This is a preprint of an article whose final and definitive form will be published in *The Australasian Journal of Philosophy* 2007; *The Australasian Journal of Philosophy* is available online at:

http://journalsonline.tandf.co.uk/.

REALIZATION AND THE METAPHYSICS OF MIND*

Thomas W. Polger
Department of Philosophy
University of Cincinnati
Cincinnati, OH 45221-0374 USA

thomas.polger@uc.edu

According to a familiar view in philosophy of mind pioneered by Hilary Putnam [1967] and elaborated by Jerry Fodor [1974], mental states or properties are *realized* by brain states or properties but are not identical to them. More generally, it is commonly held to be characteristic of the entities, states, properties, or events of the special sciences overall (not only psychology) that they are *realized* by but not identical to physico-chemical entities, states, properties, or events. Call this view *realization physicalism* (RP). The attractive feature of RP is that it promises that we can have our cake and eat it too. Special science entities and properties are not identical to physico-chemical entities and properties, and therefore cannot be reduced to or eliminated in favor of them. Yet because the special science entities and properties are realized by unimpeachably natural entities and properties, the special science entities and properties are themselves "naturalistically kosher." In this way, RP is (or is a version of) non-reductive physicalism. If it works, this is a pretty good trick. Whether it works obviously hinges on the nature of the realization relation.

You might therefore expect that a great deal has been done to clarify the realization relation. But you would be wrong. Both the formulation of RP and its defense by way of familiar multiple realizability arguments have proceeded almost entirely without discussion of the realization relation itself. Once in awhile it is noticed that realization is in need of scrutiny, but almost invariably that is left as a project for another day [e.g., Horgan 1993; Horgan and Tienson 1996; Kim 1996, 1998].[1] At last that day has come, and recently some philosophers have examined the realization relation directly [Poland 1994, Wilson 2001; Shoemaker 2001, 2003; Melnyk 2003; Polger 2004a]. Of particular interest, Carl Gillett has defended an explicit account of realization [2002, 2003]. Gillett's theory is important for three reasons. First, it is one of the most detailed defenses of a realization relation on offer. Second, Gillett's account apparently absorbs the views of Jaegwon Kim [1998] and Sydney Shoemaker [2001, 2003]. While there are important differences among the three, there are also core similarities that can usefully be discussed in terms of Gillett's more encompassing account. Third, Gillett directly addresses the implications of his account of realization for the multiple realizability arguments that are central to RP and which have recently come under renewed scrutiny [Bickle 1998; Bechtel and Mundale 1999; Sober 1999; Shapiro 2000, 2004; Clapp 2001; Polger 2002, 2004a].

In the first section of this paper I sketch some obligations of the realization relation. *Realization* is a term of art, introduced to formulate functionalism, the principal variety of RP. Simply put, an account of realization ought to make sense of the kinds of cases that defenders of RP count as realization. In the second section I outline Gillett's "dimensioned view" of realization. In the third and fourth sections I argue that Gillett's

account fails to meet its obligations. Like Kim's and Shoemaker's views, the dimensioned view of realization does not accommodate some textbook cases of realization. But unlike Kim's and Shoemaker's views, Gillett's dimensioned view counts as examples of realization some cases that should not count if RP is to be distinguished from alternative versions of physicalism in the usual ways. I conclude that the relation described by Gillett cannot be realization, and in the fifth section I offer my own alternative approach to understanding realization. Finally I revisit the implications of accounts of realization for multiple realization arguments.

Let me be clear from the start that my aim in this paper is not to defend RP, nor to formulate an account of realization that ensures that RP is correct. My goal, rather, is to defend an account of realization that reveals RP to be a substantial and interesting theory, distinct from alternative theories and worthy of attention. It is quite a different matter whether, in final consideration, RP is the best available theory of mind. I have my doubts. But, as I will repeatedly emphasize, we cannot even take RP seriously until we understand what the realization relation is supposed to be. And it will not help RP to trivialize its claims or to conflate it with its competitors. Even if RP is an imperfect theory, properly understanding it will help us understand the desiderata on a metaphysical theory of mind.

1. The Many Faces of "Realization"

What is realization? Though the point is rudimentary, let me observe that not all cases of so-called realization are of the sort at stake with RP. Consider:

(1) Sally realized that it would rain, so she brought her umbrella.

(2) The advertisement claimed that the product would help me realize my potential.

(3) Rodin's *The Thinker* was realized in wax and bronze.

(4) The role of the air-fuel mixer is realized by a fuel injector in my car.

(5) My computer currently realizes Microsoft Word.

(6) Memory fixation is realized in humans by long term potentiation of neurons.

It is plain that the relationship between me and my potential is not the same as the relationship between Sally and the (proposition that it would) rain. And neither of those seems to be the same relationship as that between computer hardware and computer programs. Some of the above examples are not of a kind with the realization of special science properties and entities that concern RP.

(1) is perhaps the most colloquial use of 'realize' but it is clearly not the sort involved in RP. The rain is not realized by Sally, and Sally is not realized by the rain. Sally "realizes" something in that she comes to believe that the rain is real. But the reality of the rain does not depend on Sally in they way that realized entities and properties depend on their realizers according to RP.[2] Example (2) also has a strong claim to be an ordinary use of 'realize' but it does not look like a *prima facie* case of RP. (2) asserts that using some product will (partially) cause me to achieve some goal. But achieving a goal is not straightforwardly a matter of me standing in some relation to my potential, whatever potentials may be. Rather, achieving (that is, "realizing") my potential is a matter of taking actions that cause or result in some goal state.[3]

Example (3) is more difficult. On one reading, (3) describes a process by which first wax and then bronze are manipulated to produce a statue. A wax model was part of the process of realizing (manufacturing) the statue. The wax is destroyed in the production of the statue, so it is not a realizer of the statue in the way relevant to RP. Similarly, one might say that many workers and much equipment were involved with realizing *The Thinker*. If this is right, then (3) is similar to (2) in using a notion of realization as "bringing about." On another reading, which is perhaps more familiar to philosophers, (3) seems to assert that realization is the relation that obtains between a statue and some materials. It says that the statue is (at a time, at least) realized by a bit of wax, and is realized (still) by a bit of bronze—by many such bits of bronze, as it turns out. The wax and the bronze are realizers, on this reading, of the same statue. On this view, realization is a relationship that holds between the wax or bronze (respectively) and *The Thinker*. Of course, the relationship between statues and the material of which they are made is of longstanding philosophical interest. Could that relation be realization? If it might be, then (3) sounds more like (4) and it is the kind of thing that RP claims. So it seems to me that an account of realization might accommodate (3), or might explain it away as a case "realization" in the sense of causal production.

By comparison, examples (4) through (6) are common in discussions of RP. Mechanical, computational, and psychological entities and properties are canonical cases of realization. Replaceable mechanical parts are perhaps the most frequent examples of realization, as in (4). Combustion engines mix air and fuel for ignition; in older vehicles this task was accomplished by carburetors, but fuel injectors are more common with newer engines. My car's engine must have an air-fuel mixer, and it

happens that role is occupied by a fuel injector. It would be no easy matter to replace my car's fuel injector with a carburetor or with some novel device, but with some ingenuity it could be done. (Several reality television shows feature people building machines with just such improvised parts.) Likewise, per (5), my computer hardware realizes or implements the programs and algorithms that it runs. Certain electrical states of the device realize computational states such as, say, storing the contents of the last *copy* operation. The electrical activity of the device is not identical to any program state of Microsoft Word, but it implements or realizes such program states.

Finally we come to just the kind of case that is at stake for RP theorists—the realization of psychological states by brain states, as in (6). RP about the mind began as the thesis that the brain should be understood as a computing machine and the mind should be understood as a program or set of programs. Although the literal computational view is out of fashion in many circles, the idea that the relationship between brain and mind is the same as that between machines and programs is quite current.

The implication of the present considerations is that not every relationship that is called "realization" is a case of realization as construed by RP. One who aims to give an account of realization will have to pick examples carefully. Realization is a non-destructive, non-causal, synchronic dependence relation.[4] Thus an account of realization that is intended to defend or explain RP need not accommodate cases like (1) and (2) as realization.[5] Example (3) is contestable; an account of realization could allow that it is a case of realization, or it might be able to discount (3). But any account of realization that is intended to defend or explain RP ought to affirm examples like (4)-(6)

as genuine cases of realization. These examples may not give us nonnegotiable conditions on realization but they do offer a sketch of the features that realization is typically taken to have.

2. The Dimensioned View of Realization

We are now in a position to evaluate a representative account of realization that Carl Gillett calls the "dimensioned view" [2002, 2003]. Gillett understands realization as a relationship among property instances. In his framework, properties are *instantiated* by individuals whereas properties are *realized* by other properties.[6] Specifically, Gillett proposes:

> Property/relation instance(s) F1-Fn realize an instance of a property G, in an individual *s*, *if and only if s* has powers that are individuative of an instance of G in virtue of the powers contributed by F1-Fn to *s* or *s*'s constituent(s), but not vice versa. [2002: 322][7]

The "powers" in question here are causal powers, as Gillett's account is framed in terms of the causal theory of properties. On his view, a system *s* instantiates a certain property G when it or its parts have the causal powers that individuate G. That is, the causal powers individuative of G may belong to its bearer *s*, or to the parts of *s*. The second disjunct, allowing realization of G by properties of a system's constituents, is Gillett's innovation. Permitting realization to span mereological levels is the basis for dubbing this the "dimensioned" view. Formulations that lack this feature, specifically those of Kim [1998] and Shoemaker [2001], construe realization as a "flat" intra-level relation. For example, Shoemaker suggests that "property X realizes property Y just in case the

conditional powers bestowed by Y are a subset of the conditional powers bestowed by X" [Shoemaker 2001: 78]. Such accounts require that the powers individuative of G be instantiated in *s* itself.

