

This is a preprint of an article whose final and definitive form will be/is published in the *Australasian Journal of Philosophy*, 2010. The *Australasian Journal of Philosophy* is available online at: <http://www.tandf.co.uk/journals/>.

THE UNREALITY OF REALIZATION¹

Chase B. Wrenn

University of Alabama

ABSTRACT: This paper argues against the *realization principle*, which reifies the realization relation between lower-level and higher-level properties. It begins with a review of some principles of naturalistic metaphysics. Then it criticizes some likely reasons for embracing the realization principle, and finally it argues against the principle directly. The most likely reasons for embracing the principle depend on the dubious assumption that special science theories cannot be true unless special science predicates designate properties. The principle itself turns out to be false because the realization relation fails the naturalistic test for reality: it makes no causal difference to the world.

1. The Realization Principle

The *realization principle* says:

There is a relation of realization, which obtains between “lower-level” and “higher-level” properties.

The principle stands behind a great deal of current philosophical work, much of it aimed at clarifying the nature of the realization relation itself [Hendel 2001; Gillett 2002; Polger 2004].

There are debates about whether there is a one-to-many mapping of higher-level properties

to the lower-level properties that realize them [Bechtel and Mundale 1999; Putnam 1967; Fodor 1974; Fodor 1997]. There are debates about which lower-level properties realize which higher-level properties.² And there are even debates about whether everything that realizes a given higher-level property is a *realizer* of that property [Rupert 2007].

At their best, philosophical discussions of realization proceed naturalistically, aiming to answer questions about the realization relation by drawing on, interpreting, and developing our best scientific views of the natural world. This is unsurprising, given what I take to be the main motivation for positing a realization relation at all: The fundamental physical facts of the world determine all the facts there are, and so the claims of the special sciences face an eliminativist threat. Why suppose our best psychological theories are literally true, when everything that happens is ultimately determined by fundamental physics? Why bother with psychology, when (in principle, anyway) physics will do? The realization principle allows physicalists to evade the challenge by conferring a measure of ontological respectability on the subject matter of the special sciences. Those sciences address phenomena *realized by* fundamental physical phenomena, but the phenomena are no less real, and the theories of them are no less true, for their ontological dependency.

I think the realization principle is false, but I do not have my sights on an all-out refutation of it in this paper. Instead, I aim to show that the realization principle is inconsistent with certain principles of naturalistic metaphysics I find compelling. A refutation of the realization principle would require a more extensive defence of those principles than I undertake here. I take the principles for granted, and I argue that those who embrace them should reject the realization principle. (Conversely, those who cleave to the realization principle anyway should reject one or more of those otherwise compelling metaphysical principles.)

The realization principle is false because realization is not a genuine relation. Rejecting the realization principle means refusing to reify realization. It does not require embracing eliminativism or reductionism about higher-level properties, or antirealism about the special sciences; nor does it require us to suppose higher-level properties are radically autonomous of lower-level properties. I am officially neutral about existence of higher-level properties and their reducibility to lower-level properties. My claim is just that there is no such relation as “realization” to connect higher- and lower-level properties. Or, to put it more carefully: the reification of such a relation is inconsistent with principles of naturalistic metaphysics I find compelling.

The following section sets out those principles, giving them a cursory and necessarily incomplete defence. Section 3 addresses some likely arguments in favour of positing a realization relation. All are unpersuasive, though, for they turn on the dubious assumption that special science theories cannot be true unless their predicates designate properties. My case against the realization principle occupies Sections 4 and 5. In Section 4, I argue that realization is unreal because it makes no causal difference in the world. In Section 5, I reply to several likely objections. Section 6 briefly addresses the question of how to understand the relations among higher- and lower-level sciences without positing a realization relation.

2. Naturalistic Metaphysics

There are two fundamental metaphysical questions. One is the question of *ontology*: What is there? The other is the question of *ideology*: What *sorts* of things are there, and what are they like? I take my cue from W. V. Quine [1951; 1983] in answering each question. (But I do not claim that Quine would agree with everything I say here.)

Quine was aware that, once we distinguish ontology from ideology, the ontological question is easy. What is there? Everything. The things I believe to exist are what the

quantifiers of my total theory of the world range over. As a naturalist, my total theory of the world is the best available total scientific view of things. The scientific everything—that is, what the best available scientific view of things means by “everything”—is everything simpliciter. Nothing outside the scope of science is real.

This is not yet to say anything about whether our scientific everything includes properties, relations, propositions, numbers, beliefs, quarks, photons, neutrinos, depressions, luminiferous ether, or dark matter. Such questions are almost always matters of ideology, not ontology. They come down to whether our best scientific view of things says some things are quarks, neutrinos, depressions, properties, ether, etc.³ We typically answer ideological questions by doing empirical science, not by purely *a priori* speculation.

The question of the reality of properties and relations can be posed either ontologically or ideologically. Ontologically, we can ask whether our best theory of the world quantifies into predicate position. If it does, it is a theory with quantifiers ranging over not only individuals, but attributes and relations as well. With Quine [1970], I am dubious that our best theory of the world is “second order” in this sense, because I think we are already ideologically committed to enough set theory to replicate the good parts of second order logic without quantifying over anything but individuals and sets of them.

Posed ideologically (and assuming a first-order theory of the world), the question of the reality of properties and relations asks whether anything in our scientific everything is a property or relation. How could we tell?

Some naturalists seem to think that we are justified in positing properties to go along with any predicates that happen to figure in a theory we accept. The idea is that we take the theory to be true, and it needs properties as the semantic values of its predicates in order to

be true. Accepting the theory then justifies us in positing properties for its predicates to designate.

