The Multiple Realizability Argument against Reductionism
Author(s): Elliott Sober
Source: *Philosophy of Science*, Vol. 66, No. 4 (Dec., 1999), pp. 542-564
Published by: The University of Chicago Press on behalf of the Philosophy of Science Association
Stable URL: http://www.jstor.org/stable/188749
Accessed: 14/02/2011 16:18

The Multiple Realizability Argument Against Reductionism*

Elliott Sober†‡
Department of Philosophy, University of Wisconsin-Madison

Reductionism is often understood to include two theses: (1) every singular occurrence that the special sciences can explain also can be explained by physics; (2) every law in a higher-level science can be explained by physics. These claims are widely supposed to have been refuted by the multiple realizability argument, formulated by Putnam (1967, 1975) and Fodor (1968, 1975). The present paper criticizes the argument and identifies a reductionistic thesis that follows from one of the argument's premises.

1. Introduction. If there is now a received view among philosophers of mind and philosophers of biology about reductionism, it is that reductionism is mistaken. And if there is now a received view as to why reductionism is wrong, it is the multiple realizability argument.[1] This

*Received March 1999.

†Send requests for reprints to the author, 5185 Helen C. White Hall, Department of Philosophy, University of Wisconsin, Madison, WI 53706.

‡My thanks to Martin Barrett, John Beatty, Tom Bontly, Ellery Eells, Berent Enç, Branden Fitelson, Jerry Fodor, Martha Gibson, Daniel Hausman, Dale Jamieson, Andrew Levine, Brian Mclaughlin, Terry Penner, Larry Shapiro, Chris Stephens, Richard Teng, Ken Waters, Ann Wolfe, and an anonymous referee for this journal for comments on earlier drafts.

1. Putnam (1967, 1975) and Fodor (1968, 1975) formulated this argument with an eye to demonstrating the irreducibility of psychology to physics. It has been criticized by Lewis (1969), Churchland (1982), Enç (1983), and Kim (1989), but on grounds distinct from the ones to be developed here. Their criticisms will be discussed briefly towards the end of this paper.

The multiple realizability argument was first explored in philosophy of biology by Rosenberg (1978, 1985), who gave it an unexpected twist; he argued that multiple realizability entails a kind of reductionism (both about the property of fitness and also about the relation of classical Mendelian genetics to molecular biology). In contrast, Sober (1984) and Kitcher (1984) basically followed the Putnam/Fodor line. The former work argues that the multiple realizability of fitness entails the irreducibility of theo-

Philosophy of Science, 66 (December 1999) pp. 542–564. 0031-8248/1999/6604-0003$2.00

argument takes as its target the following two claims, which form at least part of what reductionism asserts:

(1) Every singular occurrence that a higher-level science can explain also can be explained by a lower-level science.
(2) Every law in a higher-level science can be explained by laws in a lower-level science.

The "can" in these claims is supposed to mean "can in principle," not "can in practice." Science is not now complete; there is a lot that the physics of the present fails to tell us about societies, minds, and living things. However, a completed physics would not thus be limited, or so reductionism asserts (Oppenheim and Putnam 1958).

The distinction between higher and lower of course requires clarification, but it is meant to evoke a familiar hierarchical picture; it runs (top to bottom) as follows—the social sciences, individual psychology, biology, chemistry, and physics. Every society is composed of individuals who have minds; every individual with a mind is alive;[2] every individual who is alive is an individual in which chemical processes occur; and every system in which chemical processes occur is one in which physical processes occur. The domains of higher-level sciences are subsets of the domains of lower-level sciences. Since physics has the most inclusive domain, immaterial souls do not exist and neither do immaterial vital fluids. In addition, since the domains are (properly) nested, there will be phenomena that lower-level sciences can explain, but that higher-level sciences cannot. Propositions (1) and (2), coupled with the claim of nested domains, generate an asymmetry between higher-level and lower-level sciences.

Reductionism goes beyond what these two propositions express. Events have multiple causes. This means that two causal explanations of the same event may cite different causes. A car skids off the highway because it is raining, and also because the tires are bald (Hanson 1958).

---

retical generalizations about fitness; the latter argues for the irreducibility of classical Mendelian genetics to molecular biology. Waters (1990) challenges the specifics of Kitcher's argument; much of what he says is consonant with the more general criticisms of the multiple realizability argument to be developed here. Sober (1993) defends reductionism as a claim about singular occurrences, but denies that it is correct as a claim about higher-level laws.

2. If some computers (now or in the future) have minds, then the reducibility of psychology to biology may need to be revised (if the relevant computers are not "alive"); the obvious substitute is to have reductionism assert that psychology reduces to a physical science. Similarly, if some societies are made of mindless individuals (consider, for example, the case of the social insects), then perhaps the reduction will have to "skip a level" in this instance also.

Proposition (1) says only that if there is a psychological explanation of a given event, then there is also a physical explanation of that event. It does not say how those two explanations are related, but reductionism does. Societies are said to have their social properties *solely in virtue of* the psychological properties possessed by individuals; individuals have psychological properties *solely in virtue of* their having various biological properties; organisms have biological properties *solely in virtue of* the chemical processes that occur within them; and systems undergo chemical processes *solely in virtue of* the physical processes that occur therein. Reductionism is not just a claim about the explanatory capabilities of higher- and lower-level sciences; it is, in addition, a claim to the effect that the higher-level properties of a system are determined by its lower-level properties.[3]

These two parts of reductionism are illustrated in Figure 1. The circled *e* represents the relation of diachronic explanation; the circled *d* represents the relation of synchronic determination. Reductionism says that if (x) explains (y), then (z) explains (y); it also asserts that (z) determines (x). The multiple realizability argument against reductionism does not deny that higher-level properties are determined by lower-level properties. Rather, it aims to refute propositions (1) and (2)—(z) does not explain (y), or so this argument contends.

2. Multiple Realizability. Figure 2 is redrawn from the first chapter, entitled "Special Sciences," of Fodor's 1975 book, *The Language of Thought*. It describes a law in a higher-level science and how it might be related to a set of laws in some lower-level science. The higher-level

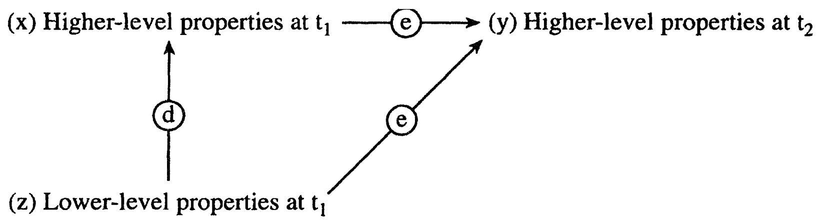

Figure 1. Relations of synchronic determination (*d*) and diachronic explanation (*e*) that may connect higher- and lower-level properties.