The flat view is easy enough to understand. A certain individual K realizes a certain property or state, e.g., being an air-fuel mixer, if and only if K has the causal powers that are individuative of [the property of] being an air-fuel mixer. This, intuitively, is what it is for K to occupy the role of air-fuel mixer. Individual K's properties simply are those that are distinctive of air-fuel mixers. In my car, it is a fuel injector that occupies that role, that realizes the air-fuel mixer. That is, it—that very thing—has the properties which are individuative of that component of my car. The flat view earns its name because realized and realizer properties must be instantiated in the same individual, at the same mereological level.

The trouble, Gillett argues, is that the flat view neglects the possibility that realized and realizer properties might be instantiated in different individuals. Specifically, it may be that some property is instantiated in an object in virtue of the instantiation of realizer properties not in that very same object but instead in the parts of that object. So realization may occur in two ways: the powers that individuate the realized property may be contributed by "horizontal" intra-individual properties or "vertical" inter-individual properties. Hence the second disjunct of Gillett's formulation, and the "dimensioned" view of realization in contrast to the old "flat" view captured by the first disjunct alone. Gillett's dimensioned view thus absorbs and expands on the views of Kim and Shoemaker. I'll call these three *causal approaches* to

realization because they all agree that realized properties are to be individuated causally.

Gillett's main example in support of the dimensioned view is the realization of the property of hardness in a diamond.[8]

(7) The hardness of a diamond is realized by the (instances of) properties of the atoms composing the diamond.

According to Gillett, hardness should be understood as a property instantiated in one individual, viz., the diamond, which is realized in it in virtue of the properties instantiated in many different individuals, viz., the carbon atoms that compose the diamond. (Diamonds are macromolecules, and so are composed of atoms.) If so, realizer and realized properties are not instantiated in the same individual.

We will return to diamonds shortly. For the moment we need only remember that the dimensioned view endorses all the cases of realization counted by the flat view and adds more. First we will consider a line of critique that applies equally to Gillett's dimensioned view of realization and to the flat causal views of Kim and Shoemaker. Then we will examine a special problem that arises for the dimensioned view.

3. The Dimensioned View Does Not Compute

When textbooks and professors explain RP they invariably appeal to computing devices, and usually to Turing machines. After all, this is how functionalism—the most prominent form of RP—is introduced into the literature.[9] In framing his functionalist hypothesis Putnam explicates functional states in terms of probabilistic automata [1967]. The basic idea of functionalism is that mental states are in some sense states of

brains, but are not identical to brain states—just as machine program states are implemented by but not identical to the states of a physical device. Functional states are individuated differently than physical states of brains or machines. In what sense, then, are mental properties or states nevertheless physical properties or states? Putnam introduced the term *realization* for this relation.[10] As Putnam used the term, realization is the relation between machines and programs and also between brains and minds. For this reason I argued that examples (5) and (6) should be accommodated by any account of realization. (6) is an example of the primary thesis of RP. And (5) is what gives substance to (6) as a theory of the nature of minds. Brains realize or implement minds just as physical devices realize or implement programs. Indeed this analogy exhausts what was said about realization for many years. If (5) is not a case of realization, then we have no idea what RP asserts.

I am not here going to say anything about the attributes or failings of functionalism, much less the probabilistic automata formulation. What I claim is that any account of realization that is supposed be relevant to RP is going to have to make sense of realization of the kinds of states about which functionalists have theorized, including realization of computational states such as those of a probabilistic automaton. But Gillett's dimensioned view focuses only on some cases of realization. It thereby neglects the archetypal case of realization of a formal program or algorithm.[11]

Consider the familiar case of a Turing machine that implements the addition function. What must be the case for a physical system to realize addition? The physical system must have states whose causal relations to one another somehow "correspond" or "map" onto the mathematical relations characterized by addition. As Robert Van

Gulick puts it: "Instantiation [i.e., realization] of such a formal machine description requires roughly that there be some mapping from the formal states, inputs, and outputs of the abstract machine table onto physical states, inputs, and outputs of the instantiating system, such that under that mapping the relations of temporal sequence among those physical items are isomorphic to the relations of formal succession among the machine table items" [1988: 80]. The formal or mathematical relations are not themselves causal relations, so their realization cannot be in virtue of the physical system (or its parts) contributing the causal powers that are individuative of them. Abstract and mathematical functions and formal relations are simply not individuated causally, so they are not even candidates for realization on Gillett's view.

Among RP theorists, the observation that I am raising as a problem for Gillett's account or realization has long been recognized. For example, Robert Cummins writes:

> We may think of the button-pressing sequences as arguments to a function *g* that gives display states as values. An adding machine *satisfies g*; that is, the arguments and values of *g* are literally states of the physical system. Addition, as was remarked above, relates numbers, not physical states of some machine, so a physical system cannot literally satisfy the plus function. What an adding machine does is *instantiate* the plus function. It *instantiates* addition by *satisfying* the function *g* whose arguments and values represent arguments and values of the addition function, or in other words, have those arguments and values as interpretations. [1989: 89][12]

In Cummins' terminology, the dimensioned and flat views seem to be accounts of the *satisfaction* relation, rather than realization ("instantiation") itself. Realization of abstract algorithms and computations cannot be a matter or having the causal powers individuative of them for they are not causally individuated.

The crux of my objection to the causal approach of Gillett, Kim, and Shoemaker requires that what is realized be causally individuated. On those views abstract algorithms like addition cannot be realized because they are not causally individuated. But mathematical and computational realization are standard examples for RP, so the causal approach must be mistaken.

The objection that I am pursuing can be pressed using other standard examples of the realization from discussions of the special sciences, such as the realization of economic properties, à la Fodor [1974]. Fodor holds that monetary exchanges are realized by physical events: "Some monetary exchanges involve strings of wampum. Some involve dollar bills. And some involve signing one's name to a check" [1974: 124]. On the other hand, some signings of one's name are monetary exchanges and others are not. And the difference need not be in the causal powers of one or another token signature. Rather, economic properties and events such as monetary value and exchanges appear to be (at least in part) intentionally individuated. Consider:

(8) A dollar bill, four quarters, or ten dimes are among the realizers of one dollar.

A forged or "twin" dollar bill may have all the contemporaneous causal powers of a genuine one dollar bill, but it is not a dollar bill. The forgery does not have the property of being worth one dollar because it does not stand in the appropriate economic relations to other financial entities and properties, among which is the intentional

property of having been produced under authorization of the United States government. Of course one might hope (and many have) that the intentional relations that help constitute economic states and properties can themselves be physically realized. That may be the case. But the theory that economic or other social sciences properties are intentionally realized does not depend on any such "naturalization" of intentionality. Economic properties could be realized in intentional relations even if intentionality is an irreducible feature of Cartesian minds.

A similar point can be made using the example of functionalism concerning intentional states themselves. Suppose:

(9) Brain state #123 realizes (in me) the belief that coffee is delicious.

The idea is that for some state of my brain to realize a belief is for it to enter into numerous relations with other states of my brain as well as (through perceptual and motor mechanisms) the world. It has to play the "belief role" in my "cognitive economy." But what is the belief role? Well, one might think that a state doesn't count as a belief about coffee unless it was implicated in causing me to occasionally engage in coffee-related behaviors, and unless together (along with relevant desires) it justified or made rational (not merely caused) certain coffee directed actions. So beliefs are (at least in part) rationally or normatively individuated.[13] For example, if I believe that coffee is delicious and that the stuff in this cup is coffee then (let us suppose) I am justified in believing that the stuff in this cup is delicious, and (given certain desires) it will be rational for me to drink the stuff in this cup. In short, my brain state must be situated in a web of physical relations that are isomorphic to the semantic, logical, and normative relations among concepts. What my brain state need not do, in order to realize the belief

that coffee is delicious, is to actually stand in the logical, semantic, and normative relations that are (partly) individuative of beliefs.

The question at hand is not whether either (8) or (9) expresses a view we ought to adopt, but whether they are even coherent. According to the causal approach they are both nonstarters because the relations individuative of the realized states and properties are not causal relations. It seems that both the flat (Kim and Shoemaker) and dimensioned (Gillett) views may have to deny that computational, arithmetical, economic, intentional and semantic properties can be physically realized. Yet these are classic examples of realized states and properties. The causal realization theorist might hold out hope that computational, arithmetical, economic, intentional, semantic, and normative relations turn out simply to be causal relations. Maybe they will. But it is at least open to some theorist to hold that to realize a belief is to occupy a role in a network of irreducibly intentional or semantic relations, or that to realize monetary value is to have a role in a network of autonomous economic relations. If that is right, then the realization relation itself does not require that realized states and properties be causally individuated.[14]

The causal approach, I maintain, construes realization too narrowly and therefore yields an impoverished account of the realization relation. The relations included in the flat and dimensioned views could be among the realization relations. But it is wrong to think of them as the only realization relations, or as characterizing realization generally. With respect to my examples of "realization" above, the flat and dimensioned views give the right answers to (1) and (2), and also (4). ((3), remember, is negotiable.) But they fall short on (5), as well as (8) and (9). We cannot ignore case (5). If

(5) is not realization, then (6) is not what RP asserts; for (6) intends to claim that the brain-mind relation is the same as the machine-program relation, and that relation is realization. If (6) is the claim of RP, then (5) must be realization as well.

