionism: instead of holding that objects at higher layers exist but are identical to sums of objects at lower layers, objects at higher layers are *eliminated* in favor of objects at lower layers.[14] The usual assumption, if this is done, is that the privileged layer is that of microphysics, following Oppenheim and Putnam. (See Schaffer 2003a for a discussion of privileged layers.)

The eliminativist argues for the elimination of higher layer objects by claiming that all or most objects do not compose sums. Because there are no sums of objects, there are no higher layer objects that they constitute, so there is no constitutive overdetermination. For thoroughgoing eliminativists, nothing except fundamental particles, or fields, or whatever is

most fundamental exists. I take eliminativism to be an interesting nonstarter, since eliminating proteins, rocks, and stars must be a position of last resort. Surely it is more acceptable to have widespread constitutive overdetermination—or even epiphenomenalism—than to accept that such higher layer objects do not exist. Eliminativism suffers from more than just radical implausibility: if there exists an infinitely descending series of layers, such that for every elimination we allow, another beckons, then we never end up with any fundamental existents. We eliminate endlessly.[15]

There are two more options for the nonreductionist to consider—some sort of widespread causal supervenience, or true symmetric overdetermination. Consider the first option: what if the sort of widespread symmetric causal overdetermination we seem to see is not true symmetric overdetermination after all? Instead of true overdetermination, there are multiple layers of causal relations, where the causal relations at higher layers nonreductively supervene on causal relations at lower layers. In this picture, there is no overdetermination because causal processes initiated by higher layer objects only cause higher layer effects and lower layer objects only cause lower layer effects.

I think such a picture has some nice intuitive appeal but is ultimately a nonstarter. Do we really think there are all these different layers of causation? As with symmetric overdetermination, this view seems to produce a lot more causation in the world than we ever knew about or wanted. Not only do we have microphysical causation, but we have numerically different processes of chemical causation, individual causation, artifact causation, mental causation, and so on. This is unattractive. But there is a more serious problem in the offing.

The way I described the problem with overdetermination and a layered world may make you think that the right way to view the problem is to hold that things caused other things on the same layer; that is, persons cause things like bouncings of tennis balls at their layer, while sums of cells cause things at theirs. But in fact, the problem with overdetermination is not layer specific—causation seems to cross layers. As Kim (1998, pp. 42–43) puts it, when you take an aspirin to relieve a headache, intuitively, you also causally intervene in the brain process upon which the headache supervenes.

Similarly, we have strong intuitions that it is possible for a person to causally intervene at the layer of fundamental particles and for a society to

causally intervene at the layer of individuals. If I remove a carbon atom from a carboxyl group, then I cause a change in a protein *and* in the sum of atoms that constitute it, and this cannot be explained away by suggesting that I am only causing something at a higher layer which supervenes on my sum of atoms causing a change in another sum of atoms. It seems importantly right to say that I am causing a change in the protein, *and* that *I* (not just my sum of atoms) am also causing a change in the sum of atoms that constitutes the protein. If so, then the claim that causation is layer specific is ad hoc and implausible. (Jonathan Schaffer, in conversation, suggests two other reasons why skepticism about causation between layers is unwarranted. First, if the universe were created by a Big Bang, then if there were no causation between levels we would lack original causes for objects at other levels, such as mental objects. Second, we would need an explanation for the robust regularities we find between levels.)

It may seem that we have reached our final option—accept constitutive overdetermination and defend it as unproblematic or at least inevitable. Once we understand how sums of lower layer objects constitute objects at higher layers, we grant that widespread symmetric overdetermination exists after all, and explain its being widespread by means of the constitutive tie between sums of objects and what they constitute. After all, when we have the sum, we have the higher order object it constitutes, and both are treated equally under the laws. So why be surprised that we have multiple causes? But accepting this option comes at a heavy cost. To see this, we must investigate the phenomenon in more depth. Constitutive overdetermination is stranger than at first it might appear.

The best way to think of cases of symmetric overdetermination is in terms of multiple causal processes: for each overdetermining cause, there is an uninterrupted process from cause to effect such that each process runs to completion. According to this account, *modulo* irrelevant interaction effects, the intrinsic character of each causal process is the same as if there were only one process: none of the overdetermining processes is changed by the addition or removal of another overdetermining process (Hall and Paul forthcoming, § 3).[16]

There is an alternative way to understand symmetric overdetermination. Instead of each cause individually bringing about the effect, all the objects we have been calling ''causes'' are conjoined to bring about the effect, even while none of these objects taken individually causes the effect. Only the objects taken together can be said to be the cause of the effect. For example,

instead of saying that each of the three golf balls is a cause of the shattering of the window, none of the individual golf balls is a cause of the window's shattering—while the sum of all three golf balls is. As Schaffer (2003b) points out, if none of these objects alone counts as a cause of the effect, it is hard to see how all of them together can count as a cause. Nothing more is added! (For an in-depth discussion of other problems with this sort of approach, see Hall and Paul forthcoming.)

