Multiple Realizability

Eric Funkhouser

Abstract: This article explains the concept of multiple realizability and its role in the philosophy of mind. In particular, I consider what is required for the multiple realizability of psychological kinds, the relevance of multiple realizability to the reducibility and autonomy of psychology, as well as further refinements of the concept that would prove helpful.

1. The Basic Idea and Its History

A kind (property, type, activity, etc.) is *multiply realized* if, in some sense to be specified, there are many different ways things are of that kind. Pain, for example, is multiply realized if there are many different ways animals are in the same kind of pain. A kind is *multiply realizable* if it is possible for it to be multiply realized.

The concept of multiple realizability (MR) was introduced by Hilary Putnam (1975; Chapters 14, 18, 19, and 21) and Jerry Fodor (1965, 1968, and 1974) in the 1960's as an objection to mental-physical identity theories in the philosophy of mind and as a welcome consequence of their functionalist alternative. Identity theories rose to prominence in the late 1950's as a result of influential papers by, among others, U.T. Place (1956), Herbert Feigl (1958), and J.J.C. Smart (1959). According to these identity theories, mental states (in particular, states of consciousness or sensations) are identical to physical states (in particular, brain states). To give the philosopher's favorite fictional example, pain is identical to the firing of C-fibers. Such identities were supposed to be analogous to theoretical identities in other domains – e.g., lightning is identical to motion of electric charges. The claim was that just as lightning is motion of electric charges, pain is the firing of C-fibers. This is a *type* (or *kind*) identity theory. For the claim was

not merely that a particular flash of lightning is motion of electric charges, but that every flash of lightning is motion of electric charges. The mental-physical identity theorists believed that a similar correlation holds between mental kinds and physical kinds, and that this correlation is best explained by identifying mental kinds with physical kinds. This seemed to be a tidy and scientifically respectable solution to the mind-body problem.

These classical versions of the identity theory are viable only if there is a simple one-to-one correspondence between mental kinds and physical kinds. This is where MR comes in. Putnam and Fodor denied that pain, or any other mental kind, correlates with only one physical kind. Rather, they argued that there are many ways that pain and other mental kinds can be *realized* in the physical. In short, mental kinds are not identical to physical kinds because there are many and diverse physical kinds that can give rise to the same mental kind. For example, even if pain in human beings were strictly correlated with the firing of C-fibers, it is likely that other kinds of physical states give rise to pain in some other animals. Putnam vividly emphasized the difficulties that confront the identity theorist by doubting that any one physical-chemical state correlates with pain in mammals, reptiles, and molluscs. (Putnam (1975), p. 436) Since the identity relation is one-to-one, but the mental kind to physical kind relation is one-to-many, the type identity theory must be false. This is the MR objection to the type identity theory.

To properly understand the concept of MR and its genesis one needs to have some understanding of the theory of mind that Putnam and Fodor offered as an alternative to identity theory and behaviorism. Putnam and Fodor argued that psychological kinds are functional kinds characterized by their causal relations to each other, sensory inputs, and

behavioral outputs. To simplify greatly, pain is the psychological kind that causes an urgent desire to change one's situation, is caused by tissue damage, and is followed by wincing. For functionalists the essence of a mental state lies with these causal roles rather than, say, its realization details (e.g., physical composition and structure). This view immediately opens up the *possibility* of MR – perhaps many different physical compositions and structures could occupy the functional role that is pain. Some have even thought that if psychological kinds are functional kinds, then it is a conceptual truth that psychological kinds are multiply realizable. After all, in general we can conceive of various mechanisms fulfilling a certain job description (functional role). More cautious functionalists have acknowledged that their functionalism only opens up the possibility of MR, and that empirical investigation is then needed to settle the issue. Block and Fodor (1972) is an early source of broadly empirical arguments for the MR of psychology, offering three such reasons: 1) brain plasticity, 2) the convergent evolution of psychological traits, and 3) the possibility of artificial intelligence.

Early discussions of functionalism often described the mind as a kind of software which can be run on different types of hardware. Those functionalists who take this computer metaphor seriously, *computational functionalists*, hold that psychological kinds are categorized by their relations to, and syntactic operations on, internal symbols (mental representations). On this picture the mind is a program that is implemented in the brain, just as more familiar computer programs are implemented in personal computers. Such computer programs can be implemented or realized in various types of hardware, and the same is thought to be true of psychological kinds. Further, computer programs can be studied and understood in ignorance of the hardware in which they are realized, as the

programs are at a higher level of abstraction. If the computer metaphor is apt, then psychology can also be developed and understood in ignorance of neuroscience. This has been a prominent methodological assumption of many cognitive scientists.