I am not claiming that mental states or properties are not or ought not be causally individuated. It is true that most current versions of RP appeal to causal criteria of individuation, and there are good reasons for thinking that mental states and properties should be causally characterized.[15] It may even be that the relation described by Gillett (or Kim, or Shoemaker) is the relation between minds and brains. Yet I maintain that the realization relation itself should not be restricted to realization of causal states and properties. Our task is not to make RP about the mind true. Our task is to understand the realization relation so that we will be a in position to evaluate RP as a general thesis. The question of whether brains realize minds just is the question of whether the relationship between brains and minds is the same as the relation between machines and programs or algorithms. Computational realization (implementation) is supposed to be the uncontroversial case, to which psychological realization is compared. If any theory claims that realization is not the relation between machines and algorithms, then that theory has the cart before the horse.

The difficulty for the causal approach to realization is plain. At this point, several objections may seem attractive, so we had best get them out of the way before going forward. First, one might object that this old idea of realizing abstract or computational functions is passé and therefore can be justifiably ignored. If nobody holds this view of realization anymore then perhaps it does not have to be accommodated by a proper account of realization. For better or for worse, however, the computationally inspired

view of realization seems to be commonly held. Consider, for example, David Marr's theory of vision. Marr's theory invokes a notion of realization that accommodates the realization of algorithms. For example, Marr writes: "In order that a [computational] process shall actually run, however, one has to *realize* it in some way and therefore choose a representation for the entities that the process manipulates.… The choice, then, may depend on the type of hardware or machinery in which the algorithm is to be *embodied* physically. This brings us to the third level, that of the device in which the process is to be *realized* physically" [Marr 1982: 23-24, emphasis added]. Marr's theory raises some questions about realization that are beyond the scope of the current discussion.[16] For the moment the salient feature is that Marr assumes that algorithms can be embodied or realized by physical systems. These algorithms will often be abstract or mathematical, not themselves formulated in causal terms. For example, Marr and Hildreth [1980] propose that the algorithm for accomplishing visual edge detection is the Laplacian of the Gaussian ($\nabla^2 G$). The Laplacian is a second order derivative; its "inputs" and "outputs" are values or numbers. Marr and Hildreth argue that the Laplacian of the Gaussian is implemented or realized by retinal ganglion and X-cells in human beings; but it could be realized by other hardware or wetware. For the algorithm be realized in any physical system a representational schema must be employed, as described by Cummins for the case of addition. This is precisely because the mathematical function is not itself individuated in terms of causal powers, so there are no distinctive powers for a realizer to have.

These days the details of Marr's theory of vision are not widely accepted, so one might think that we can do away with its commitments to realization of algorithms. But

even if the particulars of Marr's theory are dismissed, his general theoretical framework and his ideas about realization have been widely influential. Terence Horgan and John Tienson take it that Marr's account represents a standard view of realization, writing, "[t]he relationship between state types at Marr's middle (algorithm) level and state types at the lowest (implementation) level, and *also* the relationship between state types at Marr's top (cognitive-transition) level and state types at his middle level, is the relation that philosophers call *realization* and cognitive scientists call *implementation*" [1996: 23]. Indeed, Marr's work is often cited as evidence of the convergence of RP in philosophy of mind and empirical cognitive psychology—where this convergence is taken as evidence of their joint success. So even if it turns out that no one currently advocates a theory of minds that relies on realization of abstract computations, it seems that the possibility of such realization is nevertheless part of the basic understanding of the realization relation itself.[17]

This brings us to a second objection. Perhaps it is not just theories of cognition that rely on abstract computational functions that have been rejected, but also the very idea of abstract realization itself. After all, recent work on realization and RP—specifically, development of the flat view in Kim's and Shoemaker's metaphysics of mind—relies on realization of causally individuated states or properties, not abstract or computational states or properties. Perhaps we simply learned that the realization relation involves causal relations, and therefore that we need not accommodate realization of abstract computational algorithms?

For better or worse, it is not the case that the computational view has been abandoned. Ned Block, for example, writes, "According to cognitive science, the

essence of the mental is computational, and any computational state is 'multiply realizable' by physiological or electronic states that are not identical with one another" [1990: 146]. Yet it is true that among metaphysically inclined RP theorists about the mind attention has shifted to realization of causally individuated states. This, however, is a change in the account of *what is realized*, not in the realization relation. The realization relation itself is usually taken to have the same characteristics as realization of abstract machine states. As Putnam notes, "[w]hen it became clear that the formal properties of [computational] states are quite unlike the formal properties of psychological states, the original idea of functionalism quickly was replaced by an appeal to the notion of an ideal 'psychological theory.' But this ideal psychological theory was conceived of as having just the properties that formalisms for computation theory possess" [1999: 34]. As I interpret him, Putnam understands newer versions of functionalism that are specified in terms of the empirical causal laws of psychology—what Block [1980] calls Psychofunctionalism—as elaborations of the basic idea of computational functionalism. In Psychofunctionalism, laws of psychology replace formal program rules. In another variation, sometimes called Functionalism [Block 1980], the analytic truths or "platitudes" of folk psychology are taken to implicitly define the psychological "program." In either version, the psychological program is said to be realized in the brain just as Putnam earlier imagined that computational programs would be. This is the key feature that Psychofunctionalism and Functionalism share with the earlier computational versions. The program has changed, but realization is just as it always was.

It seems that we cannot dismiss abstract realization out of hand. The realization of computing machine states may have been abandoned as a theory of minds, but that does not mean that machines do not realize computational algorithms. After all, the computational analogy remains the dominant metaphor for RP. In short, we must be careful to distinguish between the realization relation and its applications. Perhaps philosophers have come to see that realization of causal (rather than abstract) properties or states provides a better version of RP, e.g., about the mind. But it is one thing to learn that realization of computational kinds does not make for a good ontology, and quite another to learn that realization is not the relation between machines and algorithms after all. Realization physicalists need not hold the view that mental states are abstract machine states; but they should not deny that abstract states can be realized by physical systems.

This brings us to a third response to my objection. It could be suspected that I am trading on an ambiguity in notions like *computation*, *program*, or *algorithm*. These may pick out abstract relations, but they may also pick out physical operations. We can concede that the mathematical function of *addition* is not causally individuated, and thus cannot be realized according to the causal approaches. But *adding* is a procedure or operation (an "algorithm" in the causal sense) and the property or state of *being an adder* (i.e., an *adding machine*) can be causally individuated. Even if addition cannot be realized, adding and adders can be realized. And that is all that RP requires.[18]

The first thing we must remind ourselves is that that mathematical and computational states are the main but not the only examples of realized states and properties that are not causally individuated. Economic, intentional, semantic, moral,

and epistemic states and properties are also sometimes said to be realized. And these concepts do not look to suffer from such ambiguities as *adder* may.

The second part of the response is to express perplexity about why some procedure would count as *adding* or some device would count as an *adder* unless they stood in some relation to the mathematical *addition* function. Of course realizing an *adder* involves having some causal powers. It does not, however, involve having the causal powers distinctive of *addition*, for this objector concedes that there addition is not causally individuated. But for the same reason there is no distinctive set of causal powers that individuate adders. What makes something an adder (an adding to machine) is that it has causal powers that stand in some relation to the mathematical *addition* function. The relation between physical devices or systems and mathematical functions is usually thought of as isomorphism, mapping, or representation.[19] So, following Van Gulick, Cummins, and Marr, we may say that when the physical states and state transitions of a device map or represent the formal states and transitions of the addition function, then the device realizes the property or state of being an adding machine. The crucial point is that although an adding machine is realized by something's having particular causal powers, it is not realized—as Gillett, Kim, and Shoemaker would require—by anything's having the causal powers individuative of adding machines.

For a fourth and final response to my objection, one might bite the bullet and assert that so-called realization of abstract computing machines or algorithms simply is not genuine realization after all. (If Putnam, Fodor, Marr, or anyone else thought otherwise, so much the worse for them!) After all, we are not trying to give an old-

fashioned philosophical analysis of the term 'realization' or the concept *realization*. 'Realization' is a bit of philosophical jargon. Why should we be beholden to historical versions of realization? In fact I think that this is Gillett's position.[20] And I suppose that if Gillett wants to stipulate the use of 'realization' in this way there is little to stop him. But I urge that we not accept this stipulation.