3. Reductionism should not be formulated so that it is committed to individualism of the sort discussed in philosophy of mind. For example, if wide theories of content are correct, then the beliefs that an individual has at a time depend not just on what is going on inside the skin of that individual at that time, but on what is going on in the individual's environment, then and earlier.

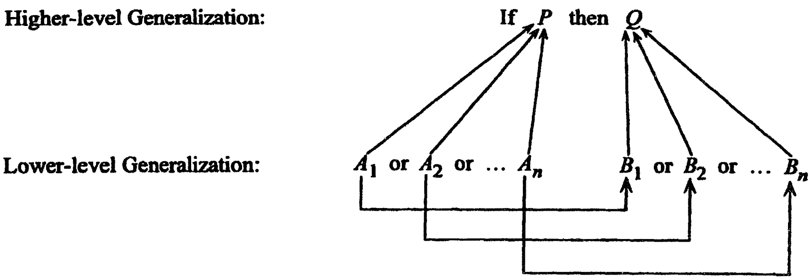

Figure 2. The lower-level properties A_i and B_j provide multiple realizations of the higher-level properties P and Q, respectively. One higher-level law and n lower-level laws are depicted, following Fodor 1975.

law is couched in its own proprietary vocabulary; P and Q are higher-level properties and the higher-level law says that everything that has P also has Q. The lower-level science provides n laws, each of them connecting an A predicate to a B predicate; the lower-level laws say that everything that has A_i also has B_i (for each $i = 1, 2, \ldots, n$).

The higher-level property P is said to be multiply realizable; $A_1, A_2, \ldots, A_n$ are the different (mutually exclusive and collectively exhaustive) realizations that P might have. Similarly, Q has $B_1, B_2, \ldots, B_n$ as its alternative realizations. What does multiple realizability mean? First, it entails the relation of simultaneous determination; necessarily, if something has A_i at time t, then it has P at t, and if it has B_i at time t, then it has Q at t. But there is something more, and it is this second ingredient that is supposed to ensure that the multiple realizability relation is anti-symmetric. An individual that has P has that property *solely in virtue* of the fact that it has whichever A_i it possesses. Because the higher-level properties are multiply realizable, the mapping from lower to higher is many-to-one. You cannot tell which of the A_i properties is exhibited by a system just from knowing that it has property P, and you cannot tell which of the B_j properties the system has just from knowing that it has Q.[4]

Two examples will make the intended meaning of multiple realizability sufficiently clear. Suppose that different types of physical system can have minds; minds can be built out of neurons, but perhaps they also can be built out of silicon chips. An individual mind—you, for

4. Although multiple realizability induces an asymmetry between P and each A_i, it does not entail that there is an asymmetry between P and the disjunctive property (A_1 or A_2 or . . . or A_n). Fodor would say that this disjunctive predicate fails to pick out a natural kind, a point that will be discussed later.

example—will have its psychological properties in virtue of the physical properties that the system possesses. But if you and someone else have some psychological properties in common, there is no guarantee that the two of you also will share physical properties; you and this other person may deploy different physical realizations of the same psychological properties. The same point can be made with respect to biological properties—you have various biological properties, and each of these is present in virtue of your possessing this or that set of physical properties. However, you and some other organism may share a given biological property even though you are physically quite different; this will be true if you and this other organism deploy different physical realizations of the same biological properties.

Since the multiple realizability relation obtains between simultaneously instantiated properties, the relation is not causal (assuming as I will that cause must precede effect). However, the diachronic laws I want to consider *are* causal—they say that a system's having one property at one time causes it to exhibit another property sometime later. The reason I will focus on causal diachronic laws is not that I think that all diachronic laws are causal, but that these provide the clearest cases of scientific explanations.[5] Thus, returning to propositions (1) and (2), we can ask the following two questions about the multiple realizability relations depicted in the second figure:

(1′) If an individual's having property P explains its having property Q, is it also true that its having property A_i explains its having property Q?
(2′) Do lower-level laws of the form "if A_i then B_i" explain the higher-level law "if P then Q"?

Let us assume that the properties described in higher-level sciences are multiply realized by properties discussed in a lower-level science. What consequences follow from this concerning reductionism?

3. The Explanation of Singular Occurrences—Putnam's Peg. Suppose a wooden board has two holes in it. One is circular and has a 1-inch diameter; the other is square and is 1 inch on a side. A cubical peg that is 15/16ths of an inch on each side will fit through the square hole, but not the circular one. What is the explanation? Putnam (1975) says that the explanation is provided by the *macro*-properties just cited of the peg and the holes. He denies that the *micro*-properties of molecules or atoms or particles in the peg and the piece of wood explain this fact.

5. Here I waive the question of whether *all* explanations are causal explanations, on which see Sober 1983 and Lewis 1986.

The micro-description is long and complicated and it brings in a welter of irrelevant detail. To explain why the peg goes through one hole but not the other, it does not matter what micro-properties the molecules have, as long as the peg and board have the macro-properties I mentioned. The macro-properties are explanatory; the micro-properties that realize those macro-properties are not. Hence, reductionism is false.

This is a delightfully simple example and argument, but it is possible to have one's intuitions run in the opposite direction. Perhaps the micro-details do not interest *Putnam*, but they may interest *others*, and for perfectly legitimate reasons. Explanations come with different levels of detail. When someone tells you more than you want to hear, this does not mean that what is said fails to be an explanation. There is a difference between explaining too much and not explaining at all.

Compare the micro-story that Putnam derides with a quite different story. Suppose someone suggested that the reason the peg goes through one hole but not the other is that the peg is *green*. Here it is obvious that a mistake has been made. If we demand that explanations be *causal* explanations, it will be quite clear why the color of the peg is not explanatory. It is causally irrelevant. This is an objective feature of the system under consideration and has nothing to do with our desire for brevity or detail.

It is possible to be misled by a superficial similarity that links the micro-story about the particles in the peg and board and the pseudo-explanation that cites the peg's color. Both of the following counterfactuals are true:

> If the particles in the peg and board had been different, the peg still would have passed through one hole but not the other, as long as the macro-dimensions were as described.
>
> If the peg had not been green, it still would have passed through one hole but not the other, as long as the macro-dimensions were as described.

If we say that causes are *necessary* for their effects (as does Lewis 1973a), we might be tempted to use these counterfactuals to conclude that the system's micro-features and the peg's color are both causally irrelevant, and hence should not be cited in a causal explanation. This proposal should be understood to mean that the effect would not have happened if the cause had not, *in the specific circumstances that actually obtained*; striking a match is not always necessary to get the match to light, but it may be necessary in various specific circumstances.