2. Denials of Multiple Realizability

Our MR argument consists of a premise stating that psychological kinds are multiply realizable and moves to a conclusion that psychology does not reduce to neuroscience or other lower level sciences. One could object by either denying the premise or denying the move to the conclusion. The MR of the mental has been widely accepted by philosophers in the wake of the functionalist's arguments and relatively few have opted to deny the premise. Resistance to the premise has, however, gradually emerged. In this section I present reasons that have been offered for denying the premise. Sections 3 and 4 consider the move from MR to non-reductionism.

The MR of psychological kinds requires sameness of psychological kind through differences in, say, neuroscientific kind. This characterization suggests two strategies for denying alleged cases of MR. First, one could argue that, despite initial appearances, the cases do not involve subjects of the same psychological kind. Second, one could concede the sameness of psychological kind but argue that the subjects are of the same neuroscientific kind as well. Both strategies have been pursued.

Thomas Polger (2002) challenged Putnam's claim that humans and octopi can have pain states of precisely the same (determinate) psychological kind. If Polger is correct, then the realization differences for their pains are irrelevant to the claim of MR. It takes both conceptual and empirical work to evaluate Polger's objection. First, a combination

of conceptual and empirical work is needed to clarify the standards for sameness of psychological kind (which should be applicable to both humans and octopi). All scientific taxonomies result from the combination of data collection and theory construction. Purely conceptual considerations can give us a proto-theory to start with and also provide constraints on theorizing, while empirical research suggests refinements of our taxonomies in light of new discoveries. Second, empirical work is required to see if humans and octopi can satisfy the same psychological kind, though differing in their realizations. (The defender of MR could deny that it requires *precise* sameness of psychological kind, though this response seems to go against the common understanding of MR.)

William Bechtel and Jennifer Mundale (1999) pursue the second strategy, arguing that neuroscientific classifications accord well with psychological classifications. Bechtel and Mundale claim that MR seems plausible only because we tend to use different grains of individuation conditions when distinguishing psychological and neuroscientific kinds. That is, a quite general psychological kind, such as pain, is said to have diverse realizations in humans and octopi. But Bechtel and Mundale claim that if realizations are individuated using an equally coarse grain, then this appearance of MR evaporates. They argue that humans and octopi that are in pain have common realizations at a quite general level of abstraction. Of course, at other levels of abstraction their realizations are of different kinds, but at those levels of abstraction their pains are of different psychological kinds as well. Bechtel and Mundale explain the correspondence between psychological and neuroscientific classifications by arguing that considerations of functional role and cross-species comparisons also guide

neuroscientific taxonomies. If psychological kinds have functional essences that can be satisfied by members of distinct species, it is then no surprise that there is a correspondence between neuroscientific and psychological classifications at a common level of abstraction.

Lawrence Shapiro (2004) also challenges MR by arguing that there is good reason to think that subjects instantiating the same psychological kind, even across species, likely have enough biological similarity to count as sharing a realization. This is because, as an empirical fact, there are quite limiting constraints on the ways in which these psychological kinds can be realized on Earth. More importantly, Shapiro argues that whether psychological kinds are multiply realized is an empirical question, rather than a (largely) *a priori* consequence of functionalism. Sometimes there is only one person for the job, and sometimes there is only one way (or few ways) to realize a functional kind.

3. Multiple Realizability and Reduction

MR is a powerful tool that has been used against reductionist theories of mind. But what is reduction? Speaking loosely, theory (or science) *A* reduces to theory (or science) *B* if, and only if, the classifications and predictive/explanatory generalizations of theory *A* can be captured, to a close approximation, by theory *B*, *and* theory *B* has superior explanatory or theoretical virtues such as greater scope, more specificity of predictions/explanations, simplicity, etc. In this section I distinguish some more specific or limited types of reduction and show how MR can be used to motivate corresponding types of non-reductionism.