This view is untenable for at least two reasons. First, positing properties for every well-formed predicate of our theory of the world leads to contradiction, as Russell's Paradox demonstrates. Second, it also will not work to posit properties for only the primitive predicates of our theory. Given a theory with one set of primitives, we can construct an equivalent theory with a different set of primitives. Primitivity is no less a product of how we choose to formulate our theories than of how the world is. The primitivity of a predicate is no proof of a property. If we are ontologically serious, and thus we mean more when we say a property exists than just that a predicate's extension is nonempty, we need a way to distinguish predicates that designate genuine properties from predicates that do not. We need a criterion of reality for properties.

The best naturalistic criterion I know is a form of the Eleatic Principle: Whatever is real has causal powers [Armstrong 1999]. So stated, of course, the principle is far too strong. It rules out the reality of all manner of causally impotent individuals, most notably mathematical abstracta. All the same, we can restrict the principle to concern only properties and relations. Any genuine property or relation must make a causal difference in the world, by making a difference to the causal powers had or manifested by its bearers.

The restricted Eleatic Principle provides only a necessary condition, but it does give us a principled way to *deny* the reality of a supposed property or relation. If a supposed property bestows no causal powers on its bearers, or if the instantiation of a supposed property makes no causal difference in the world, then that supposed property is unreal. The same goes for relations: If the satisfaction of a polyadic predicate makes no causal difference to the world, the predicate does not identify a genuine relation. By a "causal difference," I

mean a difference in the causal powers had or manifested by particulars. So the point can also be put in this way: If the satisfaction of a predicate makes no difference to what things *do* or *can do*, the predicate does not designate a real property or relation.⁴

This version of the Eleatic Principle provides a purely negative test, but it is a test that can be passed. For example, *mass* properties make a difference to the causal powers of objects that have them, and *distance* relations make a difference to the causal powers of their relata. It is the business of a theory of gravitation to specify exactly what differences these properties and relations make. As I understand our best physics, the rate at which two objects accelerate toward one another is a function of precisely which mass properties they have (i.e., of their masses) and of which distance relation they bear to one another (i.e., of how far apart they are).⁵ Mass properties and distance relations pass the Eleatic test.

But why should we accept the Eleatic criterion? Is it not too strong, incorrectly ruling against the reality of such clearly genuine properties and relations as *being a prime number* and *fatherhood*?

If there are genuine mathematical and logical properties and relations, then the Eleatic criterion might be too strong as stated. It is notable, though, that mathematics gets along just fine *without* distinguishing between properties and relations, on the one hand, and their extensions on the other. Our ideological commitment to the sets that form the extensions of predicates such as 'prime number' and 'less than' might be all we need, and in that case mathematical and logical properties and relations would not force us to relax the Eleatic criterion. On the other hand, we could always restrict the scope of the Eleatic criterion to specifically exclude mathematics and logic: Genuine non-mathematical and non-logical properties have (or bestow) causal powers.

Relations such as fatherhood fail even a version of the Eleatic criterion liberalized to make room for mathematical and logical properties. In my view, this is no great loss. To deny that there is any such relation as fatherhood is not to deny that anyone is or has a father. It is just to say that, between any father and his progeny, there is no third entity, *fatherhood* also on the scene. Fatherhood does not have causal powers. Nor does it make a difference to the causal powers had or manifested by its relata. Once we see that the posit of a fatherhood relation, as a further entity, is unnecessary for it to be true that one person is another's father, and we see that fatherhood would make no causal difference in the world, the reification of fatherhood appears pointless. Likewise, the refusal to reify it looks harmless.

Another worry, similar to the worry about *being a prime number* and *fatherhood*, is that causation itself fails the Eleatic test. And if there is no such thing as causation, then why should we have ever been interested in the Eleatic Principle at all?

This worry depends on the mistaken thought that not reifying causation means saying nothing ever causes anything. But denying that causation is part of our scientific everything means only that, between a given cause and one of its effects, there is no third *thing*, causation, also on the scene. There is no tension between the Eleatic test and causation's failure to pass it, and the test is useful even without reifying causation. It is useful in counselling against positing things that make no difference to what things do or to what they have the power to do.

That is my brief defence of the Eleatic criterion against the charge that it is too strong. The motivation for the principle itself comes from two sources, one Quinean in spirit and one decidedly *un*-Quinean.

The un-Quinean motivation is well expressed by Charles Martin [1997]. We posit properties to explain the fact that things interact with one another not holus-bolus but in virtue of the ways in which they are. When a red brick smashes through a window, we say it is because of the brick's hardness and the window's brittleness, not the brick's redness and the window's transparency. We look beyond the brick and the window to some further entities, whose interaction with one another was responsible for the event. Those further entities are properties.

We posit properties, in addition to objects, to explain the causal interactions of objects. Properties are suitable to fulfil that explanatory function only if they make a difference to the causal powers of their bearers, either by making a difference to the powers their bearers possess (what they *can* do) or to the powers they manifest (what they *do* do). But if that is our reason for positing any properties at all, it is pointless to posit properties or relations that make no causal difference.

This line of thought does not take us all the way to reifying *redness*, *brittleness*, and *hardness*. Rather, it lends plausibility to the claim that, among the things that exist are *ways things are*, i.e., properties and relations. *Which* properties and relations there are is a question it takes a great deal of science and philosophy to answer. *Hardness* and *brittleness* might make the final cut, or they might not. *Having a charge of -1*, *having spin up*, and other fundamental physical properties are likely to make the cut, but our reason for supposing they might is the same as our reason for having ever supposed *redness* and *brittleness* might make it: fundamental particles interact not holus-bolus but in virtue of the ways they are.⁶

The Quinean motivation begins with the observation that, according to our best scientific understanding of the world, everything that happens involves the redistribution of matter and energy in space and time. To posit causally inert entities is to posit entities whose

existence is irrelevant to the distribution of matter and energy in the universe. Such posits are warranted, in Quine's view, only on the grounds of indispensability to our theories of how the distribution of matter and energy in the world unfolds. Mathematical objects have no causal powers of their own, but (at least on Quine's view), they are indispensable posits of our best theory of the natural world [Quine 1980]. When it comes to properties and relations, though, as further entities over and above their bearers and the mathematical menagerie to which we are already committed, we should consider what work positing them does for us. Quine himself envisioned no work for them. I am willing to be slightly more tolerant. Perhaps there are *some* properties and relations among our scientific everything, but only those that make a difference to the distribution of matter and energy in space and time. That is, only those that pass Eleatic muster.