There are general questions that may be raised about the adequacy of this account of causation.[6] However, even if we waive these ques-

6. I will mention two. The first concerns how this theory of causation analyzes putative

tions, it is important to examine more closely how the counterfactual test connects with Putnam's argument. Let us suppose that the micro-properties of the peg and board's molecules are not necessary for the peg to go through one hole but not the other, if we hold fixed the macro-dimensions. But are the macro-dimensions necessary, if we hold fixed the micro-properties? That is, are we prepared to affirm the following counterfactual?

> If the macro-dimensions of the peg and board had been different, while the micro-properties were as described, the peg would not have passed through the one hole but not the other.

This counterfactual has a nomologically impossible antecedent. Many of us simply draw a blank when asked to assign a truth value to such assertions. The semantics of Stalnaker (1968) and Lewis (1973b) does not; it says that the counterfactual is vacuously true. However, before we interpret this as vindicating Putnam's argument, we also should note that the same semantic theory says that the following counterfactual is true as well:

> If the macro-dimensions of the peg and board had been different, while the micro-properties were as described, the peg still would have passed through the one hole but not the other.

It is hard to see how such counterfactuals can vindicate the judgment that the macro-properties are causally efficacious while their micro-realizations are not.[7]

I very much doubt that the concept of explanatory relevance means what Putnam requires it to mean in this argument. When scientists discover why smoking causes cancer, they are finding out which ingredients in cigarette smoke are carcinogenic. If smoking causes cancer, this is presumably because the micro-configuration of cigarette smoke is doing the work. If there turn out to be several carcinogenic ingredients and different cigarettes contain different ones, this does not make the molecular inquiry explanatorily irrelevant to the question of why people get cancer. The fact that *P* is multiply realizable does not mean that *P*'s realizations fail to explain the singular occurrences that

---

cases of overdetermination by multiple actual causes. Suppose Holmes and Watson each simultaneously shoot Moriarty through the heart. The theory entails that Holmes did not cause Moriarty's death, and Watson did not either. Rather, the cause is said to be disjunctive—Holmes shot him or Watson did. The second question comes from thinking about the possibility of indeterministic causation. Just as the totality of the antecedent causal facts need not suffice for the effect to occur, so the effect could have happened even if the causes had been different.

7. I am grateful to Brian Mclaughlin for drawing my attention to this line of argument.

P explains. A smoker may not want to hear the gory details, but that does not mean that the details are not explanatory.[8]

Putnam says he does not care whether we call the micro-story about the peg and the board a non-explanation, or simply describe it as a "terrible" explanation (Putnam 1975, 296). He thinks that the "goodness" of an explanation "is not a subjective matter." According to the objective concept of good explanation that Putnam has in mind, "an explanation is superior if it is more general" and he quotes with approval a remark by Alan Garfinkel—that "a good explanation is invariant under small perturbations of the assumptions" (301). What makes a more general (more invariant) explanation *objectively* better than one that is less? Putnam's answer is that "one of the things we do in science is to look for laws. Explanation is superior not just subjectively, but *methodologically*, in terms of facilitating the aims of scientific inquiry, if it brings out relevant laws" (301). My reply is that the goal of finding "relevant" laws cuts both ways. Macro-generalizations may be laws, but there also may be laws that relate micro-realizations to each other, and laws that relate micro- to macro- as well. Although "if P then Q" is more general than "if A_i then B_i ,"[9] the virtue of the

8. It is worth considering a curious remark that Putnam makes in a footnote before he introduces the example of the peg and board. He says:

> Even if it were not physically possible to realize human psychology in a creature made of anything but the usual protoplasm, DNA, etc., it would still not be correct to say that psychological states are identical with their physical realizations. For, as will be argued below, such an identification has no *explanatory* value in *psychology*. (1975, 293)

He then adds the remark: "on this point, compare Fodor, 1968," presumably because Fodor thought that antireductionism depends on higher-level properties being *multiply* realizable.

If we take Putnam's remark seriously, we must conclude that he thinks that the virtue of higher-level explanations does not reside in their greater generality. If a higher-level predicate (P) has just one possible physical realization (A_1), then P and A_1 apply to exactly the same objects. Putnam presumably would say that citing A_1 in an explanation provides extraneous information, whereas citing P does not. It is unclear how this concept of explanatory relevance might be explicated. In any event, I have not taken this footnote into account in describing the "multiple realizability argument," since Putnam's point here seems to be that *multiple* realizability does not bear on the claims he is advancing about explanation. This is not how the Putnam/Fodor argument has been understood by most philosophers.

9. I grant this point for the sake of argument, but it bears looking at more closely. Intuitively, "if P then Q" is more general than "if A_i then B_i" because the extension of P properly contains the extension of A_i. However, each of these conditionals is logically equivalent with its contrapositive, and it is equally true that the extension of not-B_i properly contains the extension of not-Q. This point is not a mere logical trick, to be swept aside by saying that the "right" formulation of a law is one that uses predicates that name natural kinds. After all, some laws (specifically, zero force laws) are typically

micro-generalization is that it provides more details. Science aims for depth as well as breadth. Some good explanations are fox-like; others are hedgehogian (Berlin 1953). There is no objective rule concerning which is better.

The claim that the preference for breadth over depth is a matter of taste is consistent with the idea that the difference between a genuine explanation and a nonexplanation is perfectly objective. In fact, it also is consistent with Hempel's (1965) view that the concept of scientific explanation should be explicated in terms of the notion of an ideally complete explanation, and that this is an objective notion. Perhaps an ideally complete scientific explanation of a singular occurrence in which an individual (or set of individuals) exhibits a multiply realizable property (or relation) would include the macro-story, the micro-story, and an account of how these are connected. If this is right, then reductionists and antireductionists alike are mistaken if they think that only part of this multilevel account deserves mention. But whatever the merits are of the idea of an ideally complete scientific explanation, we need to recognize that science in its currently incomplete state still is able to offer up "explanations." Perhaps these should be termed "explanation sketches," since they fall short of the Hempelian ideal. In any case, it remains true that science provides a plurality of such accounts of a given event. They vary in how detailed they are and in the level of organization described.[10]

Returning to Putnam's example, let us imagine that we face *two* peg-plus-board systems of the type that he describes. If we opt for the macro-explanation of why, in each case, the peg goes through one hole but not the other, we will have provided a *unified explanation*. We will have

---

stated as conditionals but their applications usually involve the predicates that occur in the contrapositive formulation. For example, the Hardy-Weinberg Law in population genetics describes how gamete frequencies will be related to genotype frequencies when no evolutionary forces are at work; its typical applications involve noting a departure from Hardy-Weinberg genotype frequencies, with the conclusion being drawn that some evolutionary forces are at work (Sober 1984). To say that the Hardy-Weinberg law has zero generality because every population is subject to evolutionary forces is to ignore the standard way in which the law is applied, and applied frequently, to nature.