3.1. Methodological Reduction

We can speak of the *methodology* of one science reducing to the methodology of another science. How is this to be understood? Let's take psychology and neuroscience as an example. If the methodology of psychology were to reduce to the methodology of neuroscience, then the *way we do* neuroscience would become the *way to do* psychology. Of course various scientific disciplines have overlapping scientific methodologies at a general level of abstraction, so a methodological reduction requires something rather specific. Namely, the classifications, generalizations, and distinctive research methods of the reducing science must become essential to the satisfactory practice of the methodologically reduced theory. In our case, a methodological reduction would require that psychologists in good standing learn neuroscience and do their theorizing and investigations in its terms and according to its specific methods.

What connection, if any, does MR have to methodological reduction? Since Putnam and Fodor it has been argued that MR ensures methodological *autonomy*, the denial of methodological reductionism. For if a diverse range of neuroscientific kinds realizes a given psychological kind, then the psychological kinds crosscut the neuroscientific kinds. It then becomes difficult to see how insights and generalizations couched in psychological vocabulary could benefit from recognition of neuroscientific classifications.

Methodological reductionists have responded that recognition of the neuroscientific realizations of psychological kinds could suggest a reworking of psychological taxonomies. Bechtel and Mundale (1999) offer research on visual processing as such an example. (Each of the "new wave" reductionists mentioned in section 4 offers similar

examples.) But even if no such alterations of psychological classifications and theory are suggested after descending to the lower level, such descent could nevertheless uncover psychological patterns that, as an empirical fact, would not have been discovered otherwise.

3.2. Explanatory Reduction

An *explanatory reduction*, in one sense, is the discovery of the ways in which underlying mechanisms give rise to some higher level kind. It should be clear that MR is compatible with explanatory reduction in this sense.

In another sense, an *explanatory reduction* obtains whenever it is shown that (nearly) everything that one theory purports to explain is better explained by another theory in good standing. (Again, the explanations of the latter theory could be better in virtue of possessing various explanatory virtues, such as greater scope, accuracy, etc.) So if it were shown that (nearly) all psychological explanations can also be given in neuroscientific terms (but not *vice versa*), then we would have an explanatory reduction of psychology to neuroscience. Early functionalists argued that MR blocks such a reduction – that is, it ensures that psychology is explanatorily autonomous. Putnam famously gave the example of a cubical peg that is capable of passing through a square, but not a round, hole. He claimed that this fact could be explained at the relatively abstract geometrical level, but not from the microstructural level. (Putnam (1975), pp. 295-296) This is because the same geometrical properties can be realized in various microstructures – MR again! And just as microstructural details cloud the picture when

searching for an explanation of the peg's behavior, perhaps neuroscientific details cloud the picture when searching for an explanation of a person's behavior.

Elliot Sober (1999) objects that physics can explain everything that higher level sciences can explain. He acknowledges Putnam's observation that the perspective from the lower level sciences often provides too much detail. But Sober argues that this does not mean that the lower level sciences are incapable of explaining the peg's behavior. They simply might offer more detailed explanations than are useful in some circumstances. But in other circumstances explanatory depth is preferred to breadth.

Jaegwon Kim (1993), Ruth Millikan (1999), and Richard Boyd (1999) have even argued that there cannot be a science of multiply realized (functional) kinds. They claim that we should not expect any generalizations/explanations to hold for such kinds, except for those captured by their functional specification. For such generalizations are blocked by the very diversity of lower level realizations. Their claim certainly runs counter to standard methodological assumptions of functionalists and cognitive scientists. The best way to refute their position, and certainly the most straightforward, is simply to discover such interesting generalizations.

3.3. Ontological Reduction

The most fundamental sense of 'reduction' is *ontological reduction*. An ontological reduction occurs whenever it is shown that the entities and laws picked out by the vocabulary of one science, with perhaps a little theoretical modification, can also be picked out by the vocabulary of another science. (Further, the reducing science has ontological priority in virtue of possessing some set of theoretical virtues.) Here 'entity'

is used in a broad sense to include both kinds/properties and substances. All physicalists agree that there is only one type of substance – physical substance – and most agree that the particular substances picked out by higher level sciences reduce to the substances picked out by lower level sciences. Disagreement emerges with respect to kinds. Several non-reductionists have argued that since psychological kinds are multiply realized by neuroscientific kinds, and since "wild" or "heterogeneous" disjunctions of neuroscientific kinds are not themselves neuroscientific kinds, then psychological kinds are not identical to neuroscientific kinds. In that case psychology does not ontologically reduce to neuroscience.