Two features of this view will be important in the arguments to follow. First, we should not simply assume that all the predicates of a theory designate genuine properties and relations, even if we accept the theory. A theory can be true even when there are no properties or relations answering to its predicates.⁷ Second, when a relation is genuine, its instantiation makes a causal difference. Genuine relations make a difference to the powers objects have or manifest, and that means they make a difference to the distribution of matter and energy in space-time.

3. Arguments for the Realization Principle

There are two main ways to argue for the realization principle. One is to point to examples of higher-level properties realized by lower-level properties. The other is to claim that we need to posit the realization relation to account for the structure of science. Neither line of argument is convincing.

First, consider arguments from examples. Sometimes, these are thought experiments, such as the famous ones concerning humans and Martians with nothing physical in common who both suffer pain. Sometimes, they supposedly depend on hard scientific fact: Octopi and humans both have eyes, though their eyes are analogous and not homologous, and they are structurally different from one another. The recent real estate bubble in the United States and the Dutch tulip bulb bubble of the 1630s are both economic bubbles, but they have nothing physical in common. Given physicalism, though, the physical world is all there is, so we need an account of how properties such as *having an eye* or *being a bubble* relate to fundamental physical reality. The answer, supposedly, is realization. There is a relation of metaphysical determination between lower and higher levels; the fundamental physical properties *realize* the properties addressed by special sciences.

Whether the arguments involve Martians, octopi, or tulip bulbs, they always depend on an appeal to intuition. In particular, they depend on the intuition that the predicates of a higher-level theory—predicates such as ‘has an eye’, ‘has pains’, and ‘is a bubble’—designate genuine properties. The intuition is powerful because we believe the theories these predicates are embedded in, and the easiest way to understand how the theories are made true posits properties as semantic values for their predicates.

But we cannot just posit properties for our predicates willy-nilly. We must always be open to the discovery that some of our cherished predications are true *even though* their predicates do not designate properties. Arguments from examples of higher and lower-level properties ignore that possibility. They assume the reality of the higher-level properties in question, simply on the strength of the false intuition that the theories we accept couldn’t be true without them. Perhaps different physical properties in seventeenth century Holland and in twenty-first century America bear the realization relation to *being a bubble*, but that is

possible only if there is such a property—such an entity in our scientific everything—as *being a bubble* to begin with.

Moreover, there are legitimate reasons to doubt the reality of higher-level properties altogether. For example, Jaegwon Kim's exclusion argument purports to show that such properties as *being an eye*, *being in pain*, or *being a bubble* fail to pass Eleatic muster (see Kim [2005: Ch. 2] for a recent elaboration). As long as we suppose these properties are not identical with their supposed lower-level realizers, there is not enough causal work to go around, and the higher-level properties appear to be epiphenomenal. But the Eleatic Principle counsels us to deny the reality of what is epiphenomenal.⁸

Naturally, philosophers have tried to find ways to allow for the reality of higher-level properties despite Kim's worries, and their efforts are instructive. The best approach I know of is that of Sydney Shoemaker and others, who conceive of realization as the *inclusion of causal powers* [Shoemaker 2001; Clapp 2001]. On this account, higher-level properties do have causal powers, and those powers constitute a *proper subset* of the higher-level properties' lower-level realizers' causal powers. The view is clearest if we suppose determinate properties realize determinables, so consider the determinable property *red* and its determinates, *crimson* and *scarlet*. The latter two properties differ in their causal powers, and they are different from one another. But each realizes *red*, because each bestows the distinctive causal powers of redness on its bearers. The causal powers of *red* are among the causal powers of *crimson* and *scarlet*, and that is what it means to say that they realize redness. Likewise, if fundamental physical properties realize *being a bubble* (or another higher-level property), on this view, the causal powers of *being a bubble* (or whatever) are among the causal powers of the fundamental physical properties that realize it.

Though it is technically neat, the inclusion account of realization does more to undermine the reification of realization than to motivate it. First, the account shows how realization claims (i.e., claims of the form ‘ L realizes H ’) could be true despite there being no realization relation in our scientific everything. If we understand ‘ L realizes H ’ as ‘The set of causal powers individuated by H is a nonempty proper subset of the set of causal powers individuated by L ’, we are committed only to the existence of L , H , causal powers, and sets of causal powers. We are not thereby committed to positing realization as part of our scientific everything.⁹

Second, given naturalism, we should think of properties as entities that bestow causal powers. But we need not suppose crimson things have the causal powers of redness bestowed on them twice, once by their crimsonness and once by their redness. In fact, we need not suppose there is any such property as redness at all. On the alternative view, we find it useful to abstract away from the differences among things with properties that are not identical but, by virtue of conferring some of the same causal powers, are similar to one another. Our predicate, ‘is red’, gives us a way to do that. Once we see that the causal powers distinctive of *red* are among the causal powers of *crimson* and *scarlet*, we no longer need to suppose that redness is a genuine property.

The inclusion account of realization does not bolster the assumption that special sciences deal with higher-level properties, and, even granting that they do, it shows how realization claims could be true without being committed to the existence of a realization relation. So, it does not strengthen the arguments from examples for the realization principle. These features of the inclusion account also point to a weakness in arguments from the relationship between higher- and lower-level sciences.