One reason, already mentioned, is that the analogy of minds and computing machines was a significant motivation for developing RP in the first place. The fundamental idea of RP is that the relationship between brains and mental states is the same as the relationship between physical devices and computational states. Gillett will have trouble making any sense of the original mind-software analogy or how it leads to current versions of the theory, because on his view the brain-mind relation is not the same as the machine-program relation. If the trouble were limited to a broken metaphor that would be of little concern. But Putnam's radical proposal was to take the metaphor as literal truth—to convert the computational analogy into a computational theory of mind. The theory is that the relationship between mind and brain is the same as (not merely analogous to) the relationship between software and hardware. Putnam calls that relation *realization*. But if the causal theorists are correct, then Putnam misused realization when he coined it, for realization (as they understand it) could never be the relationship between an abstract algorithm and its implementing hardware.

Rather than say that Putnam was mistaken about his own theory, we could decide that even the technical term 'realization' is ambiguous, so that Putnam was talking about a different relation than Gillett and the causal realization theorists. This might be conciliatory but it is not very plausible. Kim and Shoemaker (to whom Gillett

attributes the flat view) and Block and Fodor (to whom Gillett attributes the dimensioned view) are clearly intending to explicate functionalism or its descendants. They are discussing the realization relation that is at stake in RP, the relation that was introduced by Putnam. So it will be of no use to stipulate that we should use 'realization' in a new, different, more general, or more basic way.

A third reason to resist Gillett's stipulation is that it seems to commit the error of confusing the realization relation with its applications. It may be that realization of abstract states or properties does not support metaphysical theories of mental states—the cases of ultimate interest to Kim, Shoemaker, and Gillett. That is a good reason for rejecting the abstract computational theory of minds. But we should not absorb the demands of one application of realization into the relation itself. Although the RP theory of minds is the marquee case, the realization relation is invoked by theories and explanations that do not implicate metaphysical claims and their idiosyncrasies.[21] As noted earlier, realization sometimes figures in theories of intentional and semantic content, in explanations of computing machines, and is central to some views of the special sciences generally.

I have thus far been emphasizing what is left out when realization is construed only causally, and bolstering my critique against various replies. I have been using Gillett's view as the representative of the causal approach. If I am right the difficulties about the realization of computational states and properties are faced equally by the accounts of Kim and Shoemaker. But Gillett's dimensioned view also confronts some problems that the Kim's and Shoemaker's flat views do not. And to these we now turn.

4. The Dimensioned View and Non-Reductive Physicalism

The case that motivates Gillett's reasoning about realization is the hardness of a diamond, example (7). "The sciences," Gillett writes, "have given us a very precise and detailed account of how the hardness [H] of the diamond results from the properties/relations of the individual carbon atoms" [2002: 319]. Hardness is instantiated in the diamond but not instantiated in the individual atoms. And hardness contributes different powers to the diamond (e.g., glass cutting) than the properties of atoms contribute to them. "Given these differences, H cannot be identical to any of the particular properties/relations of the carbon atoms.… the properties/relations of the carbon atoms apparently 'play the causal role' of H, but not vice versa, and, consequently, it is plausible that H is *realized* by the relations/properties of the carbon atoms" [2002: 319].

Gillett's idea is that the hardness of the diamond is realized in it because of the properties of its parts, the carbon atoms. Hence, and contrary to the flat view of Kim and Shoemaker, realized and realizing properties may occur in different objects.[22] The flat view cannot explain the realization of the properties of an object by the properties of its constituents, for it assumes that realized and realizer properties are instantiated in the same individual. The dimensioned view remedies this alleged failing. But in so doing the dimensioned view draws realization too near to material composition in general. Gillett's view is that when the carbon atoms *compose* the diamond and the properties of the carbon atoms *realize* the properties of the diamond. For Gillett, realization and composition go hand in hand.

The trouble with Gillett's proposal is that it runs counter to the chief motivation for RP, which is to provide an alternative to the view that mental states are identical to or composed of brain states. The claim that mental states are realized by brains is supposed to contrast with the view that mental states are like the other macro-states of brains. This is not the case on Gillett's account. Brain-appropriate properties are realized in brains just as hardness is realized in diamonds. If this were correct, the claim of RP that mental properties are realized in brains would not distinguish mental properties from other macro-properties of brains. They are all of a kind: they are realized properties.

Early proponents of the mind-brain type-identity theory proposed that mental states be identified with "physico-chemical" states of brains. The reason for this emphasis was to distinguish the identity theory from the hypothesis that mental states are identical to psychological states of brains, which is uninformative in its plain reading and invokes dualistic "psychical" states and properties in its panpsychist or double-aspect reading. The point, for the identity theorist, is that mental states and properties should be identified with the kinds of states and properties that brains have in virtue of being physical objects of a certain sort rather than with states and properties that have been superadded to brains. These will be, roughly speaking, the intrinsic or structural states and properties of brains qua objects of the natural sciences.

Versions of functionalism deny that psychological states are intrinsic or structural states of brains. Putnam's alternative hypothesis is that mental states are functional or relational states of brains. Against the identity theory, Putnam suggests that brain states (thought of structurally or mereologically) realize mental states

(thought of functionally.) These are the kinds of states that could be realized by brains, other physical systems, or (in principle) non-physical systems. Realized states and structural states have different characteristics, and those of realized states make them especially good candidates to be the objects of the special sciences. (For example, whereas RP is usually said to be a non-reductive theory, mereological accounts of minds are typically counted as reductive.) But on Gillett's account, structural states are also realized. So they will have the same characteristics as other realized states. In short, Gillett's dimensioned view undercuts RP by assimilating ordinary macroscopic states and properties into realized states and properties.

Perhaps it is useful to put this point another way: A major benefit of RP is supposed to be that it makes special sciences entities and properties highly *realization independent*. According to Putnam it does not matter whether minds are "copper, soul, or cheese" [1975]. It does not matter what kind of stuff realizes the word processing program on my computer, so long as that stuff has the proper "formal" or "abstract" properties and relations. Yet Gillett's account makes the stuff matter a great deal. Indeed, on his account it is doubtful that any computational state or property is realized in common when Microsoft Word runs on my Macintosh and on some PC. This is because these items are, according to RP, causally heterogeneous. But on Gillett's account, in order for two machines to realize the same program they must be causally homogenous. This is entirely contrary to the motive of RP, which was to legitimize explanatory kinds that are not causally homogenous—as Fodor [1974] is at pains to argue.[23] With friends like Gillett, RP doesn't need enemies.

RP may or may not be a good theory of mind. But we'll never be able to properly evaluate it if we cannot distinguish it from its competitors, and this means we should understand realization as distinct from composition. When Gillett offers an account of realization that covers also cases of material composition—when he argues that structural properties are also realized properties [2002: 319, fn.4]—he obliterates the very distinction on which RP depends. As Van Gulick puts the concern: "If physiological properties and other physical properties of many sorts can all be interpreted as functional properties, then the functional nature of psychological states cannot be taken as evidence against the identity thesis. Moreover, the very thesis of functionalism itself is in danger of losing interest, for its appeal lay in picking out some supposedly distinctive characteristic of psychological properties" [1983: 190]. Van Gulick's point, I take it, is that the theory that identifies mental states and properties with macro states and properties of brains is the identity theory. If all brain states and properties are realized, then the contrast between identity theories and RP gets no grip. The point is quite general. If physiological properties and other physical properties of many sorts can all be interpreted as realized properties, then RP is not a distinct or distinctive doctrine after all. It seems, therefore, that friends of RP should not invoke Gillett's account of realization.

Notice that with respect to the present concern, Kim's and Shoemaker's flat views, whatever else their flaws, fare rather better. This is precisely because the flat views exclude some cases—namely, the "dimensioned" cases that Gillett includes—and thereby distinguishes realization from other constituting relations. The flat theorist can discriminate cases wherein an individual has a property in virtue of the functional or

relational role that it plays from those in which an individual has a property solely in virtue of the causal powers of its parts. The former is a case of realization; the latter is a case of ordinary material composition. RP holds that mental states and properties are of the former sort, they are realized by brain states and properties. The identity theory holds that mental states and properties are of the latter sort, they are materially composed.

Where does this leave us? In the first section I argued that, along with the flat view and other exclusively causal realization theories, the Gillett's dimensioned view wrongly excludes realization of computational, economic, semantic, and intentional states and properties. Yet these are just the sort of states and properties that RP holds are realized by physical causal systems. In this section I argued that the dimensioned view compounds its difficulties by collapsing the distinction between functionalism and the identity theory. Stipulating that the relation among properties of an individual and the properties of its parts is realization, Gillett argues that functional and structural properties are both realized. But this conclusion runs against RP. The accumulating problems for the dimensioned view suggest that it is time to go back to the drawing board and formulate a new account of realization.

5. Realization and Roles

Gillett protests that advocates of the flat view of realization take the idea of playing or occupying a role too literally [2002: 321, 2003: 593]. They do, indeed, take it literally. William Lycan declares, "Functionalism is the only positive doctrine in all of philosophy that I am prepared (if not licensed) to kill for. And I see the

'role'/'occupant' distinction (some say obsessively) as fundamental to metaphysics" [1987: 37]. The role/occupant idea is a fundamental idea of RP. The history of functionalism (and RP, if they are not simply the same) is the story of a search to figure out how best to characterize the mental role or roles. Early attempts used abstract computational specifications of the role, later attempts have employed one or another causal specification, and most recently evolutionary characterizations have become fashionable [Polger 2004a]. What all of these cases have in common is that mental roles—be they computational roles, causal roles, or teleological roles—are specified relationally or, as it is common to say, functionally. This simple observation tells us that we should think about realization in terms of roles. To occupy a role is to have the relations that are distinctive of the role. In short, to realize a property or state is to have a function:[24]

(R) Property/state instance *P* realizes property/state instance *G* iff *P* has the function $F_G(x)$.