Fodor (1974) provides the classic argument for ontological non-reduction from MR, focusing on laws in particular. He asserts that the ontological reduction of special science *S* requires that the laws of *S* be derivable from the laws of physics (or some intermediary lower level science). Of course, the laws of physics are formulated in terms of the predicates of physics. Fodor claims that the laws of *S* can be derived from the laws of physics only if there are "bridge" laws connecting predicates of *S* with those of physics. If these bridge laws are really laws, then Fodor insists that the predicates used to formulate them must pick out kinds. (Out of agreement with Fodor on this point, I have characterized ontological reduction in terms of entities and laws – kinds and laws go hand-in-hand.) But because special science kinds are typically multiply realized in the physical, the physical predicates that appear in the proposed bridge laws mark off a heterogeneous and unsystematic collection that is not a physical kind. So, there are no such bridge laws – higher level kinds and their laws do not reduce to physical kinds and their laws. Fodor still believes, however, that every special science event is identical to a

physical event. And this commitment, which he calls *token identity*, is all that he thinks is required of physicalists. Such a combination of kind-pluralism with token identities (of events, property instances, states, etc.) has been widely, but not universally, accepted by non-reductive physicalists.

4. Reductionism Strikes Back

It would be premature, however, to proclaim the death of reductionism. Several philosophers have accepted the claim that psychological kinds are multiply realized while nevertheless denying these non-reductionist conclusions. Three reasons have been offered in favor of the view that ontological reductionism is compatible with MR.

4.1. Local Reductions

First, some have argued for localized, species, or structure-specific reductions of psychological kinds to neuroscientific kinds. Patricia Churchland (1986), David Lewis (1994) and Jaegwon Kim (1993, 1998) recommend this strategy. On their view human-pain, say, reduces to one neuroscientific kind, while octopus-pain reduces to another neuroscientific kind. This strategy, however, concedes the point that the more general kind – pain *simpliciter* – does not so reduce. In fact, the advocate of local reductions will likely be an eliminativist about psychological kinds like pain *simpliciter* and might be skeptical about the very possibility of a science of multiply realized kinds. The move to local reductions also fits nicely with the objection, discussed in section 2, that human-pain and octopus-pain are not of the same psychological kind. However, those who deny that objection will likely find local reductions inadequate to justify the label 'reduction'.

And the non-reductionist will likely contend that general kinds, like pain *simpliciter*, still play an essential explanatory and predictive role in psychological theorizing.

4.2. Disjunctive Kinds

Second, one could argue that multiply realized psychological kinds are to be identified with, and thereby reduced to, neuroscientific kinds picked out by disjunctive neuroscientific predicates. This strategy purports to salvage classical type identity theories in a very straightforward manner. Given MR, such disjunctions would be quite long and heterogeneous from the neuroscientific point of view. But Kim (1993) argues that if a psychological kind is legitimate (i.e., a unity grounded in real similarity), then the kind picked out by the disjunctive neuroscientific predicate that it is coextensive with is equally legitimate as a kind. The non-reductionist's rejoinder, following Fodor (1974, 1997), will likely be that such disjunctive "kinds" are not neuroscientific kinds at all. Such disjunctions appear in no neuroscientific generalizations and might not even be projectible into generalizations (i.e., they might not even be *capable* of occurring in any scientific generalization). Their heterogeneity means that they do not possess the similarity or unity required of kinds, even though they have the common property of realizing the same psychological kind.

Disputes like this between Fodor and Kim are often clouded by two factors. First, frequently there is slippage between talk of predicates and talk of kinds (properties, types, etc.), or the invocation of a claim about predicates when kinds are relevant. Ontological reduction, in the first instance, concerns kinds rather than predicates – it is not simply a matter of pairing vocabulary across sciences. Rather, an ontological reduction obtains

only when it is discovered that the predicates of some lower level science can pick out the same kinds as are picked out by the predicates of some higher level science. We can always append logical operators to predicates to form other predicates, for example disjunctive predicates. Indeed, if token physicalism is true then every special science predicate will be co-extensive with some such disjunctive physical predicate. But, it does not (yet) follow that the disjunctive physical predicate expresses the same kind as the corresponding special science predicate, nor does it follow that the disjunctive physical predicate expresses any kind at all.