So the more plausible way of thinking about symmetric overdetermination adopts the first approach—think of it in terms of multiple causal processes, each running uninterrupted from overdetermining cause to effect. (To require that a causal process exists from each cause to the effect in cases of overdetermination rules out the possibility of conflating overdetermining causes with epiphenomena.) But although this picture is an improvement over the view that overdetermining causes are conjoined, significant conceptual problems remain.

The trouble with symmetric overdetermination in general, and by extension with constitutive overdetermination, is that it is mysterious how each cause could really be a *full* cause of the effect; it is hard to understand how exactly having multiple complete causes of an effect is supposed to *work*. There is something deeply puzzling about true overdetermination, as opposed to preemption and the like. For in true overdetermination, *multiple competing causal processes run all the way to completion*, where each cause is a *complete* cause of the effect that acts independently of the other overdetermining causes *C* (i.e., setting aside irrelevant interaction effects, the overdeterminer does not cause jointly with any of *C* and is not caused by any of *C*). Somehow, multiple causes are subsumed by the same laws and multiple processes run all the way to completion, each bringing about the effect independently of the other causes bringing about the very same effect.

To bring out the intuitive difficulty we have with this picture, consider the difference between symmetrically overdetermined effects and jointly caused effects. An effect is jointly caused if multiple causes combine to produce it. For example, Billy and Suzy each spray some paint on a wall. Billy's paint is red and Suzy's is blue, and the resulting graffiti has red parts and blue parts (and perhaps some purple parts). We can see how each cause is contributing to bringing about the effect, and it is easy to understand the causal contribution of each cause because if any of the causes of the effect

were missing, the effect would not have occurred just as it did (if Suzy hadn't sprayed her paint, the graffiti would have been red). True symmetric overdetermination does not exhibit this characteristic: each cause does enough to do all the (relevant) causal work that is actually done, so if one or more of the overdetermining causes were removed, the effect would occur just as it did when caused by all the overdeterminers.

The graffiti example illustrates how we have a deep understanding of how joint causation works, since each cause makes a clearly defined (and clearly limited) contribution to the production of the effect. But since this model does not extend to true symmetric overdetermination, how are we to understand this sort of causation? Remember, each overdeterminer generates an uninterrupted causal process running to completion from cause to effect. If each cause brings about a process that generates the effect all by itself, why aren't there multiple instances of the very same (type of) effect, one created by each overdetermining cause? (If we treat causation in terms of properties, we can ask why there aren't multiple instances of the same properties.) We seem to have too much causation to go around.

Note that the situation is not merely one where we have many *possible* causes, or overdetermination of mere sufficiency for each effect. If the nonreductionist's picture correctly describes the actual world and is consistent with our best physical theories, there is a transfer of energy, momentum, or some other conserved quantity from *each* overdetermining cause to the effect, and each of these transfers brings about the entire effect. (Describing the case as a transfer of a conserved quantity is a way of saying that in the actual world there is an uninterrupted causal process from each cause to the effect.) Return to our example of the tennis ball. According to the laws, when I hit the tennis ball, I transfer an amount of momentum, p, to the ball. But also according to the laws, when the sum of particles that constitutes me hits the tennis ball, *it* transfers an amount of momentum p to the ball. This brings a central problem to the fore—the problem of additivity of conserved quantities. For example, under the laws, momentum is additive; when there are two transfers of momentum p the total transfer is $2p$. But the tennis ball only exhibits an increase of magnitude p, not $2p$![17]

Perhaps the nonreductionist will hold that in cases of overdetermination transfers of conserved physical quantities are not additive. But in order to accept this claim, reductionists will want to know *why* and, more importantly, *how* this failure to be additive could be the case. Note that it isn't

enough just to claim that some sort of special relationship between the overdetermining causes prevents additivity. As long as there are multiple independent causal processes overdetermining the effect, there will be multiple transfers of conserved quantities to the effect, so additivity constraints seem to apply. (Perhaps in cases of compositionality the special relationship of proper part to whole eliminates the problem of additivity because strictly speaking it eliminates the overdetermination. The whole causes the effect in virtue of the proper part causing the effect; the whole only initiates a causal process in virtue of its proper part initiating that process.)[18]

A more general problem with overdetermination (more general because it applies to accounts of causation in nomically different worlds in addition to accounts of this-worldly causation) concerns our implicit intuitions about the intrinsicness of the causal relation. Consider a structure *S* of events consisting of an event *E* together with all of its causes back to some earlier time. Intuitively, it seems right to say that the causal characteristics of *S* will be fixed solely by its intrinsic character together with the governing laws. But now add some extrinsic detail *D* to the picture (*D* is external to *S*), such as an additional, overdetermining cause of *E*. By hypothesis, adding *D* does not affect the (relevant) causal character of *S*: everything that was a cause of *E* before the addition of *D* is a cause afterwards, and *E* occurs exactly as it would have occurred. But then we can put the worry about symmetric overdetermination this way: what causal contribution is *D* making to the production of *E*, and how does it make it? Simply replying that *D* is sufficient to cause *E* does not explain how *D* makes its causal contribution to the production of *E*, it merely reiterates that *D is* a cause of *E*. How can it be the case that the causal facts about *E* can be changed without changing the intrinsic causal character of *S*? (The causal facts are changed because *E* is now overdetermined.) The worry is that symmetric overdetermination seems to violate the natural intuition that the complete causal character of *S* is fixed solely by its intrinsic character plus the laws.[19]