I have in mind arguments that begin with the observation that our sciences seem to fall into a hierarchy with certain interesting features. “Higher-level” sciences tend to be autonomous of “lower-level” sciences, insofar as we cannot infer anything about how a lower-level science will describe or explain a phenomenon from higher-level descriptions or explanations of it. Higher-level sciences often deal with entities that are readily construed as *made of* entities lower-level sciences deal with. Finally, sciences produce *ceteris paribus* laws subject to disruption only “from below” or “at the same level,” not “from above.” For example, a psychological law might fail to hold in a particular case because it is pre-empted by a neurobiological factor or by another psychological factor, but psychological laws aren’t disrupted by macroeconomic factors.¹⁰ The notion of realization, and especially the notion of *multiple* realization, promises to explain why science has this structure.

But the explanation would be predicated on the idea that higher-level sciences concern higher-level properties, and an alternative understanding is available. We can suppose the predicates of higher-level sciences do not designate genuine properties at all. Instead, they enable us to keep track of properties that are similar (insofar as they confer some of the same causal powers as one another) without being exactly alike [Heil 1999; 2003]. The sort of higher-level science we would expect in that case would be one that gives us a useful but noisy picture of what is going on in the world. It would be a science whose “laws” are approximately true, subject to disruption in indefinitely many ways, and such that their disruption is generally explicable when we revert to a means of describing the world that is more fine-grained and less noisy. It is, in short, the sort of science that we have [Dennett 1991].

The availability of this alternative explanation significantly weakens the argument from the structure of science. The alternative explanation is simpler, in that it does not posit

higher-level properties, and it seems to explain the structure of science at least as well as the appeal to higher-level properties does. Reflection on the structure of science should not lead us immediately to accept the realization principle.

4. Against the Realization Principle

My criticisms of the arguments for the realization principle have involved doubts about the reality of higher-level properties. Those doubts suffice to cast doubt on the adequacy of the arguments. And if, in fact, there are no higher-level properties, the falsity of the realization principle follows immediately. Realization would be a relation without relata.

I am, however, officially neutral on the question of the existence of higher-level properties. So, I do not intend to argue against the realization principle by casting doubt on their existence. The point of the preceding discussion has been only to show that the likeliest arguments for the realization principle are not conclusive. In what follows, I will assume that there are higher-level properties, which are realized by lower-level properties. Nevertheless, the predicate ‘realizes’ does not designate an entity in our scientific everything, because the truth or falsity of realization claims makes no causal difference in the world. By the Eleatic criterion, then, there is no realization relation, and so the realization principle is false.

Realization is supposed to be an explanatory relation of metaphysical determination. So far as I can tell, that means two things. First, realization is the inverse of a metaphysical supervenience relation: Given that a lower-level property complex realizes a higher-level property, it is metaphysically impossible for something to have the lower-level complex without having the higher-level property.¹¹ Second, when a lower-level property complex realizes a higher-level property, things with the lower-level complex have the higher-level property *because* they have the lower-level complex. I will encapsulate these features of

realization by saying that, if a lower-level property complex L realizes a higher-level property H , then “ L is metaphysically sufficient for H ” and “ L explains H .”

It is obviously not contingent whether L is metaphysically sufficient for H ; given S5, if L is metaphysically sufficient for H , then it is necessary that L is metaphysically sufficient for H . But metaphysical sufficiency is cheap. My being a philosopher is metaphysically sufficient for five to be a prime number, and it is (given the essentiality of origins) metaphysically sufficient for my parents’ having existed. But my being a philosopher explains neither of those things. For L to realize H , it must not only be metaphysically sufficient for H , but it must be metaphysically sufficient in the right way to explain H .

Philosophers have given various accounts of what sort of metaphysical sufficiency makes L explain H in the right way to count as realizing H . Some might appeal to token identity: L realizes H if and only if, necessarily, all instances of L are instances of H , but it is not necessary that all instances of H are instances of L . Others (such as Shoemaker or Clapp) might appeal to the inclusion of causal powers: L realizes H if and only if the causal powers individuating H are among the causal powers individuating L , but not vice versa. Still others (such as Gillett) appeal to having one set of causal powers “in virtue” of another: An instance of L realizes an instance of H if and only if something has the powers individuating H in virtue of having the powers individuating L , or in virtue of having constituents with those powers.¹² On all of these accounts, though, it is not contingent whether L explains H ; if L explains H in the requisite way, it is necessary that L explains H in that way.

As philosophers typically think of realization, then, realization claims are not contingent. If a lower-level property complex realizes a higher-level property, it realizes that

property necessarily. And if a lower-level property complex does not realize a certain higher-level property, it could not possibly realize that property.¹³

For realization to be a genuine relation, there must be a causal difference between the case in which a given lower-level property realizes a given higher-level property and the case in which it does not. For there to be such a difference, though, there must be cases of both kinds. It must be possible for the lower-level property to realize the higher-level property, and it must be possible for the lower-level property not to realize the higher-level property. Because realization claims are not contingent, one of those cases is impossible. So, there is no such thing as the difference between cases of the two sorts, and so there is no causal difference that realization makes. Consequently, realization fails the Eleatic test for reality. It is not a constituent of our scientific everything.

Suppose, for example, that an instance of C-fibre stimulation realizes an instance of pain. It realizes that instance of pain necessarily, so there is no such case as that in which that C-fibre stimulation fails to realize that pain. There being no such case, there is no such thing as the difference between the case in which the C-fibre stimulation realizes the pain and the case in which it does not. So, the C-fibre stimulation's realizing the pain makes no causal difference in the world.

If we take a more Quinean route and identify "causal differences" with differences in the distribution of matter and energy in space-time, we can make the argument somewhat differently. If a lower-level property complex realizes a higher-level property, there is no possible distribution of matter and energy such that the lower-level complex exists without realizing higher-level property. So, there is no such thing as the "difference" its realizing the higher-level property makes to the distribution of matter and energy in space-time. So, it fails our test for being a genuine relation in our scientific everything. Realization is not real.