The basic idea is that *P* realizes *G* if and only if *P* has the *G*-function, if it plays the *G*-role. For example, something realizes the property or state of *being a heart* if and only if it has the function of *pumping blood*. A realization relation is a function-conferring relation.[25]

There are many different kinds of functions; for this reason the general formulation (R) does not restrict the notion of function, but allows that various functions could be employed. Some functions are individuated in terms of causal

powers (as the flat and dimensioned views require) but others are not. In formulations of RP it is often obscure what sort of function or functions are employed. Yet fixing the notion of function is crucial to understanding a version of RP. What it takes to have a function, or to function in a certain way, depends on the kind of function. As there are many kinds of functions, there are correspondingly many different realization relations. We could make this explicit by reformulating (R) using the function $F n_G(x)$, where n is the variety of function.

Realizing a causal function is one thing. Plausibly, having a causal function requires instantiating certain causal powers or dispositions, or being apt to do so [e.g., Cummins 1975].[26] The individuating function for the role occupied by carburetors is probably causal. Above I argued that the individuating function for being a state of Microsoft Word or being an adding machine is not causal. If Ruth Millikan [1984, 1989] is right, the individuating functions for hearts and beliefs are etiological. Realizing an etiological property or state does not require any particular causal powers; instead it requires having a certain history [e.g., Millikan 1989, Neander 1991].[27] This is why realizing an etiological property does not fit either the flat or dimensioned views in any obvious way, for a history does not contributed any causal powers to the realized property.

My proposal is captured by (R). Realization is understood as the having of a function. The realization relation, then, is whatever (often complex) relation it takes to have a kind of function. The account is deliberately open-ended about what relations count as functional relations. Each functional kind carries its own realization relation. They are unified, insofar as they are unified at all, by being functional relations. If all

functions have something in common, then that is what unifies realization relations. If, as I suspect, functions are a heterogeneous lot, then so too are the realization relations. For this reason I prefer to think of my account as denying that there is a single unified realization relation.

There are a number of advantages to my functional account of realization. My account has the resources to distinguish RP theories like functionalism from non-RP theories like the identity theory. Although my account may make functions abundant, once we fix the variety of function employed in a particular version of RP it is then a substantial question to ask when or whether two things or properties stand in the given functional relation to one another. It may be true that every thing realizes every finite state machine. But it is certainly not true that every thing realizes even one etiological function.

Obviously it is important that my approach to realization accommodates the realization of algorithms by machines. Having a computational function, understood as a formal algorithm, does not appear to be a matter of having specific causal powers; rather it seems to involve having some causal powers or other that are isomorphic with or map onto abstract or semantic relations. This is how Cummins and Van Gulick, describe the realization relation. The abstract versions of realization may or may not form the basis of an RP account of the nature of minds. They may nevertheless be part of an account of computing devices, or have practical value in psychological explanation [Marr 1980, Jackson and Pettit 1990]. On the other hand, it may turn out that they are not very useful at all. If so, these faults do not have to do with computational realization failing to be a case of realization. Take for example the worry

mentioned above, that on an abstract computational theory every rock counts as a realization of every finite state machine [Putnam 1988; Searle 1990; Chalmers 1996; Scheutz 1998, 1999].[28] If this is correct and we do not think that every system also counts as having some mental states, then we should not adopt the abstract computational theory of minds. We may, furthermore, decide that any such account is too weak to be part of an interesting or useful theory of computation. But, far from being a reason to deny that abstract functions are realized, the objection assumes that abstract realization is coherent. This example is supposed to be a *reductio ad absurdum* of a certain theory of mind, not of the account of realization.

From my perspective, the "flat" view of realization that Gillett attributes to Kim and Shoemaker is a special case of (R) that employs causal functions. On that variation, having the causal function $F_G(x)$ is a matter of having the causal powers individuative of *G*. In order to accommodate Gillett's "dimensioned" view, we need only revise (R) to allow that either *P* or the properties of *P*'s bearer's parts can have the causal function.[29] Of course this is an *ad hoc* maneuver. A better solution would be to discover a notion of causal function according to which *P* has function $F_G(x)$ when *P*'s bearer's parts stand in relations O_1-O_n. No doubt there are many, and if so they can be used to formulate instances of (R). My account is flexible in that any notion of function can be employed, and it will bring along its own realization relation. But the account is also substantive. If notions of function are to do any work (e.g., in a theory of mind, economics, truth, or ethics) they will have to be defended.

Of course we are not at present defending RP, we are merely exploring how to formulate it. Advancing a version of RP is another matter. That is a hard task, and its

success will no doubt hinge on which version is adopted, on making a careful choice of which functions to employ. At this stage it is important to leave those options open. My account makes it easy to formulate RP's overarching thesis, as well as specific versions. But it also makes it a substantial job to defend the merits of particular versions, to show that they accomplish all that RP hopes.

The failures of some applications of (R) are not failures of an account of realization relation itself. This is important. If RP is a serious thesis, then 'realization' is not just the name for whatever relation actually holds between minds and brains. Unfortunately, advocates of RP never have said much about the nature of the realization relation. But they do give examples. I claim that realization must be understood in terms of these prototypical examples: computer programs, adding machines, carburetors, hearts, and the like. If we want to understand RP then we must take these examples seriously.

The proposed approach to realization is more flexible than either the flat or dimensioned view. Recall the cases with which we began:

(4) The role of the air-fuel mixer is realized by a fuel injector in my car.

(5) My computer sometimes realizes Microsoft Word.

(6) Memory fixation is realized in humans by long term potentiation of neurons.

My account of realization supports example (4), just like the causal views. But it also accommodates example (5), which the causal views do not. I have my doubts about whether or not (6) is part of a good theory of minds [Polger 2004a]. But on my account

realization is at least a candidate for being the brain-mind relation. And that is as it should be.

Ultimately I suppose that the merit of any account of realization will be judged in its theoretical fruitfulness. So it is an advantage that treating realization functionally allows us to make sense of and evaluate novel applications of the realization relation. Two examples are Frank Jackson's and Philip Pettit's moral functionalism [1995, 1996; Jackson 1998] and Michael Lynch's alethic functionalism [2001, 2004]. These theorists explicitly take themselves to be extending the realization theory that originated with functionalist theories of mind, and I see no reason to deny the connection.

Finally, the functional approach to realization also provides a useful perspective on the debates within philosophy of mind and metaphysics concerning multiple realizability. To these we now turn.

6. Multiple Realization Reconsidered

Gillett is motivated in part by the goal of making sense of the doctrine of multiple realization (MR). Roughly speaking, MR is the idea that special sciences entities and properties can be realized by physically heterogeneous realizers.[30] Noting that few if any specific accounts of the realization relation have been suggested, Gillett observes, "it appears plausible that providing such an account of *realization* would be a necessary step in any precise understanding of *multiple* realization" [2002: 592]. Offering the dimensioned view, he argues that Lawrence Shapiro's [2000] critique of multiple realization arguments illicitly relies on the flat view of realization. According to Gillett, the dimensioned view (a) can avoid Shapiro's critique of MR and (b) should be

attributed to Shapiro's targets, Ned Block and Jerry Fodor [1972]. In this way Gillett defends Block and Fodor by representing the debate over multiple realization as tacitly depending on substantive accounts of realization. "The question of the correct account of realization is thus inextricably bound-up with the issue of the nature of MR," he concludes [2003: 600]. We are now in a position to see what is wrong with Gillett's approach, and why the functional account of realization should be preferred.

Consider an example that is contested by Shapiro and Gillett: two corkscrews that differ only in that one is made of aluminum and the other of steel. Shapiro argues that the two corkscrews should not count as different realizations of one kind of thing. His reasoning is that because the difference between being made of steel and being made of aluminum is not relevant to individuating the objects as corkscrews, this is not a case of MR. This is a case of distinct corkscrews, but not of distinct realizations. Shapiro argues that Block and Fodor [1972], in contrast, would count the two corkscrews as distinct realizations. According to Gillett, in so arguing Shapiro tacitly invokes the flat view of realization because he focuses our attention on the properties of the corkscrews rather than the properties of their parts. The causal powers individuative of corkscrews are the same in each case; but in one case those powers are contributed by steel and in the other by aluminum. If we adopt the dimensioned view, we notice that the parts have different properties and thus conclude (with Block and Fodor) that the two corkscrews present a genuine case of MR after all.

As Gillett sees the disagreement, Shapiro assumes that the relevant differences between realizations must occur at the level of realized properties. Thus two corkscrews that both have the same corkscrew-individuative causal powers would not count as

multiply realized. This neglects the possibility that the relevant differences between realizers may occur in their constituents; the constituents of the corkscrews are relevant if their corkscrew-individuative causal powers have physically different sources. (But, Shapiro would say, the idea of RP is that realized properties are independent of and "screened off" from the properties that realize them.)