These considerations help to bring out the fact that, no matter how we try to gloss the way it is supposed to work, true symmetric causal overdetermination is *strange*. It is a real mystery how symmetric overdetermination can be understood on a deep level, and how we are to fit this sort of causation into our more general picture of how causes bring about effects. Since it is reasonable to assume that every macroscopic object is constituted by

sums of particles, we can now see that the problem for the nonreductionist amounts to a problem with explaining just about every instance of macroscopic causation in the world. If constitutive overdetermination is supposed to occur almost everywhere in the actual world (since almost everything actual is constituted by or constitutes something else), then we lack a decent understanding of the way causation works in almost every actual-world case. To simply bite the bullet and accept constitutive overdetermination is to accept that actual causation is fundamentally mysterious in a hitherto unrecognized way.

This problem extends past the worry that we have made ordinary causation mysterious, since it also creates trouble for analyses of causation. Extant analyses of causation do not adequately explain symmetric overdetermination even if they give the right answer, namely, that each overdeterminer counts as a cause. If overdetermination cannot be set aside as peripheral, the task of developing an analysis of ordinary causation is made significantly harder. Since an acceptable analysis of ordinary causation is needed for *fully* informative treatments of many central topics in philosophy—free will, action, decision theory, reference, perception, and laws of nature, for example—worsening the prospects for an acceptable analysis of ordinary causation means worsening the prospects for complete versions of these theories as well.

Note that I am happy to grant that symmetric overdetermination is metaphysically possible, and perhaps even actual. The claim is not that symmetric overdetermination is metaphysically incoherent, nor that we have no reason to believe in overdetermining entities. (See Sider 2003 for a critical discussion of such claims.) Rather, the point is that constitutive overdetermination (and perhaps other sorts of true symmetric overdetermination) should be thought to be mysterious and problematic, and if a philosophical theory implies this sort of overdetermination, it implies widespread mystery. Accepting that a few localized instances of causation are mysterious is unpleasant. Accepting that almost all instances of ordinary causation are mysterious is unacceptable.

If there is no better option, then perhaps accepting that the world has constitutive overdetermination is the best option nonreductionists have. I used to think it *was* the best option—I thought that constitutive overdetermination, while extremely problematic, was simply a feature of the world that nonreductionists had to accept. After all, at least it didn't make objects

epiphenomenal or causally inefficacious (e.g., by holding that higher or lower order objects are continuously preempted by lower order objects).[20] But it *is not* the best option.

To see this, we need to trace our chain of reasoning back to an early assumption we made about the nature of the world. How did the problem arise in the first place? It arose when nonreductionists combined the thesis that objects are not identical with their constituting sums, with a model of the world as layered. If material objects are numerically distinct from the sums of material objects that constitute them, without additional guidance to their ontology this suggests that they are somehow stacked on top of their constituting sums within the same space-like extra-dense layers of cake like some sort of weird Escher building. While the view is not quite substance pluralism (since the nonreductionists' layers are not ontologically independent), it is awfully close. Rejecting the identity claim made by the reductionist without explaining how to reinterpret the model of the layered world leads to this situation.

The idea that nonreductionists must endorse a ''stacking'' or ''dense'' interpretation of the layered model is the source of objections to nonreductionism that complain that if a statue is not identical to the sum of its bronze particles, ''why don't the two together weigh twice as much?'' (Lewis 1986a, p. 252). Since contemporary nonreductionists should deny that higher layer objects share the spatiotemporal parts of their constituting sums of lower layer objects because of differences in their properties—for example, differences between properties of the larger spatiotemporal parts of a person and properties of the sums of particles that constitute those parts—the need to make sense of how the world is layered becomes even more pressing. We need additional guidance as to the nature of the constitution relation if we are to avoid deep ontological quagmires.

I suspect some version of this mysterious picture of extra-dense layers has been subtly attributed to nonreductionists by many philosophers, especially by reductionists. And the trouble with this (aside from the conceptual difficulties surrounding such a picture), as we have seen, is that if multiple distinct objects related by constitution are each causally sufficient for the effect in question, then it seems that we must admit massive amounts of constitutive overdetermination.