Given that there are higher-level properties, true realization claims turn out to be necessary. Their necessity means that they make no difference to the causal powers of anything or to the distribution of matter and energy through space-time.¹⁴ Consequently, the realization relation fails the Eleatic test for reality. There is no such thing as realization. Or, to be more cautious: The reification of realization is inconsistent with plausible naturalistic metaphysics.

5. Defences of the Realization Principle

This section addresses three ways one might try to defend the realization principle against the preceding arguments. The first challenges the premise that realization claims are not contingent. The second purports to show that the realization relation passes the Eleatic test after all. The third contends that realization deserves to be exempt from the Eleatic test because it is a “metaphysical” relation. None succeeds.

5.1 Defence 1: Realization is contingent.

The case against the realization principle depends on the premise that realization claims are not contingent. An obvious countermove would be to insist that realization claims are contingent after all, and so there is a difference between the cases in which they are true and the cases in which they are false.

There are two ways in which realization claims might be contingent. First, for a given lower-level property complex to realize a given higher-level property, some further, enabling conditions might be necessary, and it might be contingent whether the enabling conditions obtain. Second, even given that all the right enabling conditions obtain, one might claim it is metaphysically possible for the lower-level property complex to exist without realizing the higher-level property. This would require claiming that L could realize H without being metaphysically sufficient for H , even given that all the right enabling conditions obtain.

The point of positing realization is to bolster nonreductive physicalist realism. Any physicalism worthy of the name, though, holds that the totality of fundamental physical facts (e.g., the overall distribution of matter and energy through space-time) metaphysically determines all the facts there are. To hold that realization claims are contingent in the second way is really to abandon physicalism altogether. It is to maintain that there are some contingent facts—viz., the facts about what realizes what—that can vary independently of the low-level physical facts. But if we give up physicalism in order to save the realization principle, we have given up the very metaphysical view that made realization an attractive notion in the first place.

The first way in which realization claims might be contingent is much more plausible. It is almost a commonplace that nothing realizes anything unless the “background conditions” are right; C-fibres firing in a Petri dish wouldn’t realize pain even if they would realize pain in a functioning organism. Even if realization claims are contingent in this way, though, the realization principle is no better off.

To pass the Eleatic test, realization must make a difference to the causal powers of particulars. Suppose L realizes H , but only if contingent enabling conditions C obtain. There is thus a causal difference between worlds where L realizes H and otherwise similar worlds where L does not realize H : C obtains in the former but not the latter. However, we cannot call this a difference *that realization makes*. L ’s realizing H does not make the enabling conditions obtain. If anything, the enabling conditions make L realize H . The enabling conditions are what make the causal difference in this case (or, L ’s obtaining in those conditions), not realization itself. Given L , H , and the enabling conditions, there is no further work for the realization relation to do.

In addition, as mentioned above (n. 11), we can always just include the enabling conditions as part of the lower-level complex that realizes *H*. Then it is clearly not contingent that the lower-level complex realizes *H*, and there is no difference realization makes over and above that made by its relata. Note that the realization relation itself is not part of the “enabling conditions.” If it were, it would be something all enabling conditions have in common, no matter what they enable to realize what, but we should expect no such commonality. There is no reason to suppose the conditions enabling an anatomical configuration to realize *being an eye* have anything at all in common with those enabling a pattern of buying and selling to realize *being an economic bubble*.

The only physicalistically plausible way realization claims might be contingent, then, is irrelevant to the argument against the realization principle. It does not provide a way for realization to make a causal difference. Nevertheless, there is an important point to be taken from this defence of the realization principle: Realization claims might report on contingent features of the world (e.g., the obtaining of appropriate enabling conditions) even though realization itself is not a piece of our scientific everything.

5.2 Defence 2: If C-fibre stimulation didn't realize pain, we couldn't cause pain by stimulating C-fibres; so, realization makes a causal difference.

This defence aims to show that realization claims make a causal difference after all, and there are several ways to formulate it. One is to say that realization makes a causal difference by making an existential difference:¹⁵ But for its being realized by a lower-level property, a higher-level property (or property instance) would not even exist. The suggestion is that realization thus gets credit for some of the causal difference the instantiation of the higher-level property makes. Another way of putting it is that realization gives us a causal handle on higher-level properties we would not have otherwise. We can manipulate

something's higher-level properties by manipulating its lower-level properties, and that could not happen if the lower-level properties did not realize the higher-level properties. Yet a third way of putting the point is this: For any L and H , if L realizes H , it is impossible for there to be an instance of L without an instance of H . But if L does not realize H , it may be possible for there to be an instance of L without an instance of H . So, if L realizes H , we can cause instances of H by causing instances of L , but we have no such guarantee if L does not realize H . Thus realization appears to make a causal difference after all.

These lines of thought do seem to get something right. We can infer from ' L realizes H ' that it is metaphysically impossible to have an instance of L without an instance of H , and we cannot make that inference from ' L does not realize H '. So, we can infer that, if L realizes H , then causing an instance of L will always cause an instance of H . But that is not enough to establish that realization makes a causal difference.

Consider the mathematical relation of being less than, $<$. Like all mathematical relations, $<$ is causally inert. Yet, from ' $x < y$ ', we can infer that whatever causes y beans to fall out of a bean bag will also cause more than x beans to fall out of that bag. Such an inference is insufficient to show that $<$ makes a causal difference, and it is equally insufficient to show that realization makes one.

For realization to make a causal difference, it is not enough for realization claims and their negations to have different implications, or even for them to have different implications concerning causation. Rather, they must nontrivially partition the metaphysically possible distributions of causal powers to objects (or of matter and energy through space-time). And as we have seen, the necessity of realization claims prevents them from doing that.

There is no meaningful metaphysical contrast between the case in which a given lower-level property realizes a given higher-level property and the case in which it does not. To think there is one is to confuse the metaphysical contrast between L 's realizing H and L 's not realizing H with one of two other contrasts, neither of which is relevant to the reality of realization.