How do we adjudicate between these two different judgments regarding which properties are relevant to MR? Gillett's suggestion seems to be that the question is resolved once we fix the realization relation. On my view, this is another example of putting the cart before the horse. We must fix our MR judgments and then formulate an appropriate account of realization.

It may seem odd to settle on examples of multiple realization without first understanding realization. (Shouldn't we understand individual cases before we delve into multiplicity?) I, too, once thought that we would need an account of realization to make sense of multiple realization. But in fact the two ideas are rather different. Realization is a metaphysical relation. Multiple realization is a claim about the diversity of things that can have common states or properties, i.e., about the generality of certain explanatory kinds. The latter does not carry substantial commitments about the former.

On closer examination we see that this must be the case if the MR arguments for RP are going to work. RP is a substantial metaphysical claim about the nature of psychological and other special sciences properties—that they are realized by but not identical to "physico-chemical" properties. The primary argument for RP employs the claim that mental states are MR in order to block identity theories. But if MR itself embodied a substantial commitment to the nature of realization, then it would be

question-begging to invoke it in an argument for the claim that psychological properties are realized. For MR to do its job it must not be a substantial metaphysical thesis. Instead, I have argued, MR is properly understood as a thesis about the generality of explanations and explanatory kinds [Polger 2002, 2004a]. If so, then Shapiro and others concerned with MR have been, after all, correct that they need not provide specific accounts of realization. What Shapiro offers is the hunch that corkscrew explanations will not treat steel and aluminum corkscrews as different enough to count as MR. In contrast, if we have a waiter's corkscrew in the one hand and a lever-style corkscrew in the other, these two will count as MR of the kind *corkscrew*. This is because although they both open bottles, they open bottles in different ways. They are treated as different from the point of view of a hypothetical special science of corkscrews. (Perhaps they figure in different generalizations, or some such.) Of course for some purposes one might care about the difference between steel and aluminum corkscrews, for example if I need something to pry a part off my car, or to hold open a heavy door. Shapiro is betting that these functions are not part of corkscrew science, which restricts itself to the behavior of corkscrews under normal bottle-opening conditions. Corkscrew science, as imagined by Shapiro, is typical of the special sciences in that it is relatively autonomous from lower-level sciences, and thus from lower-level facts including many or most facts about the composition of corkscrews.

Return now to more familiar examples of multiple realization:

(10) Pain is realized in human beings, dogs, and octopi.

(11) Hearts are realized in humans and worms.

These examples are compatible with the usual philosophical uses of 'realization'. But (10) and (11) do not presuppose any metaphysical account of realization. Different kinds of creatures appear to experience pain. The class of things that experience pains cross-cuts some biological classes, such as species. In this sense, pain is multiply realizable—it is realized in biologically distinct kinds of things. But whether it is realized differently in various kinds of things is another matter. Increasingly there is evidence that all pain sensory systems share common features. If so, pain experiences may be identified with states or processes in any such system. Indeed there is good reason to think that this style of explanation is plausible [Polger and Sufka, 2006]. If pain can be identified with a brain process that is common across pain-experiencing species, then it is not realized in the RP sense.[31] This raises the paradoxical possibility that some property may be multiply realizable (because it can occur in different kinds of systems) but not realized (because it is type identical to a property that is common to all the systems that instantiate it.) This situation is not so strange. Remember: realization is a metaphysical relation, and multiple realizability is a thesis about the generality of explanatory kinds. If you want to know whether hearts are multiply realizable, you will find this out not by consulting an account of realization but rather by seeing what biologists are prepared to count as hearts. Next we find out what if anything hearts have in common.[32] Only then we can determine what relationship being a heart bears to being a physical object of a certain kind—whether that relation is identity or realization. Gillett should agree if he is serious about accounting for a relation that is tacitly used by scientists, be they chemists or psychologists.

I am not suggesting that scientists will tell us what the nature of the realization relation is. Statements of MR, like (10) and (11), do not make claims about ontological dependence relations, they make claims about the generality of explanatory kinds. They tell us what some systems will do in certain circumstances (viz., they behave like they have pains and hearts, respectively) but they do not tell us about the natures of the metaphysical status of pains or hearts.[33] The assertion that the property or kind *eye* is multiply realizable is not a claim about the metaphysical relations of the kind eye to physical kinds, it is a claim about whether the kind *eye* is restricted to mammals or can be applied more widely. So too with *belief* and *pain*. Having decided the extension of *eye* or *belief* or *pain*, we may then inquire whether that property or kind is coextensive with some physical and physico-chemical kind. If eyes are coextensive with some physical kind L, then perhaps eyes are identical to L's. But if eyes are not coextensive with any physical kinds, so the MR argument goes, then eyes are not identical to a physical kind.[34] If so, eyes must be physically realized. In that case an account of realization should be able to explain the realization of eyes. This is the usual path to RP, whether about eyes or economic properties or beliefs. The crucial point is that, contra Gillett, we determine the extent of multiple realization prior to formulating an account of realization.

The special sciences are usually understood to be realization independent. As I have noted, Putnam is at pains to point out that functionalism is compatible even with dualism. From the point of view of functionalist psychology it does not matter what minds are made of. In contrast, on Gillett's account the stuff does matter. The hardness of a diamond is realized not by the diamond's properties but by the properties of the

stuff of which the diamond is composed, viz., the carbon atoms.[35] As a result, Gillett's dimensioned view is at odds with RP's commitment to the autonomy of the special sciences. Gillett is lead to this view in part because he mistakenly builds into the thesis of MR a substantial commitment to an account of realization. But that is exactly what MR cannot be if it is to be a premise in any non-question-begging argument for RP.

It is true that realization and multiple realization are inextricably bound up. But the dependence runs the opposite direction that Gillett supposes. We need to know about multiple realization in order to formulate a good account of realization, rather than vice versa.

7. Conclusion

On the account of realization that I have advanced, there is no single and substantive realization relation. Realization is a matter of having a function, so there are as many realization relations as there are varieties of function. In this sense realization relations are easy to come by. Each notion of function dictates the relation or relations that count as its realization relations. Realization relations are cheap but they are not free. It is further question whether any of the realization relations can do the work that they are called upon to do for RP. I have argued that no available notion of realization can do everything that functionalists about the mind demand [2004a]. I do not know whether RP makes a good account of things economic, computational, sociological, moral, or alethic.

8. Notes

* I have had the pleasure of discussing this paper and its ancestors with many people. I am especially grateful to Tom Bontly, José Luis Bermúdez, Carl Craver, Eric Funkhouser, Carl Gillett, John Heil, Jaegwon Kim, Michael Lynch, Brendan O'Sullivan, Larry Shapiro, John Symons, Gene Witmer, and Chase Wrenn. I would also like to thank the anonymous referees for this journal, and audiences for versions of this paper that were presented at the American Philosophical Association, Southern Society for Philosophy and Psychology, the Society for Philosophy and Psychology, the University of Cincinnati, and Washington University in St. Louis. Work on this paper was supported in part by the Charles P. Taft Research Center at the University of Cincinnati.

[1] I am aware that there have been, over the years, many discussions of realization in the context of explicating one or another particular theory [e.g., Field 1978, Heil 1992, Tye 1994, Chalmers 1994, 1996]. However most of these discussions, upon examination, either say nothing at all about realization itself (instead focusing on the question of whether, given system A is a realization or realizer of F, a similar system B is also an F realization or realizer), focus on particular cases only, or else explain realization by analogy with some case that is assumed to be uncontroversial—usually computer programs being realized by machines.

[2] Of course if RP is correct than it may be that Sally (or some part of her) realizes the belief that it is raining. But that is not what is asserted by (1) on its face. I suppose one might argue that (1) is an elliptical formulation of (1*) Sally realized (or has realized in her) the belief that it will rain, so she brought her umbrella. But (1) seems to be saying something about the content of Sally's belief and its rational connection with her

behavior, rather than about the metaphysical status of her belief. For example, we would still say (1), but not (1*), if beliefs are identical to brain states rather than realized by them, or if beliefs were Cartesian objects.

[3] There are two specific differences between "realizing" one's potential and realization according to RP. First, "realizing" one's potential is a matter of causing ones potential to obtain, but RP's realization relation is usually presumed to be a non-causal dependency. A mother does not realize her child, though she may be said to cause the child to come into being. Second, realization is usually taken to involve relational or extrinsic rather than internal or intrinsic changes. For example, the relationship between a living tree and a chair (which is made of wooden parts, crafted from the tree's wood) is not usually described as realization. It may be said that the parts of the chair "realize" the chair, though usually we say that the parts compose the chair. Even in that sense it seems odd to say that the tree "realizes" the chair. After all, the tree does not compose the chair, even if the chair is composed of (part of) the material which composes or composed the tree. Compare: this large bean bag realizes a chair, or this tree (fallen in the forest) realizes a bench. The beanbag and the fallen tree do seem to be apt to realize chairs or benches in the way that RP understands realization. This kind of realization involves a relation sometimes described as "playing a role." The beanbag and the fallen tree play the role of providing seating. The standing tree does not. We might change the tree (by cutting it down) and subsequently it could realize a bench. Or (by sawing it, milling it, and so forth) we could manufacture a chair from some of the material from the tree. These changes are sufficiently radical that we might seriously question whether the tree is still in existence. If not, then it is certain that the tree does not realize the chair in the

RP way, for realization is typically supposed to be contemporaneous and nondestructive. Yet even if the tree survives it would be odd to claim that the tree itself realizes a chair. Maybe chairs are realized by trees, but this would at least be a contentious case of realization.