Nonreductionists need to produce a clear account of how the layered world model works according to their views—an account that does not

imply widespread overdetermination and does not endorse epiphenomenalism, substance dualism, or some sort of emergentism. Below, I sketch an account of the ontology of constitution that provides a working understanding of the relation of constitution and thus a working interpretation of the layered world model. By providing an outline of how objects can overlap when they are related by constitution, I outline an interpretation of the model of the world that points nonreductionists away from constitutive overdetermination. (For a more developed account of this ontological approach and of overlap between coincident objects, see Paul 2002, 2006, and my manuscript "The Ontology of Objects.")

The first step is to characterize the fundamental ontology of the objects of the world. Instead of thinking of objects in primarily spatiotemporal mereological terms, think of them in terms of their properties. Objects are sums of properties, so they have properties as parts.[21] According to this account of the objects of the world, objects like persons are fusions that include property instances such as *having mass m* and *having shape s*, and so forth, as are the sums that constitute them.

When sum *A* constitutes higher layer object *B*, *A* and *B* overlap with respect to many of their property instances: they literally share such instances as overlapped parts, especially those that we can dub *material* instances like *having mass m* or *having shape s*. If we do not assume that objects are individuated by their matter or region, we can see that we can have numerically distinct objects that differ with respect to some of the property instances (such as *de re* modal property instances) they include even if they share their material and place. The idea is that objects related by constitution can share material property instances such as *having mass m* or *including particle p*, even if they do not share *all* of their property instances.[22]

For example, the person and the sum of cells that constitutes him overlap with respect to many of their material instances, but they do not overlap with respect to all their property instances, even if they share their matter and spatiotemporal region. The person includes the property instance of *being handsome* while the sum of his cells does not, and the sum of his cells includes the property instance of *having n cells essentially* while he does not. Likewise, a protein and the sum of microparticles that constitutes it share some, but not all, of their property instances. The protein includes the property instance of *having mass m*, and the sum of particles

also includes *having mass m*. But the protein includes the property instance of *having particle p accidentally*, while the sum of microparticles includes an instance of *having particle p essentially*. These objects differ in the property instances they include, and hence they are absolutely, or numerically, different even while they spatiotemporally and materially overlap. When objects only partially qualitatively overlap, the objects share only some (qualitative) property instances and in this sense are only "partly" identical.

The nonreductionist now has the basic tools to construct a (partial) working interpretation of the model of the layered world. The layers are organized into a hierarchy governed by the relations of constitution and composition such that sums of objects at lower layers constitute objects at higher layers—sums of particles constitute atoms, sums of atoms constitute molecules, and so on.[23] When objects constitute other objects, they are partly identical in that they partly overlap by sharing many of their property instances. This last point is of essential importance: the objects overlap with respect to some of their property *instances*; if the person and his constituting sum of cells overlap with respect to property instance *P*, then they *share* the very same property instance. Since the property instances had by constituting and constituted objects are literally shared, the nonreductionists' interpretation shows how the world is not extra-dense with respect to shared property instances in a constitutional hierarchy—any more than a shared office wall is twice as dense because it is part of two different offices.

Now that we have a nonreductionist interpretation of objects related by constitution, we can assess the overall picture for the possibility of constitutive overdetermination. First, we need to say a bit more about the way causation is supposed to work in the world. There are two features of any successful reductive account of causation that will be important for our solution. For *c* to cause *e*, *c* must lawfully entail *e*, given the background conditions and the laws. In other words, some sort of nomological sufficiency is required for causation, at least for the cases I am concentrating on. Also, causation involves property instances—either as constituents of events and objects, or as causes and effects in their own right.

In Paul 2000, I argue that it is best to take property instances as the causal relata. But even if you insist that events or objects must be the causal relata, property instances must play a central role. Events or objects are causal actors only in virtue of their property instances, for property instances (or exemplifications, etc.)[24] are what are subsumed by laws. Ideally, for nomic

subsumption of cause and effect, for event *c* with property *p* and event *e* with property *q*, *c* lawfully entails *e* iff *c*'s exemplification of *p* is subsumed by the antecedent of the right law or laws that entail a consequent subsuming *e*'s exemplification of *q*.

Why does this matter? I want to draw attention to the role of property instances in causation because I think it is essential to see that when we say an object or event is a cause of an effect, we usually mean that the object or event is a cause in virtue of some of its distinguishing properties being involved in the requisite nomic subsumption. Not all of its properties must be involved, but enough of the *important* ones need to be involved in order to say that *c* caused *e*. This means either that property instances are causes and effects, or that objects and events are causes and effects in virtue of exemplifying certain properties.

Either we need to take property instances as causes and effects, or we need to be clear about what properties of an event or object are causing what. Imagine a sphere that is spinning while being heated. The spinning is a cause of the sphere's motion. The heating is a cause of the sphere's high temperature. As it turns out, the hot, spinning sphere touches a flag, which flutters and then bursts into flame. The hot, spinning sphere is a cause of the fire, but the breeze it created also caused the flag to flutter. If we say that the sphere causes the fire and causes the fluttering, we must be careful to say that the sphere causes the fire in virtue of its temperature properties, while it causes the fluttering in virtue of its motion properties.