10. Putnam's argument also has implications about the explanatory point of citing distal and proximate causes of a given effect. Imagine a causal chain from C_d to C_p to E. Suppose that C_d suffices for the occurrence of C_p, but is not necessary, and that the only connection of C_d to E is through C_p. Then Putnam's argument apparently entails that C_p explains E, and that C_d is either not an explanation of E, or is a terrible explanation of that event. But surely there can be an explanatory point to tracing an effect more deeply into the past. And surely it does not automatically increase explanatory power to describe more and more proximate causes of an effect.

explained similar effects by describing similar causes. However, if we choose a micro-explanation, it is almost inevitable that we will describe the two systems as being physically different, and thus our explanation will be *dis*unified. We will have explained the similar effects by tracing them back to different types of cause. Putnam uses the terms "general" and "invariant" to extol the advantages of macro-explanation, but he might just as well have used the term "unified" instead. In claiming that it is a matter of taste whether we prefer the macro- or the micro-explanation, I am claiming that there is no objective reason to prefer the unified over the disunified explanation. Science has room for both lumpers and splitters. Some people may not be interested in hearing that the two systems are in fact different; the fact that they have the same macro-properties may be all they wish to learn. But this does not show that discerning differences is less explanatory. Indeed, many scientists would find it more illuminating to be shown how the same effect is reached by different causal pathways.

In saying that the preference for unified explanation is merely a matter of taste, I seem to be contradicting a fundamental fact about scientific inference—that it counts in favor of the plausibility of a theory that the theory unifies disparate phenomena. Actually, no such consequence follows from what I am saying. Here, it is essential to distinguish the *context of justification* from the *context of explanation*.[11] When two theories are evaluated in the light of the evidence available, the fact that one is unified and the other is disunified is epistemologically relevant. In a wide range of circumstances, the unified theory can be expected to be more predictively accurate than the theory that is disunified, when they fit the data about equally well (Forster and Sober 1994). Whether a theory is unified is relevant to deciding whether we should accept it. However, the problem addressed by the multiple realizability argument is not about acceptance. We are supposed to assume that the macro-story and the micro-story are both *true*. Given this, we now are asked to decide which provides the better explanation of why the systems behave similarly. Unification is relevant to acceptance, but unification is not objectively relevant to deciding which accepted statements we should use in formulating explanations. The latter is simply a matter of taste—do we want more details or fewer? The context of justification and the context of explanation are different.

11. The distinction between justification and explanation was clearly drawn by Hempel (1965), who points out that why-questions can be requests for evidence or requests for explanation. This distinction supplements the familiar logical empiricist distinction between the *context of discovery* and the *context of justification*.

4. The Explanation of Laws—Fodor's Horror of Disjunctions. Whereas Putnam discusses the explanation of singular occurrences, Fodor uses the idea of multiple realizability to argue that laws in a higher-level science are not explained by laws in a lower-level science. This shift introduces some new considerations. Although many, if not all, explanations of singular occurrences are causal, the most familiar cases of explaining laws do not involve tracing effects back to their causes. Laws are usually explained by deriving them from "deeper" laws and initial condition statements; the explained laws and the explaining laws are true at the same time, so it is hard to think of the one as causing the other.

To understand Fodor's antireductionist position, let us consider the following derivation of a higher-level law:

> If A_i then B_i (for each $i = 1, 2, \ldots, n$).
> If A_1 or A_2 or . . . or A_n , then B_1 or B_2 or . . . or B_n.
> P iff A_1 or A_2 or . . . or A_n.
> Q iff B_1 or B_2 or . . . or B_n.
>
> ---
>
> If P then Q.

The first premise describes a set of lower-level laws; the second premise follows from the first. The third and fourth premises state bridge principles that connect a property discussed in a higher-level science with its multiple, lower-level, realizations. By assumption, the premises are true and the conclusion follows from the premises. Why, then, is this derivation not an explanation of the higher-level law?

Fodor's answer is not that the premises involve concepts that come from the higher-level science. Given that the higher-level science and the lower-level science use different vocabularies, any derivation of the one from the other must include bridge principles that bring those different vocabularies into contact (Nagel 1961). Rather, Fodor's reason is that laws cannot be disjunctive. Although he grants that each statement of the form "if A_i then B_i " is a law, he denies that the second premise expresses a law. For the same reason, the third and fourth premises also fail to express laws. To reduce a law, one must explain why the proposition is not just true, but is a law; this is supposed to mean that one must derive it solely from lawful propositions. This is why Fodor thinks that multiple realizability defeats reductionism.

Even if laws cannot be disjunctive, why does the above derivation fail to explain why "if P then Q" is a law? After all, the conclusion will be nomologically necessary if the premises are, and Fodor does not deny that the premises are necessary. Are we really prepared to say that the truth and lawfulness of the higher-level generalization is *inexplicable*, just because the above derivation is peppered with the word "or"? I confess that I feel my sense of incomprehension and mystery palpably subside when I contemplate this derivation. Where am I going wrong?

It also is not clear that laws must be nondisjunctive, nor is it clear what this requirement really amounts to. Take a law that specifies a quantitative threshold for some effect—for example, the law that water at a certain pressure will boil if the ambient temperature exceeds 100°C. This law seems to be disjunctive—it says that water will boil at 101°C, at 102°C, and so on. Of course, we have a handy shorthand for summarizing these disjuncts; we just say that any temperature "above 100°C" will produce boiling water. But if this strategy suffices to render the law about water nondisjunctive, why can't we introduce the letter α to represent the disjunction "A_1 or A_2 or . . . or A_n" and β to represent the disjunction "B_1 or B_2 or . . . or B_n"? It may be replied that the different disjuncts in the law about water all bring about boiling by the same type of physical process, whereas the different physical realizations A_i that the higher-level property P might have are heterogeneous in the way they bring about the B_i's that are realizations of Q.[12] The point is correct, but it remains unclear why this shows that laws cannot be disjunctive.

Disjunctiveness makes sense when it is understood as a *syntactic* feature of sentences. However, what does it mean for a proposition to be disjunctive, given that the same proposition can be expressed by different sentences? The problem may be illustrated by way of a familiar example. Suppose that the sentence "every emerald is green" and the sentence "every emerald is grue and the time is before the year 2000, or every emerald is bleen and the time is after the year 2000" are equivalent by virtue of the definitions of the terms "grue" and "bleen" (Goodman 1965). If laws are language-independent propositions of a certain type, and if logically equivalent sentences pick out the same proposition, then both sentences express laws, or neither does. Nothing changes if green is a natural kind whereas grue and bleen are not.