The second factor that clouds such disputes is the fact that claims about kinds are often advanced without a sophisticated enough account of kinds to back them up. For example, Kim (1993) considers a psychological predicate, P, with only three kinds of physical realization. Let N be the disjunction of these three physical predicates. Kim (1993) claims that N expresses a kind to the very same degree that P expresses a kind. Here, nomological co-extensiveness with a kind is taken to be sufficient for a kind identity. Fodor (1997) objects that the open-ended disjunctions of physical predicates which are co-extensive with psychological kinds are not proper candidates for laws because they are "gerrymandered" and, more generally, open disjunctions are not acceptable predicates for laws. It might be better to attack Kim more directly, however. While N is true of all and only the same things as P, P nevertheless expresses a kind that is at a different level of abstraction than the kind (if any) expressed by N. The claim is that Kim's nomological co-extensiveness does not establish sameness of kind. This response is most clearly available to the property realist, such as the trope theorist, who can admit that N and P are true of all and only the same things while denying that the set

of N property instances is co-extensive with the set of P property instances. This is possible because the N-properties and P-properties could be at different levels of abstraction, and the trope theorist admits abstract particulars into her ontology.

4.3. "New Wave" Reduction

Third, the classical picture of reduction inspired by Nagel (1961), which requires bridge laws that are often interpreted as inter-theoretic kind identities or bi-conditionals, could be replaced with an alternative theory of reduction which purports to be compatible with MR. Such "new wave" theories of reduction have been outlined and advocated by Robert Richardson (1979), Clifford Hooker (1981), Patricia Churchland (1986), and John Bickle (1998). This has become a popular route for those who claim to be reductionists, especially when coupled with a commitment to localized reductions. One worry about revisionary accounts, however, is that they might save reduction in name only, by defending a sense of reduction that Fodor and other non-reductionists never intended to deny. Here, as elsewhere, we should be careful that our disputes are not merely verbal.

5. Refining the Concept

Surprisingly, most of the discussion of MR through the 1990's unfolded without its participants articulating an account of *realization* or the relevant sense of *multiple* realizations. In recent years greater attention has been given to offering a more rigorous characterization of these concepts.

5.1. 'Realization'

Robert Wilson (2001) correctly notes that there are two facets to the standard view of realization. (Wilson goes on to argue against each.) First, realizers are (synchronically) *sufficient* for the kinds they realize. Second, realizers *constitute* that which they realize. On the standard view, if a certain brain state realizes pain in human beings then that brain state is sufficient for human-pain and constitutes human-pain whenever that brain state occurs. The sufficiency clause makes realization akin to supervenience and other asymmetric necessitation relations. Such relations also preserve basic physicalist principles in that their obtaining ensures that the physical truths determine the mental truths. The constitution clause separates realization from relations such as supervenience, however. Supervenience is merely a logical claim of law-like variation and is compatible with various theories of mind, including substance dualism. The constitution clause ensures that mental-in-physical realization rules out substance dualism by guaranteeing that the mental is there *with* or *in* the physical. As such, realization is a better relation than supervenience for capturing the physicalist's views on the mental-physical relationship. Accordingly, it has succeeded supervenience as the critical relation for understanding both reductive and non-reductive theories of mind. Whether or not the standard view is correct, it does help to give an intuitive understanding of what most mean by 'realization'.

Returning to the computer metaphor offers another route to understanding realization. A computer can be conceptualized at many different "levels" – e.g., machine language, assembly language, higher level languages, etc. These are different levels of abstraction. A realized kind (e.g., software) similarly can be thought of as being at a higher level of abstraction than its realizations (e.g., the hardware on which it is run).

There have been two prominent approaches to the metaphysics of realization. On the first approach, functional properties are thought of as second-order properties. Namely, a functional property is the property of having some property or other that meets a given functional specification. The first-order properties that meet the functional specification are the realizers. So, the realization relation is that relation which holds between such second-order and first-order properties. This account is endorsed by Kim (1993, 1998), but traces back to David Lewis (see his (1994), for example) and others.

The second approach, compatible with the first, focuses on the causal powers of the realized and realizing kinds. On this view realization is understood as inheritance of causal powers. Namely, realized kinds inherit a proper subset of the causal powers of their realizers. Advocates of this approach also typically hold that the essence of a kind lies with its causal powers. This approach is most prominent in Sydney Shoemaker (2001), and is also found in Kim (1993) and Carl Gillett (2002, 2003).