The role of property instances in causation is relevant to whether constitutive overdetermination is implied by the nonreductionist view. Under the model of the world that the nonreductionist proposes, the way the world is arranged involves a number of numerically different but literally overlapping objects (and events that include these objects as constituents) that share their material, their spatiotemporal region, and in general many of their nonmodal property instances. In the relevant cases, the objects *share* many, if not all of, their nonmodal properties, and they do this by sharing *instances* of these properties. In particular, they share many of their material nonmodal property instances; instances of properties of having particular colors, weights, masses, shapes, temperatures, and so forth. Such nonmodal property instances typically figure in laws of nature and are involved in causing many of the effects we observe. For example, when I weigh my cat on the bathroom scale, the scale reads 8 pounds. My cat

exemplifies (includes) the property of *weighing 8 pounds*, and this property is part of a lawfully sufficient condition for the scale's reading 8 pounds. So, my cat's standing on the scale is a cause of the scale's reading 8 pounds, in virtue of her having the property instance of *weighing 8 pounds*.[25]

Now, of course, the sum of cells that constitutes my cat also exemplifies (includes) the property of *weighing 8 pounds*, as does the sum of molecules that constitutes the sum of cells. But do we have constitutive overdetermination of the reading on the scale? No, because my cat and her sums *share* the causally efficacious property instance of *weighing 8 pounds*. My cat and her sums are numerically distinct objects that are not entirely distinct—and a common part of the objects, the property instance of *weighing 8 pounds*, is what does the causing. Only *one* causal process is initiated here, since only one property instance is subsumed by the (relevant) laws. Since a part of the shared portion of these entities is what is causally relevant we can say that the cat and the sum of cells, the sum of molecules, and so on, each caused the reading on the scale but there is no overdetermination.

Likewise for the protein and its sum of particles, and other objects in the hierarchy of constitution. It was only when we thought of the different objects as somehow stacked or piled up in some way with duplicated property instances or parts that we seemed to get constitutive overdetermination. Of course, not all of the property instances of the objects or events we are considering will be causally relevant: modal property instances, and perhaps some other property instances may not be causing anything in some cases. But we do not need *every* property instance of an event or an object to be causally efficacious in order to say that the event or object was a cause—just the important ones. In the cases of putative constitutive overdetermination, we need the property instances that figure in the relevant laws of nature to be the causes, not *de re* property instances of being essentially such and such or relational property instances such as *being an art object*. The important property instances are the ones that we thought all along were doing the causing—we thought all along that it was my cat's property of *weighing 8 pounds* that was responsible for the reading on the scale.

If we thought that it was the cat's *being accidentally 8 pounds* that caused the reading on the scale, then we wouldn't have worried that the reading was overdetermined in the first place, since the property instance of *being accidentally 8 pounds is not* shared by the cat and the sum of her cells. But

we don't think the cat's being *accidentally* 8 pounds, as opposed to her being 8 pounds *simpliciter*, is a cause of the reading. Whether she is *accidentally* 8 pounds or whether she is *essentially* 8 pounds has nothing to do with my causal judgment—only whether she *is* 8 pounds.[26] Likewise for my cat's property of *being gorgeous*: my cat is gorgeous but the sums that constitute her are not, and so anything caused by her gorgeousness (the *oohs* and *aahs* of my dinner party guests upon beholding her, for example) is not overdetermined.

The same goes for the causal action of a protein: let us say that a few of the empirical properties of a protein cause a scientist to get an interesting result using NMR or IR spectroscopy. Of course, the sum of fundamental particles that composes the protein also includes these properties. It is the shared property instances of the protein and its sum—for example, the property of *including particle p*—that cause the spectroscopic readings. What *does not* cause the readings is the protein's property of *accidentally including particle p*. Nor does the sum's property of *essentially including particle p* affect the readings: these properties simply are not causally relevant to the effect in question, and they are not shared by the objects. (Some might even claim that *de re* modal properties are never causally active. Then all the shared properties of the objects could be the causally relevant properties, and we would not have to worry about excluded properties in our causal story. Questions arise about how we know about such properties if they are causally inactive—the same questions we might have about how we know about numbers, abstract objects, and so on.)

We need to consider one final worry. Have we simply dodged the overdetermination bullet by talking about objects when the real problem lies with the properties of objects? Can the reductionist move the debate from relations between objects to relations between properties, and argue, for example, that the properties of *being a sum of particles* and *being a protein* overdetermine the spectroscopic readings? Presumably, the justification for such a move would involve the claim that the nonreductionist must hold that property instances at different layers are just as irreducible as objects at different layers, and since property instances are the real causes and effects, we need to pay special attention to them when investigating possible overdetermination.