Although Fodor (1975) does not mention grue and bleen, it is fairly clear that his thinking about natural kinds—and his horror of disjunctions—both trace back to that issue.[13] Goodman (1965) held that law-

12. Fodor (1998, 16) says that a disjunction may occur in a bridge law if and only if the disjunction is "independently certified," meaning that "it also occurs in laws at its own level." The disjunction in the law about boiling presumably passes this test.

13. See, for example, Davidson's (1966) discussion of "all emeroses are gred" and also Davidson 1970.

like generalizations are confirmed by their positive instances, whereas accidental generalizations are not. The statement "all emeralds are green" is supposed to be lawlike, and hence instance confirmable, in virtue of the fact that "emerald" and "green" name natural kinds (or are "projectible"); "all emeralds are grue," on the other hand, is supposed to be non-lawlike, and so not confirmable by its instances, because it uses the weird predicate "grue." However, subsequent work on the confirmation relation has thrown considerable doubt on the idea that all and only the lawlike statements are instance confirmable (see, e.g., Sober 1988).

If P and (A_1 or A_2 or . . . or A_n) are known to be nomologically equivalent, then any probabilistic model of confirmation that takes that knowledge into account will treat them as *confirmationally* equivalent. For example, if a body of evidence confirms the hypothesis that a given individual has P, then that evidence also confirms the hypothesis that the individual has (A_1 or A_2 or . . . or A_n). This is a feature, for example, of Bayesian theories of confirmation (on which, e.g., see Howson and Urbach 1989 and Earman 1992). Disjunctiveness has no special meaning within that framework.

Fodor (1975, 21) concedes that the claim that laws must be nondisjunctive is "not strictly mandatory," but then points out that "one denies it at a price." The price is that one loses the connection between a sentence's expressing a law and the sentence's containing kind predicates. "One thus inherits the need for an alternative construal of the notion of a kind"; I am with Fodor when he says that he does not "know what that alternative would be like" (22). Fodor is right here, but his argument is prudential, not evidential. Like Pascal, Fodor is pointing out the disutility of denying a certain proposition, but this is not to show that the proposition is true.

The multiple realizability argument against the reducibility of laws is sometimes formulated by saying that the disjunctions that enumerate the possible realizations of P and Q are "open-ended." This would defeat the derivation described above—the third and fourth premises would be false—but it is important to see that the rules of the game now have changed. The mere fact that P and Q are multiply realizable would no longer be doing the work. And if the point about "open-endedness" is merely epistemological (we now do not *know* all of the physical realizations that P and Q have), it is irrelevant to the claim that higher-level sciences are reducible *in principle*.[14]

14. Moreover, the multiple realizability argument is not needed to show that the thesis of *reducibility in practice* is false; one can simply inspect present-day science to see this.

5. Probabilistic Explanations. The multiple realizability argument is usually developed by considering deterministic laws. However, laws in many sciences are probabilistic. How would the argument be affected by assuming that P and Q are probabilistically related, and that the A_i and the B_i are too?

Suppose that A_1 and A_2 are the only two possible realizations that P can have, and that B_1 and B_2 are the only two realizations that Q can have (the points I'll make also hold for $n > 2$). Suppose further that the probabilistic law connecting P to Q has the form

$$\Pr(Q \mid P) = p.$$

Then it follows that

$$p = \Pr(Q \mid P) = \Pr(Q \mid A_1)\Pr(A_1 \mid P) + \Pr(Q \mid A_2)\Pr(A_2 \mid P).$$

If we substitute $p_1 = \Pr(Q \mid A_1)$ and $p_2 = \Pr(Q \mid A_2)$ into this expression, we obtain

$$p = (p_1)\Pr(A_1 \mid P) + (p_2)\Pr(A_2 \mid P).$$

The probability (p) described in the higher-level law is a *weighted average* of the two probabilities p_1 and p_2; the weighting is determined simply by how often systems with P happen to deploy one micro-realization rather than the other.

It is not inevitable that $p = p_1 = p_2$. For example, suppose that smoking (P) makes lung cancer (Q) highly probable and that cigarette smoke always contains one of two carcinogenic ingredients (A_1 or A_2), which are found only in cigarette smoke. It can easily turn out that one of these ingredients is more carcinogenic than the other.[15] This means that there can be an important difference between higher-level and lower-level explanations of the same event—they may differ in terms of the probabilities that *explanans* confers on *explanandum*. To see why, let us add one more detail to the example. Suppose that lung cancer can be realized by one of two types of tumor (B_1 or B_2) growing in the lungs. Given this, consider an individual who has lung cancer. How are we to explain why this person has that disease? One possible reply is to say that the person smoked cigarettes. A second possibility is to say that the cancer occurred because the person inhaled ingredient A_1. Putnam's multiple realizability argument entails that the second suggestion is either no explanation at all, or is a "terrible" explanation. I suggest, however, that it should be clear to the unjaundiced eye that the second explanation may have its virtues. Perhaps A_1 confers on lung

15. If laws must be time-translationally invariant, then it is doubtful that "$\Pr(Q \mid P) = p$" expresses a law, if P is multiply realizable (Sober 1999).

cancer a different probability from the one entailed by A_2 ($p_1 \neq p_2$), and so the first account entails a different probability of cancer than the second ($p \neq p_1$). Furthermore, perhaps A_1 and A_2 confer different probabilities on the two tumors B_1 and B_2 and these tumors respond differently to different treatments. The additional details provided by the micro-explanation are not stupid and irrelevant. They make a difference—to the probability of the *explanandum*, and to much else.[16] Perhaps it is a good thing for cancer research that the multiple realizability argument has not won the hearts of oncologists.

6. Inference to the Best Explanation. I suspect that the multiple realizability argument has exerted so much influence because of a widespread misunderstanding concerning how *inference to the best explanation* works. The rough idea behind this mode of inference is that one should accept or reject hypotheses by deciding whether they are needed to explain observed phenomena. This inferential procedure seems to bear on the issue of reductionism as follows: We *now* need statements formulated in higher-level sciences because present day physics is not able to tell us how to understand societies, minds, and living things. However, if reductionism is correct, then these higher-level statements will not be needed once we have an ideally complete physics, and so they *then* should be rejected. But surely an ideally complete physics would not make it reasonable to reject all statements in higher-level sciences. This means that those statements must be needed to explain something that statements in an ideal physics could not explain. The multiple realizability argument presents itself as a diagnosis of why this is so.