Another important distinction is made in Shoemaker (1981), where he distinguishes *core realizations* and *total realizations*. Core realizations make the most significant, or at least salient, contributions to bringing about their realized kinds. However, core realizations are not strictly sufficient for their realized kinds. Instead, core realizations must be placed in the appropriate context and supplemented with other conditions in order to successfully bring about those realized kinds. For example, a certain piece of paper is the core realization of a U.S. dollar. But that piece of paper must be properly situated within a larger institutional setting — involving the U.S. Treasury, etc. – in order for that piece of paper to be worth one dollar. The total realization, which is sufficient for this task, is the combination of the core realization and these supplemental conditions.

This is an important distinction to keep in mind. One may wonder whether the more interesting MR claims cover core or total realizations. (I would think the more interesting claims involve core realizations, with total realizations playing the role of background conditions.) Wilson (2001) also exploits the core/total realization distinction in arguing for a context sensitive account of realization, in contrast to the standard view sketched above.

5.2. 'Multiple'

The MR of a kind is supposed to be a non-trivial claim with real methodological, explanatory, and/or ontological consequences. At the outset we offered a loose characterization of MR as the existence of many different (possible) ways to be of a kind. But this characterization is much too loose for any serious purpose – many kinds can be realized in various ways without meeting our intuitive standards for MR. What counts as a "way" and how many or diverse must these ways be for a kind to be multiply realizable?

Shapiro (2004) tackles these questions for functional kinds. On his view the essence of a functional kind is its role in bringing about some outcome or exercising some capacity. Certain properties – the *R-properties* – are particularly relevant to this performance, and these R-properties are identified by the science of that functional kind. Shapiro claims that it is only differences in these R-properties that can ground claims of MR. For example, the waiter's corkscrew and the double-lever corkscrew are distinct realizations of the corkscrew kind. This is because they differ in the properties relevant to explaining how they remove corkscrews. If there were a science of corkscrews, it

would discern this difference. Shapiro offers a revisionary account of MR, in that on his view compositional differences are typically not sufficient to ground MR – e.g., aluminum and steel waiter's corkscrews are not distinct realizations.

Eric Funkhouser (2006, Forthcoming) argues that Shapiro's conclusion regarding R-properties gets things backwards. Differences that can be discerned from the perspective of a particular science cannot ground the MR of its kinds. Rather, differences that *cannot* be discerned from the perspective of a particular science ground their MR. Psychology is thought to be multiply realized because psychological kinds can differ in their realizations, though these differences are indiscernible (and irrelevant) to psychological generalizations. It is MR in this sense that is supposed to ground its autonomy. Compare: Because crimson and scarlet are colors that can be distinguished by the science of color, they are not distinct realizations of the kind red. Instead, they are determinates of a common determinable.

It should also be noted that kinds are multiply realized *relative* to other kinds. This fits well with its connection to reduction, as theories are reduced to other theories and not reduced *simpliciter*. So, a kind might be singularly realized with respect to one kind, but multiply realized with respect to another kind. For example, water is singularly realized relative to the chemical kind H_2O, but multiply realized relative to lower levels of abstraction.

6. Further Applications and Future Research

MR is widely thought to hold in special sciences besides psychology, especially when such sciences invoke functional kinds. Indeed, Fodor (1974) did not argue for any claim

unique to psychology, and he used economics as a prominent example. David Hull (1974) and Philip Kitcher (1984) have also argued for the irreducibility of biological theories (e.g., Mendelian genetics) on the basis of MR. The concept has even been employed against reductionist theories outside of science – e.g., Michael Lynch (2004) has used it to argue against reductionism about truth.

Future research will likely continue to develop the central issues involved with both the premise and the move to the conclusion of our MR argument. With respect to the premise, further clarification of both 'realization' and 'multiplicity' is needed. Then, more empirical work and theorizing about the taxonomies of the different sciences is needed to determine whether psychology meets this refined account of MR. With respect to the conclusion, the connections between MR, autonomy, and reduction (in their various senses) need to be more carefully delineated. This is no small project. Finally, it would be nice to have a good explanation of how sciences that generalize over multiply realized kinds are possible, even if we know that they are actual.*

*I wish to thank Brie Gertler and an anonymous referee for comments that significantly improved this paper.

References

Bechtel, William and Jennifer Mundale. 1999. "Multiple Realizability Revisited: Linking Cognitive and Neural States", *Philosophy of Science* 66, pp. 175-207.

Bickle, John. 1998. *Psychoneural Reduction: The New Wave* (Cambridge, MA: The MIT Press).