Fortunately for the nonreductionist, shifting the terms of the debate to properties merely shifts the terms of the solution to properties. First,

recognize that there are many more predicates than properties and property instances, but allow that there may be a property instance of *being protein P* if this property instance is really just a conjunction or fusion of ontologically more fundamental property instances. Consider protein *Pro*, constituted by sum of molecules *Mol*. The nonreductionist should hold that the property instance of *being protein Pro* is a complex property instance that is really just a conjunction of many more fundamental property instances such as *having shape s*, *having mass m*, *having n molecules accidentally*, and so on. Now consider the property instance of *being sum Mol*: it is a complex property instance that is just a conjunction of many more fundamental property instances such as *having shape s*, *having mass m*, *having n molecules essentially*, and so on. Just as *Pro* and *Mol* partly overlap or share their property instances, the property instance of *being Pro* and the property instance of *being Mol* share some of their conjuncts. (Further, as a conjunctive property instance such as *being Pro* is just the fusion of the conjuncts, likewise the property instance of *being Mol*.) And just as with the solution in terms of objects—it is the shared (instances of) conjuncts that are responsible for the problematic cases of putative constitutive overdetermination.

This should be no surprise when we consider our ontology of objects: to say that an object has the property instance of *being protein Pro* is to say that there is an object that is protein *Pro*, and to say that an object has the property instance of *being sum Mol* is to say there is an object that is the sum *Mol*.[27] Since *Mol* constitutes *Pro*, *Mol* and *Pro* partly overlap and thus both can count as non-overdetermining causes of effects. For the same reasons and in the same way, the property instance of *being Mol* constitutes the property instance of *being Pro*, and so both property instances can count as non-overdetermining causes of effects. (The move parallels the debate in philosophy of mind over mental causation, and the neo-Kripkean nonreductionist about the mental can defend a similar solution there: mental property instances include conjuncts such as *being accidentally embodied*, and physical property instances include conjuncts such as *being essentially embodied*, but they overlap with respect to many of their material conjuncts.) The often-raised objection about a mental entity causing an effect ''in virtue of being mental'' fails here for the same reason that a parallel objection about my hitting the tennis ball ''in virtue of being me'' fails. In each case, a nontrivial proper part is the cause (some of my material

instances, some of the mental entity's material instances) and this is sufficient for me and for the mental entity to be causes.

Whether we focus on objects or properties, once we realize that it is the involvement of certain property instances that determine whether one thing causes another, and that in the cases of constitutive overdetermination we have considered, the causally important or relevant property instances are shared, we can see why causal responsibility is shared, not overdetermined. When I, my constitutive sum of cells, constitutive sum of molecules, and so on, hit the tennis ball, a *shared* part (a shared property instance) causes the tennis ball to bounce off my racquet at 100 miles per hour. The problem of constitutive overdetermination is merely an artifact of a flawed interpretation of the layered model of the world.

Acknowledgments

I am particularly grateful to Jonathan Schaffer for comments and discussion. I would also like to thank audiences at the 2003 Inland Northwest Philosophy Conference on Explanation and Causation, and the 2003 Annual Conference of the Australasian Association of Philosophy.

Notes

1. The original Oppenheim-Putnam program was characterized in terms of reductions involving nomic equivalence and the construction of bridge laws between terms of theories of objects in different layers, and does not explicitly address issues concerning constitution. For the purposes of this argument I characterize reductionism more directly, in terms of an identity claim about objects. Debates about reduction of objects intersect with debates about the reduction of properties, especially in philosophy of mind and philosophy of science. Whether ''layer,'' ''level,'' or ''order'' is used to discuss the Oppenheim-Putnam model varies from author to author. I have chosen ''layer,'' but I don't think much turns on the nomenclature.

2. Proteins are constructed from amino acids joined by peptide bonds, and most amino acids include carboxyl groups (OH–C = O, or –COOH).

3. Most reductionists will also argue for redescriptions of the sum and the higher order object such that there is no difference in properties after all (e.g., Lewis 1986a). I don't hold out much hope for these strategies; see Fine 2003 and Paul 2006 for discussion.

4. I defend this point in detail in Paul 2006.

5. Some contemporary reductionists will hold that mereological composition is identity, so that there really is just one ontological layer; whereas other reductionists will hold that mereological composition is merely analogous to identity, but such that wholes are ''nothing over and above'' their parts. It is unclear to me how the latter sort of reductionist will represent the layered world.

6. Kim (2002, p. 7) suggests that work by emergentists like C. Lloyd Morgan is the origin of the layered approach.

7. Defenders of nonreductionism who try to address these issues include Lowe (1998), Wiggins (2001), and Paul (2006). Rea (1997) gives an excellent survey of the early debate.

8. Kim (1984, 1998), Baker (1995), and Merricks (2001) discuss the problem of symmetric overdetermination for other nonreductionist views. Baker and Kim are important representative examples of those who focus on issues involving the reduction of the mental to the physical. Merricks is concerned about overdetermination in cases where parts compose wholes.