This line of argument rests on a misunderstanding of inference to the best explanation. If you think that A_1 is one of the micro-realizations that P has, then you should not view "P causes Q" and "A_1 causes Q" as competing hypotheses (Sober 1999). The evidence you have may justify accepting both. Inference to the best explanation is a procedure that belongs to the context of justification. Once you have used that technique to accept a variety of different hypotheses, it is perfectly possible that your set of beliefs will furnish several explanations of a given phenomenon, each perfectly compatible with the others. Some of those explanations will provide more details while others will provide fewer. Some may cite proximal causes while others will cite causes that

16. This argument would not be affected by demanding that a probabilistic explanation must cite the positive and negative causal factors that raise and lower the probability of the *explanandum* (see, e.g., Salmon 1984). Cigarette smoke may raise the probability of lung cancer to a different extent than inhaling A_1 does, and so the two explanations will differ in important ways.

are more distal. The mistake comes when one applies the principle of inference to the best explanation a *second* time—to the set of hypotheses one *already* believes, and rejects hypotheses that one does not "need" for purposes of explanation. Inference to the best explanation is a rule for deciding what to believe; it is not a principle for retaining or eliminating beliefs that one already has perfectly good evidence for accepting. If hypotheses in higher-level sciences can be accepted on the basis of evidence, they will not be cast into the outer darkness simply because physics expands.

It is worth bearing in mind that the phrase "inference to 'the' best explanation" can be misleading. The hypothesis singled out in such inferences is not the best of all explanations (past, present, and future) that could be proposed; it is merely the best of the competing hypotheses under evaluation. Hypothesis testing is essentially a contrastive activity; a given hypothesis is tested by testing it *against* one or more alternatives (Sober 1994). When psychological hypotheses compete against each other, inference to the best explanation will select the best of the competitors; of necessity, the winner in this competition will be a psychological hypothesis, because all the competitors are. Likewise, when physicalistic explanations of a behavior compete against each other, the resulting selection will, of course, be a physicalistic explanation. It is perfectly consistent with these procedures that a given phenomenon should have a psychological *and* a physicalistic explanation. Both reductionists and antireductionists go wrong if they think that the methods of science force one to choose among hypotheses that, in fact, are not in competition at all.[17]

17. This point bears on an argument that Fodor (1998) presents to supplement his (1975) argument against reductionism. I am grateful to Fodor for helping me to understand this new argument. Fodor compares two hypotheses (which I state in the notation I have been using): (i) "if (A_1 or A_2 or . . . or A_n), then Q" and (ii) "if (A_1 or A_2 or . . . or A_n) then P (because the A_i's are possible realizations of P), and if P then Q." Fodor points out that the latter generalization is logically stronger (19); he then claims that it is sound inductive practice to "prefer the strongest claim compatible with the evidence, all else being equal" (20). Since we should accept the stronger claim instead of the weaker one, Fodor concludes that reductionism is false.

I have three objections to this argument. First, I do not think that the two generalizations are in competition with each other. If one thinks that the first conditional is true, and wants to know whether, in addition, it is true that the A_i's are realizations of P, then the proper competitor for this conjecture is that at least one of the A_i's is *not* a realization of P. Second, even if the two hypotheses were competitors, Fodor's Popperian maxim is subject to the well-known "tacking problem"—that irrelevant claims can be conjoined to a well-confirmed hypothesis to make it logically stronger. Fodor, of course, recognizes that *H&I* is not always preferable to *H*, ceteris paribus; however, he thinks that a suitably clarified version of the maxim he describes is plausible and that it will have the consequence he says it has for the example at hand. I have my

7. Two Other Criticisms of the Multiple Realizability Argument. The multiple realizability argument, when it focuses on the explanation of singular occurrences, has three premises:

> Higher-level sciences describe properties that are multiply realizable and that provide good explanations.
>
> If a property described in a higher-level science is multiply realizable at a lower level, then the lower-level science will not be able to explain, or will explain only feebly, the phenomena that the higher-level science explains well.
>
> If higher-level sciences provide good explanations of phenomena that lower-level sciences cannot explain, or explain only feebly, then reductionism is false.
>
> ---
>
> Reductionism is false.

I have criticized the second premise, but the first and third have not escaped critical scrutiny (see, e.g., Lewis 1969, Churchland 1982, Enç 1983, and Kim 1989; Bickle 1998 provides a useful discussion). I will consider these other objections separately.

Philosophers with eliminitivist leanings have criticized the first premise. They have suggested that if "pain," for example, is multiply realizable, then it probably does not have much explanatory power. Explanations that cite the presence of "pain" will be decidedly inferior to those that cite more narrow-gauged properties, such as "human pain," or "pain with thus-and-such a neural realization." Philosophers who advance this criticism evidently value explanations for being deep, but

---

doubts. It is illuminating, I think, to compare this inference problem to a structurally similar problem concerning intervening variables. If the A_i's are known to cause Q, should one postulate a variable (P) that the A_i's cause, and which causes Q? I do not think that valid inductive principles tell one to prefer the intervening variable model over one that is silent on the question of whether the intervening variable exists, when both models fit the data equally well (see Sober 1998 for further discussion). Third, even if the stronger hypothesis should be accepted in preference to the weaker one, I do not see that this refutes reductionism (though it does refute "eliminativist reductionism"). After all, the reductionist can still maintain that "if P then Q" is explained by theories at the lower level.

Notice that Fodor's argument does not depend on whether the A_i's listed are some or all of the possible realizations that P can have; it also does not matter whether the modality involved is metaphysical or nomological. Notice, finally, that this argument concerns inductive inference (the "context of justification," mentioned earlier), not explanation, which is why it differs from the argument of Fodor 1975.

not for being general. I disagree with this one-dimensional view, just as I disagree with the multiple realizability argument's single-minded valuation of generality at the expense of depth. Higher-level explanations often provide fewer explanatory details, but this does not show that they are inferior *tout court.*

It might interest philosophers of mind who have these worries about multiply realized psychological properties to consider the multiply realized properties discussed in evolutionary biology. In cognitive science, it is difficult to point to many present-day models that are well-confirmed and that are articulated by describing multiply realizable properties; this is mostly a hoped-for result of scientific advance. However, in evolutionary biology, such models are extremely common. Models of the evolution of altruism (Sober and Wilson 1998), for example, use the concept of fitness and it is quite clear that fitness is multiply realizable. These models have a useful generality that descriptions of the different physical bases of altruism and selfishness would not possess.