[4] I am grateful to Eric Funkhouser for pressing me to make these features explicit.

[5] For the moment I'll stick with the weaker claim that an account of realization need not endorse examples like (1) and (2). A stronger claim would be that there are some cases of "realization" that should not be counted by an account of realization that is suitable for RP. Indeed I'll argue just that in Section 4.

[6] I mention this in part because "instantiate" is sometimes used as a synonym for "realize" among advocates of RP; also, "implement" and "occupy [the role of]" seem to be used interchangeably with "realize" by many philosophers. For example, Kim writes, "We are constantly reminded that any mental state, say pain, is capable of 'realization,' 'instantiation,' or 'implementation' in widely diverse neurological structures" [1992: 1].

Because Gillett requires that realization be a relation between property instances, his account is prima facie at odds with talk about the realization of states (e.g., pain), events (e.g., edge detection), processes (e.g., adding), or objects (e.g., carburetors.) This is not a problem with Gillett's view on which I will dwell. I am prepared to allow that the difference is only superficial and could be finessed; although I doubt that Gillett himself would welcome the finessing. But it is worth noting that Kim, for example, talks about the realization of "states" and "structures." So it is not obvious that either he or

Shoemaker mean to limit realization to property instances explicitly. Gillett's view may depart from the Kim and Shoemaker views more than he lets on.

[7] Gillett's formulation is less than perspicuous. It can be streamlined as follows:

> Property/relation instance(s) F1-Fn realize an instance of a property G, in an individual *s*, *if and only if* F1-Fn are properties of *s* or *s*'s constituent(s) and F1-Fn contribute the powers that are individuative of an instance of G to *s*, but not vice versa.

This revision eliminates the "in virtue of" locution from Gillett's formulation. Even with this adjustment, it is hard to see how the account explains realization. We can readily understand how the properties of *s* can contribute powers to it—they simply are its powers. But how can *s* get powers on account of the properties of other things, even its own parts? Certainly it is not generally true that things have the powers that their constituents have. Atoms can enter into covalent bonds but tables cannot. On the other hand, if it is even sometimes possible for other things to contribute powers to *s*, why does it matter whether the entities that instantiate F1-Fn are parts of *s* so long as they are capable (somehow) of contributing powers to *s*?

As a matter of fact, I am willing to grant that when A1-An compose B, then B gets its powers from A1-An. But, first, that is no explanation of how composition works. And, second, I do not suppose that this relation explains realization. In contrast, Gillett appears to be helping himself to the very relation that he hoped to explain: the relation between the properties of some thing(s) and those of a distinct object whose properties they realize.

[8] This example is also used by Tye [1994: 137]. But Tye takes it that the realizer of hardness in a diamond is a crystalline structure, rather than the properties of individual atoms.

[9] I take it that all variations of metaphysical functionalism are versions of RP. Regarding kinds of functionalism, see Polger 2004a.

[10] Kim [1998] locates the origin in Putnam's "Minds and Machines" [1960].

[11] There is a well-known ambiguity in familiar explications of Turing machines, between abstract and physical machines [see Lycan 1974 or Polger 2004a]. Here I am focusing on abstract computational functions, as will be clear from my examples. If computational functions are not abstract in the appropriate way, then think of mathematical functions instead. They, too, are sometimes said to be physically realizable. More on this shortly. (See also note 17.)

[12] Note that Cummins and Van Gulick each use 'instantiate' as a synonym for 'realize' rather than in Gillett's more restricted way.

[13] Versions of this idea are explored by Sellars [1956], Davidson [1963, 1970], Dennett [1971, 1984], and McDowell [1994], among others.

[14] Melnyk [2003] does not require that realization involve having only a causal function. Lewis [e.g., 1970] regularly discusses causal realization (in particular, regarding theories of mind) but his general account is not limited to realization of causally individuated roles.

[15] This is especially attractive if one wants to maintain a hard nosed variety of physicalism. But in the present context it is useful to remember that Putnam was quite clear that non-physical properties could in principle be realizers.

[16] See Polger 2004b.

[17] I will continue to use the computational example because it is salient in philosophy of mind. But if computational relations are not abstract in the relevant sense, then recall that we could also make the point with other examples of formal or non-causal realized properties as well, such as economic, intentional, or semantic properties.

[18] This objection was put to me by Chase Wrenn, José Luis Bermúdez, and several anonymous referees for this journal. I am grateful for their assistance. See also note 11, above.

[19] As Fodor [1974] expresses the idea, "type identification for psychological states can be carried out in terms of 'total states' of an abstract automaton which *models* the organism whose states they are" [1974: fn.8, my emphasis]. So the physical states with which psychological states are to be identified are individuated in terms of formal states of an abstract automaton. To realize a psychological state is to be a physical state that is modeled by (isomorphic to) an abstract automaton state.

[20] Gillett says that he intends to explicate "the notion of realization implicit in scientific theorizing, and not any folk concept" [2003: 594, fn. 9]. He argues that he is pursuing "empirical analysis" of the practices of scientists. Of course, even if we determine that scientists do speak of "realization" in their explanations, the lesson of §1 is that we cannot assume that the scientific use of 'realization' corresponds to the relation invoked by RP.

[21] For this reason, I take issue with Gillett's claim that realization should be understood as a fundamental metaphysical relation.

[22] Does a carburetor realize the air-fuel mixer in a car, or do its parts? (Or is the air-fuel mixer's realization over-determined? Is there a realization exclusion problem?) Let us for now set aside the question of when to invoke the inter-level clause in the dimensioned view.

[23] We will return to this topic when we consider how Gillett treats cases of multiple realization, in §6, below.

[24] This account originates in Polger 2004a. A similar account of realization is offered in Melnyk 2003.

[25] Following Gillett, I have formulated (R) as a relation among properties or states. Two remarks are in order. First, there is no problem in reformulating (R) in terms of entities, or in terms of entities and properties. This will help us to make sense of the RP practice of talking about the realization of entities, as when we say that a carburetor realizes the air-fuel mixer. (Otherwise we must say that an instance of *carburetorness* realizes an instance of *air-fuel-mixerness*, which is at least awkward.) A related point concerns Gillett's restriction of 'instantiation' to the relation between objects and properties. If $F_G(x)$ names a property then that property may be taken to be essential of a entity kind picked out by the functional property. Then *P*'s having that property (i.e., implementing that function, or occupying that functional role) makes *P* an instance of that kind. For that reason, *F* could be said to "instantiate" (i.e., be an instance of) a $F_G(x)$ object. This explanation seems to vindicate the RP use of 'instantiate' along with 'implement' and 'occupy' as synonyms for 'realize'.

[26] On this sketch, the parts of a thing do not individually play its causal role, so causal realization looks "flat" or intra-level. Perhaps there are inter-level notions of function; if

so, some realization relations may be "dimensioned" or inter-level, after all. There is nothing in my account that prevents some functions from taking multiple arguments, or taking n-tuples as arguments. For example, if there is a functional role for being a binary star system, then perhaps it can be satisfied by a pair of stars that have the function of orbiting one another. (Such an account may make for good explanations even if it does not make for good ontology.)

[27] To be precise, the individuating function for Millikan's "proper" functional entities does not involve *current* causal powers (either of the thing or of its parts), but instead involves the causal history of the thing. The crucial point is that being a heart (realizing an entity of the kind, lion heart, say) has nothing to do with what causal powers an entity currently has—for lion hearts can fail to have the causal powers that are typical of lion hearts (they can be broken) and some things that have those powers ("twin" or "swamp" lion hearts) can fail to be hearts.

[28] As Searle puts it: "For any program there is some sufficiently complex object such that there is some description of the object under which it is implementing the program. Thus for example the wall behind my back is right now implementing the Wordstar program, because there is some pattern of molecule movements which is isomorphic with the formal structure of Wordstar. But if the wall is implementing Wordstar then if it is a big enough wall it is implementing any program, including any program implemented in the brain" [1990: 27].

[29] In both variations, we will need to help ourselves to an expansive notion of causal function according to which all of a thing's causal powers contribute to its function. Shoemaker [1984] argues that such a functional characterization of all properties can be

given by Ramsifying the causal theory of properties (CTP), yielding what he calls CTP-functionalism. For more on this topic, see Polger 2004a. There I also argue against admitting a notion of function that includes all causal relations, on pain of making functionalism trivial.

[30] For more variations, see Polger 2002, 2004a; Shapiro 2000, 2004.

[31] In that case, (10) and (11) are to be interpreted along the line of (3), earlier: They say that pains or hearts have been or can be brought about (brought to exist, be "realized") in different creatures.

[32] If we are interested in whether something is multiply realiz*able* (not just multiply realiz*ed*) then we may need to find out what biological theories say about some possible heart candidates (not just actual hearts.)

[33] I am grateful to Chase Wrenn for urging me to clarify this point.

[34] For the purposes of evaluating multiple realizability we are concerned with the extensions of kinds, so accidental or in-world coextension may not be enough to settle the question. We may have to know about the distribution of the kinds under certain counterfactual conditions.