The first is the epistemic contrast between ' L realizes H ' and ' L does not realize H '. Like any meaningful claim, a realization claim has different implications from its negation. Moreover, although we are guaranteed that not both a realization claim and its denial are *metaphysically* possible, in our ignorance they might both be *epistemically* possible. Because the claims have different implications, it could then be epistemically possible that it is metaphysically possible to have an instance of L without H , and it could also be epistemically possible that it is metaphysically *impossible* to have L without H . There is a contrast here, but it is a contrast among epistemic possibilities, not metaphysical ones. It is an epistemic difference that *knowing* a realization claim makes, not a causal difference the *truth* of a realization claim makes.

A second contrast that might seem to help the realization principle is between (a) pairs of properties such that the first realizes the second and (b) pairs such that the first does not realize the second. Pairs in the former class have it in common that there can be no instance of their first property without an instance of their second. Pairs in the latter class do not have that in common.

We might say that whether it is possible to have L without H "depends on" which class includes the pair $\langle L, H \rangle$, but it would be a mistake to think of that as a difference realization makes. First, the relevant sort of "depending" is epistemic, not metaphysical. It is the same as the sense in which we might say that whether or not O.J. is guilty depends on

whether the glove fits. Second, membership in the class of pairs such that the first realizes the second does not *make* it impossible to have L without H . If the vocabulary of “making” is sensible at all here, it is the impossibility of having L without H that makes $\langle L, H \rangle$ a member of that class. The contrast between these classes, which is just the contrast between the extension of ‘realizes’ and its anti-extension, is only to be expected because realization claims are neither meaningless nor vacuous. Still, it is not a causal difference made by the truth of a realization claim, and so it is not enough to show that realization passes the Eleatic test.

5.3 Defence 3: Realization is a metaphysical relation, and metaphysical relations should be exempt from the Eleatic criterion.

A final defence of the realization principle grants that realization makes no causal difference, but it claims that realization is a *metaphysical* relation, and so it can be real without making a causal difference. I suspect this view is popular.

What is wrong with the view depends on what is meant by ‘metaphysical’. Maybe metaphysical relations are just logical relations. Then at least one of their relata must be something representational, like a proposition, a sentence, or a theory.¹⁶ If realization is a metaphysical relation in this sense, then it either relates our theories to the world or it relates some of our theories to one another. Either way, it is not a relation between higher-level *properties* and lower-level properties. At best, it is a relation between higher-level *predicates* and lower-level properties. But then the realization principle is still false; it says realization is a relation between lower-level and higher-level properties.

On the other hand, metaphysical relations might be non-logical (so they can be relations among properties) while also being non-natural (so they need not have causal

powers to be real). If realization were such a relation, it would be immune to my criticisms. But thinking of realization in that way is pointless.

The point of positing a realization relation was to show how special science properties could be part of the natural world after all. If we treat realization as a non-logical, non-natural relation, then special science properties qualify as natural because they bear a *supernatural* relation to physical properties. If the goal is to show that the natural world is the only world, that our scientific everything is *everything*, then this is not progress. Moreover, if we are willing to treat realization as a supernatural relation, we have already admitted that there is something supernatural. We might as well then allow that special-science properties are likewise supernatural (or superphysical, anyway), in which case we find ourselves no longer in need of positing a realization relation. Making the realization relation non-logical and non-natural saves the relation at the expense of the motivation for positing it.

Perhaps the claim that realization is “metaphysical” means that it is *purely and irreducibly modal*.¹⁷ It would then be a mistake to look for differences the truth of a realization claim makes within worlds. We should look for differences *across* worlds. For example, the claim that *L* realizes *H* might mean just that *L* is metaphysically sufficient for *H*. In that case, it is not at all worrisome that the truth of a realization claim makes no causal difference. Realization claims report facts about what is possible or necessary, not about what is actually the case. But we are justified in applying the Eleatic principle only to claims about what is actually the case. Realization, on this view, fails the Eleatic test because it is metaphysical and modal, not because it is unreal.

Let us suppose realization claims are purely modal; they purport only to tell us something about the distributions of properties across worlds, not to carry any information about how things are in any particular world. We are still left with no good reason to reify

realization as a *relation*. By hypothesis, it makes no causal difference and explains no causal interactions. Moreover, we need not reify realization for realization claims to be true; they are, on this proposal, made true by the distribution of properties across worlds. A reified realization relation appears to do no work, so there is no good reason to posit it.¹⁸

6. Special Sciences without Realization

Rejecting the realization principle means not reifying realization, that is, not treating it as an entity in our scientific everything. It requires neither antirealism about the special sciences nor reductionism nor eliminativism about their subject matter. Nor does it require supposing the special sciences address a radically autonomous realm of objects and properties.

The unreality of realization is consistent with all those views, but it is also consistent with a view that honours the important insights garden variety nonreductive physicalists express using claims such as:

(A) Psychological properties are realized by neurological properties.

(B) Psychological properties are multiply realizable.

Those insights can be expressed in at least two ways without reifying realization.

One way is to follow John Heil's [1999; 2003] suggestion and express the important insights in terms of predicates and truthmaking. In place of (A), we say that the truthmakers of psychological predications are always also truthmakers of neurological predications, but not vice versa. In place of (B), we say that psychological predications have a range of possible truthmakers, and there are no non-psychological predications (e.g., those of neurology or of quantum physics) with precisely the same range of possible truthmakers.

The background idea is that special science predicates class things together not in virtue of

shared properties but in virtue of their having *similar* properties, where properties are similar in virtue of bestowing some (but not all) of the same causal powers.