The third premise in the multiple realizablity argument also has come in for criticism. Perhaps *pain* is multiply realizable, but *human pain* may not be. And if *human pain* is multiply realizable, then some even more circumscribed type of pain will not be. What gets reduced is not pain in general, but specific physical types of pain (Nagel 1965). The multiple realizability argument is said to err when it assumes that reductionism requires *global* reduction; *local* reduction is all that reductionism demands. To this objection, a defender of the multiple realizability argument might reply that there are many questions about reduction, not just one. If human pain gets reduced to a neurophysiological state, but pain in general does not, then reductionism is a correct claim about the former, but not about the latter. If psychology provides explanations in which pain—and not just *human* pain—is an *explanans*, then reductionism fails as a claim about *all* of psychology.

Scientists mean a thousand different things by the term "reductionism." Philosophers have usually been unwilling to tolerate this semantic pluralism, and have tried to say what reductionism "really" is. This quest for univocity can be harmless as long as philosophers remember that what they call the "real" problem is to some degree stipulative. However, philosophers go too far when they insist that reductionism requires local reductions but not global reductions. There are many reductionisms—focusing on one should not lead us to deny that others need to be addressed.

8. A Different Argument Against a Different Reductionism. Although the multiple realizability argument against reductionism began with the arguments by Putnam and Fodor that I have reviewed, more recent appeals to multiple realizability sometimes take a rather different form. The claim is advanced that higher-level sciences "capture patterns" that would be invisible from the point of view of lower-level science. Here the virtue attributed to the higher-level predicate "*P*" is not that it *explains* something that the lower-level predicate "A_i" cannot explain, but that the former *describes* something that the latter does not. The predicate "*P*" describes what the various realizations of the property *P* have in common. The disjunctive lower-level predicate "A_1 or A_2 or . . . or A_n" does not do this in any meaningful sense. If I ask you what pineapples and prime numbers have in common and you reply that they both fall under the disjunctive predicate "pineapple or prime number," your remark is simply a joke. As a result, "if *P* then *Q*" is said to describe a regularity that "if (A_1 or A_2 or . . . or A_n) then (B_1 or B_2 or . . . or B_n)" fails to capture.

Whether or not this claim about the descriptive powers of higher- and lower-level sciences is right, it involves a drastic change in subject. Putnam and Fodor were discussing what higher- and lower-level sciences are able to *explain*. The present argument concerns whether a lower-level science is able to *describe* what higher-level sciences *describe*. I suspect that this newer formulation of the multiple realizability argument has seemed to be an elaboration, rather than a replacement, of the old arguments in part because "capturing a pattern" (or a generalization) has seemed to be more or less equivalent with "explaining a pattern" (or a generalization). However, there is a world of difference between describing a fact and explaining the fact so described. This new argument does not touch the reductionist claim that physics can explain everything that higher-level sciences can explain.

9. Concluding Comments. Higher-level sciences often provide more general explanations than the ones provided by lower-level sciences of the same phenomena. This is the kernel of truth in the multiple realizability argument—higher-level sciences "abstract away" from the physical details that make for differences among the micro-realizations that a given higher-level property possesses. However, this does not make higher-level explanations "better" in any absolute sense. Generality is one virtue that an explanation can have, but a distinct—and competing—virtue is depth, and it is on this dimension that lower-level explanations often score better than higher-level explanations. The reductionist claim that lower-level explanations are *always* better and the antireductionist claim that they are *always* worse are both mistaken.

Instead of claiming that lower-level explanations are always better than higher-level explanations of the same phenomenon, reductionists might want to demure on this question of better and worse, and try to build on the bare proposition that physics in principle can explain any singular occurrence that a higher-level science is able to explain. The level of detail in such physical explanations may be more than many would want to hear, but a genuine explanation is provided nonetheless, and it has a property that the multiple realizability argument has overlooked. For reductionists, the interesting feature of physical explanations of social, psychological, and biological phenomena is that they use the same basic theoretical machinery that is used to explain phenomena that are nonsocial, nonpsychological, and nonbiological. This is why reductionism is a thesis about the *unity* of science. The special sciences unify by abstracting away from physical details; reductionism asserts that physics unifies because everything can be explained, and explained *completely,* by adverting to physical details. It is ironic that "unification" is now a buzz word for antireductionists, when not so long ago it was the *cri de coeur* of their opponents.

To say that physics is capable in principle of providing a complete explanation does not mean that physical explanations will mention everything that might strike one as illuminating. As noted above, the explanations formulated by higher-level sciences can be illuminating, and physics will not mention *them*. Illumination is to some degree in the eye of the beholder; however, the sense in which physics can provide complete explanations is supposed to be perfectly objective. If we focus on *causal* explanation, then an objective notion of explanatory completeness is provided by the concept of *causal completeness*:

$$\Pr(\text{higher-level properties at } t_2 \mid \text{physical properties at } t_1 \,\&\, \text{higher-level properties at } t_1) = \Pr(\text{higher-level properties at } t_2 \mid \text{physical properties at } t_1).$$

To say that physics is causally complete means that (a complete description of) the physical facts at t_1 *determines* the probabilities that obtain at t_1 of later events; adding information about the higher-level properties instantiated at t_1 makes no difference.[18] In contrast, multiple

18. Let M = all the higher-level properties a system has at time t_1. Let P = all the physical properties that the system has at t_1. And let B = some property that the system might have at the later time t_2. We want to show that

$$\Pr(M \mid P) = 1.0$$

entails

$$\Pr(B \mid P) = \Pr(B \mid P \,\&\, M).$$

First note that $\Pr(B \mid P)$ can be expanded as follows:

$$\begin{aligned}\Pr(B \mid P) &= \Pr(B \,\&\, P)/\Pr(P)\\ &= [\Pr(B \,\&\, P \,\&\, M) + \Pr(B \,\&\, P \,\&\, \text{not-}M)]/\Pr(P)\\ &= [\Pr(B \mid P \,\&\, M)\Pr(P \,\&\, M) + \Pr(B \,\&\, \text{not-}M \mid P)\Pr(P)]/\Pr(P)\\ &= \Pr(B \mid P \,\&\, M)\Pr(M \mid P) + \Pr(B \,\&\, \text{not-}M \mid P)\end{aligned}$$

realizability all but guarantees that higher-level sciences are causally incomplete:

$$\Pr(\text{higher-level properties at } t_2 \mid \text{physical properties at } t_1 \ \& \ \text{higher-level properties at } t_1) \neq \Pr(\text{higher-level properties at } t_2 \mid \text{higher-level properties at } t_1).$$

If A_1 and A_2 are the two possible realizations of P, then one should not expect that $\Pr(Q \mid P \& A_1) = \Pr(Q \mid P \& A_2) = \Pr(Q \mid P)$ (Sober 1999).