[35] Gillett defends the dimensioned view from the charge that it makes MR trivial by counting any physical differences in realizers as cases of MR. He offers the example of two aluminum corkscrews which differ only in that one contains "a trace element" that "does not chemically bond with the aluminum or change the metallic structure of the aluminum atoms, but it does absorb a certain wavelength of light giving this corkscrew a yellow tinge" [2003: 598-99]. He rightly concludes that the dimensioned view need not count these as case of MR, for the presence of the trace element is irrelevant to the

causal powers that the parts contributes to the corkscrew. But I wonder whether this case only works because this element does not chemically bond with the aluminum, and thus is not a genuine part of the corkscrew but merely an imperfection in it. If the element bonded with the aluminum (creating an aluminum alloy) then would this not be just like the aluminum/steel example, and would the dimensioned view not have to claim that this was a different realization of the corkscrew? I do not see how Gillett can avoid this result.

9. References

Bechtel, W. and J. Mundale. 1999. Multiple Realizability Revisited: Linking Cognitive and Neural States, *Philosophy of Science*, 66: 175-207.

Bickle, J. 1998. *Psychoneural Reduction: The New Wave*, Cambridge, MA: MIT Press.

Block, N. (ed.) 1980. *Readings in Philosophy of Psychology, Volume One*, Cambridge, MA: Harvard University Press.

Block, N. 1990. Can the Mind Change the World?, in *Meaning and Method: Essays in Honor of Hilary Putnam*, G. Boolos (ed.), (Cambridge, UK: Cambridge University Press).

Block, N. and J. Fodor. 1972. What Psychological States Are Not, *Philosophical Review* 81: 159-181.

Chalmers, D. 1994. On Implementing a Computation, *Minds and Machines* 4 (4): 391-402.

Chalmers, D. 1996. Does a Rock Implement Every Finite-State Automaton?, *Synthese* 108: 309-333.

Clapp, L. 2001. Disjunctive Properties: Multiple Realizations, *Journal of Philosophy*, 98: 111-136.

Cummins, R. 1989. *Meaning and Mental Representation*, Cambridge, MA: The MIT Press.

Davidson, D. 2001 [1963]. Actions, Reasons and Causes, in D. Davidson, *Essays on Actions and Events* (Oxford: Clarendon Press, 2001).

Davidson, D. 2001 [1970]. Mental Events, in D. Davidson, *Essays on Actions and Events* (Oxford: Clarendon Press, 2001).

Dennett, D. 1971. Intentional Systems, *Journal of Philosophy*, LXVIII: 87-106.

Dennett, D. 1984. *Elbow Room: The Varieties of Free Will Worth Wanting*, Cambridge, MA: MIT Press.

Field, H. 1980 [1978]. Mental Representation, *Erkenntnis* 13, 1: 9-61. Reprinted in N. Block (1980).

Fodor, J. 1980 [1974]. Special Sciences, or the Disunity of Science as a Working Hypothesis. *Synthese* 28: 97-115. Reprinted in N. Block (1980).

Gillett, C. 2002. The Dimensions of Realization: A Critique of the Standard View, *Analysis* 64, 4: 316-323.

Gillett, C. 2003. The Metaphysics of Realization, Multiple Realizability, and the Special Sciences, *Journal of Philosophy,* Nov. 2003: 591-603.

Heil, J. 1992. *The Nature of True Minds*, New York: Cambridge University Press.

Heil, J. 1999. Multiple Realizability, *American Philosophical Quarterly* 36, 3: 189-208.

Horgan, T. 1993. From Supervenience to Superdupervenience: Meeting the Demands of a Material World, *Mind* 102: 555-586.

Horgan, T. and J. Tienson. 1996, *Connectionism and the Philosophy of Psychology*, Cambridge, MA: MIT Press.

Jackson, F. 1998. *From Metaphysics to Ethics: A Defense of Conceptual Analysis*, Oxford: Oxford University Press.

Jackson, F. and P. Petit. 1990. Program Explanation: A General Perspective, *Analysis* 50: 107-117.

Jackson, F. and P. Pettit. 1995. Moral Functionalism and Moral Motivation, *Philosophical Quarterly* 45(178): 20-40.

Jackson, F. and P. Pettit. 1996. Moral Functionalism, supervenience and reductionism, *Philosophical Quarterly* 46(182): 82-86.

Kim, J. 1992. Multiple Realization and the Metaphysics of Reduction, *Philosophy and Phenomenological Research*, 52: 1-16.

Kim, J. 1996. *Philosophy of Mind*, Boulder, CO: Westview.

Kim, J. 1998. *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*, Cambridge, MA: MIT Press.

Kim, J. 2005. *Physicalism, or Something Near Enough*, Princeton: Princeton University Press.

Lewis, D. 1970. How to Define Theoretical Terms, *Journal of Philosophy* 67 (13): 427-446.

Lycan, W. 1974. Mental States and Putnam's Functionalist Hypothesis, *Australasian Journal of Philosophy* 52: 48-62.

Lycan, W. 1987. *Consciousness*, Cambridge, MA: The MIT Press.

Lynch, M. 2001. A Functionalist Theory of Truth, in M. Lynch (ed.), *The Nature of Truth: Classic and Contemporary Perspectives*. Cambridge, MA: The MIT Press.

Lynch, M. 2004. Truth and Multiple Realizability, *Australasian Journal of Philosophy* 82: 382-408.

Marr, D. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*, San Francisco: W.H. Freeman.

Marr, D. and E. Hildreth. 1980. Theory of Edge Detection, *Proceedings of the Royal Society of London*, B 207: 187-217.

McDowell, J. 1994. *Mind and World*, Cambridge, MA: Harvard University Press.

Melnyk, A. 2003. *A Physicalist Manifesto: Thoroughly Modern Materialism*, Cambridge: Cambridge University Press.

Millikan, R. 1984. *Language, Thought, and Other Biological Categories*, Cambridge MA: MIT Press.

Millikan, R. 1989. In Defense of Proper Functions, *Philosophy of Science* 56: 288-302.

Neander, K. 1991. Functions as Selected Effects: The Conceptual Analyst's Defense, *Philosophy of Science* 58: 168-184.

Poland, J. 1994. *Physicalism: The Philosophical Foundations*, New York: Oxford University Press.

Polger, T and K. Sufka. 2006. Closing the Gap on Pain: Mechanism, Theory, and Fit, in *New Essays on the Nature of Pain and the Methodology of its Study*, M. Aydede (ed.). Cambridge, MA: The MIT Press.

Polger, T. 2002. Putnam's Intuition, *Philosophical Studies*, 109, 2: 143-170.

Polger, T. 2004a. *Natural Minds*, Cambridge, MA: The MIT Press.

Polger, T. 2004b. Neural Machinery and Realization, *Philosophy of Science*. 71 (5): 997-1006.

Putnam, H. 1975 [1967]. The Nature of Mental States, in H. Putnam (Ed.) *Mind, Language and Reality: Philosophical Papers, Volume 2*. (New York: Cambridge University Press, 1975).

Putnam, H. 1975. Philosophy and Our Mental Life, in H. Putnam, *Mind, Language and Reality: Philosophical Papers, Volume 2*. New York: Cambridge University Press.

Putnam, H. 1988. *Representation and Reality*, Cambridge, MA: The MIT Press.

Putnam, H. 1999. *The Threefold Cord: Mind, Body, and World*, New York: Columbia University Press.

Scheutz, M.1998. Implementation: Computationalism's Weak Spot, *Conceptus* 31 (79): 229-239.

Scheutz, M.1999. When Physical Systems Realize Functions..., *Minds and Machines* 9 (2): 161-196.

Searle, J. 1990. Is the Brain a Digital Computer?, *Proceedings and Addresses of the American Philosophical Association* 64: 21-37.

Sellars, W. 1956/1997. *Empiricism and the Philosophy of Mind*, Cambridge, MA:, Harvard University Press.

Shapiro, L. 2000. Multiple Realizations, *The Journal of Philosophy*, 97, 635-654.

Shapiro, L. 2004. *The Mind Incarnate*, Cambridge, MA: The MIT Press.

Shoemaker, S. 2001. Realization and Mental Causation, in Gillett, C. and B. Loewer (eds.). 2001. *Physicalism and Its Discontents*. Cambridge, UK: Cambridge University Press.

Shoemaker, S. 2003. Realization, Micro-Realization, and Coincidence, *Philosophy and Phenomenological Research*, LXVII (1): 1-23.

Sober, E. 1999. The Multiple Realizability Argument Against Reductionism, *Philosophy of Science* 66: 542-564.

Tye, M. 1994. Naturalism and the Problem of Intentionality, *Midwest Studies in Philosophy*, XIX: 122-142.

Van Gulick, R. 1983. Functionalism as a Theory of Mind, *Philosophy Research Archives*: 185-204.

Van Gulick, R. 1988. Consciousness, Intrinsic Intentionality, and Self-Understanding Machines, in Marcel and Bisiach (Eds), *Consciousness in Contemporary Science*. (New York: Oxford University Press, 1988.)

Wilson, R. 2001. Two Views of Realization, *Philosophical Studies*, 104: 1-30.