9. I am assuming that epiphenomenalism is not an option.

10. There are related worries about overdetermination due to the nonidentity of objects and events. Sider (2003) rightly notes that this sort of overdetermination deserves more attention. The solution I propose at the end of this paper can also be used to resolve worries in cases where objects partly constitute events; I address the issue in more detail in my manuscript, ''The Ontology of Objects.''

11. There are exceptions; see Schaffer 2003b and Hall and Paul forthcoming.

12. See, for example, Lewis 1986b, p. 194.

13. Merricks (2001, 2003) argues that this sort of overdetermination is troubling; Sider (2003) argues that it is not. A discussion of this issue will take us too far afield; in any case, this sort of overdetermination is faced by reductionists and nonreductionists alike.

14. Van Inwagen (1990) and Merricks (2001) are recent defenders of a partly eliminativist position. Cian Dorr (2002) defends a thoroughgoing eliminativism. A different kind of skeptic might argue for eliminativism from the bottom up, claiming that, for example, rocks exist but the particles that supposedly compose them do not.

15. Block (1990, 2003), Schaffer (2003a), Sider (1993), and Sosa (1999) all point out this problem for the eliminativist.

16. I am helping myself to an intuitive notion of a causal process here.

17. There is a parallel problem with effects, since each overdetermining cause overdetermines many effects (all those related in the constitutional hierarchy). My hitting the tennis ball, my constitutive sum of cells hitting the tennis ball, and my

constitutive sum of particles hitting the tennis ball, etc., each overdetermines the bouncing of the tennis ball at 100 mph. But each object overdetermining the bouncing of the tennis ball also overdetermines the bouncing of the tennis ball's constitutive sum of molecules at 100 mph and the bouncing of the tennis ball's constitutive sum of particles at 100 mph, etc. Under the laws, my hitting the tennis ball only imparts momentum *p*, but somehow, many (otherwise causally unrelated) entities receive an increase of *p* from the hit—the tennis ball, its constitutive sum of molecules, its constitutive sum of particles, etc. For ease of exposition I have ignored this complexity. (Note: don't think that the solution to the problem of additivity is solved by the complexity of the overdetermination. If anything the problem is worsened, since if there is causal overdetermination, each object—myself, and each of my constitutive sums—somehow imparts momentum *p* to each of many objects—the tennis ball and each of its constitutive sums. If there were a one-to-one mapping of transfer of momentum—for example, if there were no causation across layers—there would be no problem of additivity, but only because there would be no overdetermination. I've already shown above that rejecting causation across layers is unacceptable.)

18. If there is true symmetric overdetermination in cases of compositionality, such that the thing that is a proper part initiates a causal process sufficient for the effect, and the whole initiates another causal process sufficient for the effect, the problem of additivity might return. To explain why there is still no additivity problem, perhaps one could argue that the non-additivity of conserved quantities of parts and wholes is implicit in the laws or at least the practice of science. In other words, the claim is that in the special case where parts compose wholes, science implies that the conserved quantities are not additive after all. If such a view is to be made plausible, an explicit account of how science can be thought to imply non-additivity needs to be provided. (I am indebted to Timothy Williamson and Terry Horgan for discussion of this point.)

19. This perspective was first defended in Hall 2004, and is further developed in Hall and Paul forthcoming, section 3.

20. If there can be infinitely descending preempted chains this creates interesting problems for a counterfactual analysis of causation; see Lewis's postscript to ''Causation'' in Lewis 1986b, for a discussion.

21. I introduce and develop this approach in Paul 2002 and Paul 2006. Since spatiotemporal parts are not shared but spatiotemporal regions are, spatiotemporal parts are individuated by more than the regions they occupy. This follows from ordinary nonreductionist reasoning about property differences between the parts of spatiotemporally coinciding objects.

22. Defenders of constitution need to explain why *de re* modal and other property differences do not supervene on shared matter and location; see Paul 2006.

23. Reductionists who hold that composition is not strictly identity will agree with the nonreductionist that composition and constitution are different relations. The complete interpretation of the layered world model for this sort of reductionist and for the nonreductionist will be complex. For starters, both will need to say more about how composition and constitution together define the overall hierarchy of objects. On the nonreductionist view I am outlining here, particles compose sums of particles, but only partly compose atoms: particles compose sums of particles, sums of particles constitute atoms; atoms compose sums of atoms, sums of atoms constitute molecules; and so on. See my manuscript ''The Ontology of Objects'' for more discussion.

24. Property instances are located properties. Rather than assume a particular account of properties, I will speak interchangeably of properties had by events or objects, properties exemplified by events or objects, and property instances. What matters for my purposes here is that the property is located or exemplified at a location, not whether the property is a trope or an instance of a universal. If a property is had by an object or event then the property is located, because objects and events are located.