Is physics causally complete in the sense defined? It happens that causal completeness follows from the thesis of simultaneous determination described earlier (Sober 1999). This fact does not settle whether physics *is* causally complete, but merely pushes the question back one step. Why think that the physical facts that obtain at a given time determine all the nonphysical facts that obtain at that time? This is a question I will not try to answer here. However, it is worth recalling that defenders of the multiple realizability argument usually assume that the lower-level physical properties present at a time determine the higher-level properties that are present at that same time. This commits them to the thesis of the causal completeness of physics. If singular occurrences can be explained by citing their causes, then the causal completeness of physics insures that physics has a variety of explanatory completeness that other sciences do not possess. This is reductionism of a sort, though not the sort that the multiple realizability argument aims to refute.

REFERENCES

Berlin, Isaiah (1953), *The Hedgehog and the Fox*. New York: Simon and Shuster.
Bickle, John (1998), *Psychoneural Reduction: The New Wave*. Cambridge, MA: MIT Press.
Churchland, Paul (1982), "Is 'Thinker' a Natural Kind?", *Dialogue* 21: 223–238.
Davidson, Donald (1966), "Emeroses by Other Names", *Journal of Philosophy* 63: 778–780. Reprinted in *Essays on Actions and Events*. Oxford: Oxford University Press, 1980, 225–227.
———. (1970), "Mental Events", in L. Foster and J. Swanson (eds.), *Experience and Theory*. London: Duckworth. Reprinted in *Essays on Actions and Events*. Oxford: Oxford University Press, 1980, 207–225.
Earman, John (1992), *Bayes or Bust?: A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.
Enç, Berent (1983), "In Defense of the Identity Theory", *Journal of Philosophy* 80: 279–298.
Fodor, Jerry (1968), *Psychological Explanation*. Cambridge, MA: MIT Press.
———. (1975), *The Language of Thought*. New York: Thomas Crowell.
———. (1998), "Special Sciences—Still Autonomous After All These Years", in *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA: MIT Press.
Forster, Malcolm and Elliott Sober (1994), "How to Tell When Simpler, More Unified, or

From this last equation, it is clear that if $\Pr(M \mid P) = 1.0$, then $\Pr(B \mid P) = \Pr(B \mid P \& M)$.

Less *Ad Hoc* Theories Will Provide More Accurate Predictions", *British Journal for the Philosophy of Science* 45: 1–35.

Goodman, Nelson (1965), *Fact, Fiction, and Forecast*. Indianapolis: Bobbs-Merrill.

Hanson, N. Russell (1958), *Patterns of Discovery*. Cambridge: Cambridge University Press.

Hempel, Carl (1965), "Aspects of Scientific Explanation", in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.

Howson, Colin and Peter Urbach (1989), *Scientific Reasoning: The Bayesian Approach*. La Salle: Open Court.

Kim, Jaegwon (1989), "The Myth of Nonreductive Materialism", *Proceedings and Addresses of the American Philosophical Association* 63: 31–47. Reprinted in *Supervenience and Mind*. Cambridge: Cambridge University Press, 1993.

Kim, Sungsu (unpublished), "Physicalism, Supervenience, and Causation—a Probabilistic Approach".

Kitcher, Philip (1984), "1953 and All That: A Tale of Two Sciences", *Philosophical Review* 93: 335–373. Reprinted in E. Sober (ed.), *Conceptual Issues in Evolutionary Biology*. Cambridge, MA: MIT Press, 1994, 379–399.

Lewis, David (1969), "Review of *Art, Mind, and Religion*", *Journal of Philosophy* 66: 22–27. Reprinted in N. Block (ed.), *Readings in Philosophy of Psychology*, vol. 1. Cambridge, MA: Harvard University Press, 1983, 232–233.

———. (1973a), "Causation", *Journal of Philosophy* 70: 556–567. Reprinted with a "Postscript" in D. Lewis, *Philosophical Papers*, vol. 2. Oxford: Oxford University Press, 1986, 159–213.

———. (1973b), *Counterfactuals*. Cambridge, MA: Harvard University Press. Revised edition 1986.

———. (1986), "Causal Explanation", in D. Lewis, *Philosophical Papers*, vol. 2. Oxford: Oxford University Press, 214–240.

Nagel, Ernest (1961), *The Structure of Science*. New York: Harcourt Brace.

Nagel, Thomas (1965), "Physicalism", *Philosophical Review* 74: 339–356.

Oppenheim, Paul, and Hilary Putnam (1958), "Unity of Science as a Working Hypothesis", in H. Feigl, G. Maxwell, and M. Scriven (eds.), *Minnesota Studies in the Philosophy of Science*, Minneapolis: University of Minnesota Press, 3–36.

Putnam, Hilary (1967), "Psychological Predicates", in W. Capitan and D. Merrill (eds.), *Art, Mind, and Religion*. Pittsburgh: University of Pittsburgh Press, 37–48. Reprinted as "The Nature of Mental States" in *Mind, Language, and Reality*. Cambridge: Cambridge University Press, 1975, 429–440.

———. (1975), "Philosophy and our Mental Life", in *Mind, Language, and Reality*. Cambridge: Cambridge University Press, 291–303.

Rosenberg, Alexander (1978), "The Supervenience of Biological Concepts", *Philosophy of Science* 45: 368–386.

———. (1985), *The Structure of Biological Science*. Cambridge: Cambridge University Press.

Salmon, Wesley (1984), *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

Sober, Elliott (1983), "Equilibrium Explanation", *Philosophical Studies* 43: 201–210.

———. (1984), *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. Cambridge, MA: MIT Press. 2nd edition, University of Chicago Press, 1994.

———. (1988), "Confirmation and Lawlikeness", *Philosophical Review* 97: 93–98.

———. (1994), "Contrastive Empiricism", in *From a Biological Point of View*. Cambridge: Cambridge University Press, 114–135.

———. (1998), "Black Box Inference: When Should an Intervening Variable be Postulated?", *British Journal for the Philosophy of Science* 49: 469–498.

———. (1999), "Physicalism from a Probabilistic Point of View", *Philosophical Studies* 95: 135–174.

Sober, Elliott and David S. Wilson (1998), *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.

Stalnaker, Robert (1968), "A Theory of Conditionals", in N. Rescher (ed.), *Studies in Logical Theory.* Oxford: Blackwell, 98–112.

Waters, Kenneth (1990), "Why the Antireductionist Consensus Won't Survive the Case of Classical Mendelian Genetics", *PSA 1990.* E. Lansing, MI: Philosophy of Science Association, 125–139. Reprinted in E. Sober (ed.), *Conceptual Issues in Evolutionary Biology.* Cambridge, MA: MIT Press, 1994, 402–417.