A second way is to follow Daniel Dennett [1991] and take seriously the idea that our theories are not mere records of what there is and what it does, but tools we use to get by in the world. As such, they are answerable to the way the world is, but they are also answerable to our needs and limitations. Our theories represent the world with varying amounts of signal and noise. For some purposes, detailed, relatively noiseless representations are needed. For other purposes, such representations are inappropriate because they omit certain patterns that emerge only noisily in nature or because using them would render our problems intractable. We cannot use noisier theories to do the work of less noisy ones, and we cannot use less noisy theories to do the work of noisier ones.

On this approach, the insight behind (A) is that psychology gives us coarser grained, noisier representation of the same states of affairs of which neurology gives us finer grained, less noisy representations. The insight behind (B) is that the patterns psychology addresses emerge noisily in nature; they do not correspond to patterns addressed through the less noisy representations of neuroscience.

I suspect that the Heil-inspired and Dennett-inspired approaches are compatible with one another. On either approach, one might be tempted to adopt the vocabulary of realization. Properly understood, a claim such as (A) or (B) need not commit us to reifying realization. Their mentions of realization can be paraphrased away and reformulated as claims about the semantics and commitments of our theories in relation to one another, rather than deep metaphysical claims. Nevertheless, the temptation to continue speaking of realization is worth avoiding. The main theoretical purpose of claims such as (A) and (B) is

to formulate and ground versions of nonreductive physicalist realism, but they are ill-suited to the task if we do not reify realization.

They are ill-suited for two reasons. First, on the Heilian or Dennettian understandings, claims such as (A) and (B) would turn out to be neutral with respect to the reality of higher-level properties. They might be useful in formulating or grounding *nonreductivism*, but they are not poised to help ground *realism* about those properties at all.

Second, one way the notion of realization helps to ground nonreductive physicalist realism is by identifying genuine similarities, both among the physical substrates of higher-level properties, and across instances of realization. Despite their lack of any common intrinsic features, the physical substrates of human and Martian pain both *realize pain*, and that commonality is enough to bind them into a natural kind. Moreover, reifying realization lets us claim that there is an important metaphysical respect in which C-fibre firing is to *being in pain* as certain instances of buying and selling are to *being a bubble* and as certain electronic events in my computer are to *Microsoft Word*. But if we understand claims such as (A) and (B) consistently with the unreality of realization, they lend no support to these metaphysical claims. At most, they point out semantic similarities among our theories, not metaphysical similarities among their subject matter.

Semantics is not metaphysics, and the notion of realization is part and parcel of a specific metaphysical outlook. Pressing the notion of realization into semantic service, rather than just giving it up and adopting semantic vocabulary instead, would only invite confusion or worse.

7. Conclusion

The realization principle is false, or at least inconsistent with highly plausible principles of naturalistic metaphysics, including the Eleatic criterion of reality for properties

and relations. The arguments in favour of the principle are not compelling, and the supposed relation of realization makes no difference to anything's powers or to the distribution of matter and energy in space-time. Given the principles of naturalistic metaphysics, then, we should not posit a realization relation. Without reifying realization, we can still capture the important insights usually expressed in terms of realization. When we do so, those insights become less suited to elaborating and defending nonreductive physicalist realism about special science properties. Whether it is nonreductivism, physicalism, or realism that must give way is a question I leave for future consideration.

REFERENCES

- Armstrong, D. 1999. A Naturalist Program: Epistemology and Ontology, *Proceedings and Addresses of the American Philosophical Association* 73/2: 77–89.
- Bechtel, B. and Mundale, J. 1999. Multiple Realizability Revisited: Linking Cognitive and Neural States, *Philosophy of Science* 66: 175–207.
- Clapp, L. 2001. Disjunctive Properties: Multiple Realizations, *The Journal of Philosophy* 98/3: 111–136.
- Clark, A. and Chalmers, D. 1998. The Extended Mind, *Analysis* 58: 10–23.
- Dennett, D. 1991. Real Patterns, *The Journal of Philosophy* 88/1: 27–51.
- Fodor, J. 1974. Special Sciences, *Synthese* 28/2: 97–115.
- Fodor, J. 1997. Special Sciences: Still Autonomous After All These Tears, *Noûs* 31 (Supplement) *Philosophical Perspectives 11: Mind, Causation, and World*: 149–163.
- Gillett, C. 2002. The Dimensions of Realization: A Critique of the Standard View, *Analysis* 62/4: 316–323.
- Heil, J. 1999. Multiple Realizability, *American Philosophical Quarterly*, 36/33: 189–208.
- Heil, J. 2003. *From an Ontological Point of View*, Oxford: Oxford University Press.
- Hendel, G. 2001. Realization, *Crítica: Revista Hispanica de Filosofía* 33/98: 41–70.

- Kim, J. 2005. *Physicalism, or Something Near Enough*, Princeton: Princeton University Press.
- Martin, C. B. 1997. On the Need for Properties: The Road to Pythagoreanism and Back, *Synthese* 112/2: 193–231.
- Polger, T. 2004. Neural Machinery and Realization, *Philosophy of Science*. 71/5: 997–1006.
- Putnam, H. 1967. Psychological Predicates, in *Art, Mind and Religion*, ed. W. H. Capitan and D. D. Merrill, Pittsburgh: University of Pittsburgh Press: 37–48.
- Quine, W. V. 1951. Ontology and Ideology, *Philosophical Studies*, 2/1: 11–15.
- Quine, W. V. 1970. *Philosophy of Logic*, Englewood Cliffs, NJ: Prentice-Hall.
- Quine, W. V. 1980. On What There Is, in *From a Logical Point of View*, 2nd ed, Cambridge, MA: Harvard University Press: 1–19.
- Quine, W. V. 1983. Ontology and Ideology Revisited, *The Journal of Philosophy* 80/9: 499–502.
- Rupert, R. 2007. Realization, Completers, and *Ceteris Paribus* Laws in Psychology, *British Journal for the Philosophy of Science* 58: 1–11.
- Shoemaker, S. 2001. Realization and Mental Causation, in *Physicalism and its Discontents*, ed. Carl Gillett and Barry Loewer, Cambridge: Cambridge University Press: 74–98.