25. A different sort of worry involves the possibility of not having enough joint causation. What if two materially distinct particles (particles that do not share their matter and are not related by constitution or composition), each including the property of *having charge c*, jointly cause effect *e*? How can this be possible if the particles share the causally relevant property? In response we must remember that it is property *instances*—not properties—that are causes. If there are two materially distinct particles, they include numerically distinct property *instances* of *having charge c* even if they both have the property of *having charge c* (suppose the instances can be individuated even if the particles have the same location). Each instance initiates a causal process, so *e* is jointly caused. I discuss the different ways to individuate objects with the same properties in ''The Ontology of Objects.'' (I am indebted to Adam Elga for raising the worry about not having enough causation.)

26. So for these reasons we do not fall victim to the epiphenomenalism objection to Davidson's treatment of mental causation—the important causal properties of our higher layer objects are not epiphenomenal.

27. The claim here is that certain sorts of claims about properties function as veiled identity claims. For example, to say that I have the property of *being Laurie* is just to say that I *am* Laurie. If I am identical to a fusion of properties *LP*, then the property of *being Laurie* is just that fusion of properties *LP*.

References

Baker, L. 1995. ''Metaphysics and Mental Causation.'' In *Mental Causation*, edited by J. Heil and A. Mele. New York: Oxford University Press.

Block, N. 1990. ''Can the Mind Change the World?'' In *Meaning and Method*, edited by G. Boolos. Cambridge: Cambridge University Press.

Block, N. 1997. ''Anti-Reductionism Slaps Back.'' *Mind, Causation, World: Philosophical Perspectives* 11: 107–133.

Block, N. 2003. ''Do Causal Powers Drain Away?'' *Philosophy and Phenomenological Research* 67: 110–127.

Boyd, R., P. Gasper, and J. Trout, eds. 1991. *The Philosophy of Science*. Cambridge, Mass.: MIT Press.

Dorr, C. 2002. ''The Simplicity of Everything.'' Ph.D. dissertation, Princeton University.

Fine, K. 2003. ''The Non-Identity of a Material Thing and Its Matter.'' *Mind* 112: 195–234.

Hall, N. 2004. ''Two Concepts of Causation.'' In *Causation and Counterfactuals*, edited by J. Collins, N. Hall, and L. Paul. Cambridge, Mass.: MIT Press.

Hall, N., and L. Paul. Forthcoming. *Causation and Its Counterexamples: A User's Guide*. Oxford: Oxford University Press.

Kim, J. 1984. ''Epiphenomenal and Supervenient Causation.'' In *Midwest Studies in Philosophy*, vol. 9, edited by P. French, T. Uehling, and H. Wettstein. Minneapolis: University of Minnesota Press.

Kim, J. 1993. ''The Nonreductivist's Troubles with Mental Causation.'' In J. Kim, *Supervenience and Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press.

Kim, J. 1998. *Mind in a Physical World*. Cambridge, Mass.: MIT Press.

Kim, J. 2002. ''The Layered Model: Metaphysical Considerations.'' *Philosophical Explorations* 5: 2–20.

Lewis, D. 1986a. *On the Plurality of Worlds*. Oxford: Blackwell Publishers.

Lewis, D. 1986b. *Philosophical Papers II*. New York: Oxford University Press.

Lowe, E. 1998. *The Possibility of Metaphysics: Substance, Identity, and Time*. Oxford: Clarendon Press.

Merricks, T. 2001. *Objects and Persons*. Oxford: Oxford University Press.

Merricks, T. 2003. ''Replies.'' *Philosophy and Phenomenological Research* 67: 727–744.

Oppenheim, P., and H. Putnam. 1958. ''Unity of Science as a Working Hypothesis.'' In *Minnesota Studies in the Philosophy of Science: Volume II*, edited by H. Feigl, M. Scriven, and G. Maxwell. Minneapolis: University of Minnesota Press. Reprinted in Boyd 1991.

Paul, L. 2000. ''Aspect Causation.'' *Journal of Philosophy* 97: 235–256.

Paul, L. 2002. ''Logical Parts.'' *Noûs* 36: 578–596.

Paul, L. 2006. ''Coincidence as Overlap.'' *Noûs* 40: 623–659.

Paul, L. ''The Ontology of Objects.'' Unpublished manuscript.

Rea, M. 1997. *Material Constitution: A Reader*. Lanham, Md.: Rowman and Littlefield.

Schaffer, J. 2003a. ''Is There a Fundamental Level?'' *Noûs* 37: 498–517.

Schaffer, J. 2003b. ''Overdetermining Causes.'' *Philosophical Studies* 114: 23–45.

Sosa, E. 1999. ''Existential Relativity.'' In *Midwest Studies in Philosophy*, vol. 23, edited by P. French and H. Wettstein. Boston: Blackwell Publishers.

Sider, T. 1993. ''Van Inwagen and the Possibility of Gunk.'' *Analysis* 53: 285–289.

Sider, T. 2003. ''What's So Bad about Overdetermination?'' *Philosophy and Phenomenological Research* 67: 719–726.

van Inwagen, P. 1990. *Material Beings*. Ithaca: Cornell University Press.

Wiggins, D. 2001. *Sameness and Substance Renewed*. Cambridge: Cambridge University Press.