Block, Ned and Jerry Fodor. 1972. "What Psychological States are Not", *Philosophical Review* Vol. 81, No. 2, pp. 159-181.

Boyd, Richard. 1999. "Kinds, Complexity and Multiple Realization," *Philosophical Studies* 95, pp. 67-98.

Churchland, Patricia. 1986. *Neurophilosophy* (Cambridge, MA: The MIT Press).

Feigl, Herbert. 1958. "The 'Mental' and the 'Physical'", *Minnesota Studies in the Philosophy of Science*, Vol. 2, pp. 370-497.

Fodor, Jerry. 1965. "Explanations in Psychology", in *Philosophy in America*, ed. Max Black (Ithaca, NY: Cornell University Press).

Fodor, Jerry. 1968. *Psychological Explanation: An Introduction to the Philosophy of Psychology* (New York, NY: Random House).

Fodor, Jerry. 1974. "Special Sciences, or the Disunity of Science as a Working Hypothesis", *Synthese* 28, pp. 97-115.

Fodor, Jerry. 1997. "Special Sciences: Still Autonomous After All these Years", *Philosophical Perspectives* 11, pp. 149-163.

Funkhouser, Eric. 2006. "The Determinable-Determinate Relation", *Nous* 40.3, pp. 548-569.

Funkhouser, Eric. Forthcoming. "A Liberal Conception of Multiple Realizability", *Philosophical Studies*.

Gillett, Carl. 2002. "The Dimensions of Realization: A Critique of the Standard View", *Analysis* 62.4, pp. 316-323.

Gillett, Carl. 2003. "The Metaphysics of Realization, Multiple Realizability, and the Special Sciences", *Journal of Philosophy* 100, pp. 591-603.

Hooker, Clifford. 1981. "Towards a General Theory of Reduction (Parts I-III)", *Dialogue* 20, pp. 38-59, 201-236, and 496-529.

Hull, David. 1974. *Philosophy of Biological Science* (Englewood Cliffs, NJ: Prentice-Hall).

Kim, Jaegwon. 1993. "Multiple Realization and the Metaphysics of Reduction," reprinted in his *Supervenience and Mind*, (New York, NY: Cambridge University Press).

Kim, Jaegwon. 1998. *Mind in a Physical World* (Cambridge, MA: The MIT Press).

Kitcher, Philip. 1984. "1953 and All That: A Tale of Two Sciences", *Philosophical Review* 93, pp. 335-373.

Lewis, David. 1994. "Lewis, David: Reduction of Mind", in *A Companion to the Philosophy of Mind*, ed. Samuel Guttenplan (Oxford: Blackwell Publishers).

Lynch, Michael. 2004. "Truth and Multiple Realizability", *Australasian Journal of Philosophy* Vol. 83, No. 3, pp. 384-408.

Millikan, Ruth Garrett. 1999. "Historical Kinds and the 'Special Sciences'", *Philosophical Studies* 95, pp. 45-65.

Nagel, Ernest. 1961. *The Structure of Science* (London: Routledge and Kegan Paul).

Place, U.T. 1956. "Is Consciousness a Brain Process?", *British Journal of Psychology* 47, pp. 44-50.

Polger, Thomas. 2002. "Putnam's Intuition", *Philosophical Studies* 109, pp. 143-170.

Putnam, Hilary. 1975. *Mind, Language and Reality: Philosophical Papers, Volume 2* (New York, NY: Cambridge University Press).

Richardson, Robert. 1979. "Functionalism and Reductionism", *Philosophy of Science* 46, pp. 533-558.

Shapiro, Larry. 2004. *The Mind Incarnate* (Cambridge, MA: The MIT Press).

Shoemaker, Sydney. 1981. "Some Varieties of Functionalism", *Philosophical Topics* 12, pp. 93-119.

Shoemaker, Sydney. 2001. "Realization and Mental Causation", in *Physicalism and Its Discontents*, eds. Carl Gillett and Barry Loewer (New York, NY: Cambridge University Press).

Smart, J.J.C. 1959. "Sensations and Brain Processes", *Philosophical Review* 68, pp. 141-156.

Sober, Elliott. 1999. "The Multiple Realizability Argument Against Reductionism", *Philosophy of Science* 66, pp. 542-564.

Wilson, Robert. 2001. "Two Views of Realization", *Philosophical Studies* 104, pp. 1-31.