¹ This paper resulted from work done at John Heil’s 2006 *Mind and Metaphysics* NEH Summer Seminar at Washington University in St. Louis. An early version of it was presented in a special symposium on realization at the 2007 meeting of the Southern Society for Philosophy and Psychology. I owe thanks to all the participants in both events for helpful discussions, and I owe particular thanks to Ken Aizawa, Torin Alter, Jason Ford, Carl Gillett, John Heil, Nicholas Helms, Pete Mandik, John Post, Gene Witmer, Michelle Wrenn, Tad Zawidzki, and two anonymous referees.

² This is a common interpretation of debates about “localization” within the cognitive sciences, as well as the debate over the “extended mind hypothesis.”

³ This is an important point that naturalists all too often ignore. In my view, physics never stopped quantifying over the luminiferous ether, because it never *started* quantifying over the ether. When our physical theories changed from asserting to denying that there is an ether, the change was ideological. Changes in the ontology of a theory are changes in its scope or in its domain of application, not changes in what the theory says. Changes in what the theory says are ideological.

⁴ Eleatists and Quineans are often contrasted with one another, for good reason. Eleatists ordinarily believe in properties and relations but are suspicious of mathematical objects. Quineans ordinarily believe in mathematical objects but are suspicious of properties and relations. But Quineans should be open to the *ideological* questions of whether there are

properties and of which ones there are, and they should be open to the Eleatic answers to those questions—provided that the Eleatic answer is the one our best scientific view of the world commits us to.

⁵ A better understanding of physics might require thinking of gravitational forces in terms of masses' relations to a field rather than their distances from one another, and a complete physics might dispense with the causal relevance of distance relations altogether. I mean for distance to provide an example of how a relation might pass the Eleatic test. Any other relation that plays a role in the causal laws of physics could serve as well.

⁶ One might worry that *hardness* and *brittleness* are just the sorts of multiply realizable properties whose reality it is the business of this paper to deny. The worry would be misplaced, though, because that is not the business of this paper. I am officially neutral on the question whether such properties as *hardness* and *brittleness* exist, and I am even neutral on the question whether they are multiply realized. My claim is that there is no such entity as the realization relation; even if *brittleness* is multiply realized by properties *F*, *G*, and *H*, there is no fifth thing, the realization relation, that connects *brittleness* to *F*, to *G*, and also to *H*. See Sections 4 and 6 below.

⁷ For example, pure mathematics can get by with just two predicates, '=' and '∈', neither of which corresponds to any mathematical object.

⁸ Kim's argument is familiar and space is short, so I will not discuss its details here.

⁹ Even apart from the inclusion account, this is a general problem for arguments from examples to the realization principle. The examples give us, at most, true realization claims. Once we give up the assumption that all our predicates designate genuine properties and relations, the existence of true realization claims is far from enough to settle the question whether 'realizes' designates a genuine relation.

¹⁰ Whether or not we *are* in a depression is a macroeconomic matter. Whether or not Jones *thinks* we are in a depression is a psychological factor. The depressed state of the economy cannot cause psychological laws ordinarily true of Jones to fail, but Jones's thinking the economy is depressed might.

¹¹ On some views, it is metaphysically possible for something to have a lower-level realizer of a higher-level property without having the higher-level property, because the "contextual" or "background" conditions are not right. But it is necessary, on these views, that the higher-level property be on the scene whenever the conditions are right and the lower-level property complex is there. So, this complication will be irrelevant to most of what I say here. We can take the "background" conditions to be part of the lower-level property complex itself.

¹² For my purposes, it is not necessary to take a side on the question whether the fundamental cases of realization are cases in which one property realizes another or cases in which one property instance realizes another.

¹³ Every account of realization either analyzes the claim that *L* realizes *H* as the claim that something is necessarily so (e.g., as the claim that, necessarily, if *x* has *L* then *x* has *H*) or as a claim that happens to be necessary (as in the case of the causal inclusion account). In the first case, '*L* realizes *H*' means $\Box p$ (where what '*p*' stands for depends on which theory of realization is in play). Given S5, $\Box p$ implies $\Box\Box p$, so '*L* realizes *H*' implies 'Necessarily, *L* realizes *H*'. In the second case, '*L* realizes *H*' is identified with a claim that, if true, is necessarily true. (For example, the causal powers that individuate a property are essential to it, so if the powers individuating *H* are among the powers individuating *L*, then they

are necessarily among them.) Thus, in the second case too, '*L* realizes *HP*' implies 'Necessarily, *L* realizes *HP*'. I thank an anonymous referee for urging me to make this point explicit.

¹⁴ It might be contingent that a lower-level property complex *L exists* or is instantiated, and its existence or instantiation might make a causal difference to the world. The point here is that there is no further causal difference *L*'s realization of some higher-level property *H* makes. All the causal difference is already made by *L* itself, and realization adds nothing.

¹⁵ I thank Carl Gillett, Jason Ford, and an anonymous referee for suggesting versions of this objection to me.

¹⁶ Only representations have logical properties, and at least one relatum of a logical relation must have logical properties. If there were such a property as *being either square or non-square*, it would not be a logical property, but a property whose ascriptions to entities are logically true. Logical truth, if there is such a property, *would* be a logical property.

¹⁷ I thank an anonymous referee for encouraging me to consider this idea of what 'metaphysical' might mean.

¹⁸ The same might be said for various other metaphysical "relations," possibly including supervenience, causation, instantiation, and identity. If they make no causal difference, they fail the Eleatic test and we should not reify them. It is not clear, though, what purpose would be served by reifying these relations in the first place. Those concepts' most important jobs do not require there to be entities answering to them. (I thank anonymous referees for encouraging me to acknowledge this point.)