MIDWEST STUDIES IN PHILOSOPHY, XII

Causal, Experimental, and Structural Realisms

PAUL HUMPHREYS

1. DISCOVERY AND SCIENTIFIC EMPIRICISM

Our topic is scientific realism. It has been said that the opposition between realism and empiricism is old[1] and that is presumably why, if the pull of empiricism is sufficiently strong, realism may be hard to hold. This is a plausible claim in such areas as mathematical realism and moral realism. But the opposition is not a natural one, as is the opposition between empiricism and rationalism, and what is striking is that with all the talk of scientific realism, there is no corresponding talk of *scientific* empiricism. The principal thesis of this paper is that with the correct characterization of scientific empiricism, a minimalist scientific realism is supported by, rather than opposed to, scientific empiricism, and that the empiricism which is supposed to be at odds with scientific realism is an inappropriate kind of empiricism for science.

We must mark our place in the ever more complex network of realisms, antirealisms, and naturalistic denials of both. Much of the dispute over scientific realism has concerned, in one way or another, linguistic issues. The referential status of theoretical terms, the truth value of theoretical statements, the underdetermination of limit science, the correct demarcation between theoretical and observational terms, explanatory realism (when explanation is construed as a relation between sentences)—all these take the key feature of science to be its theorizing. Within this arena, the debate favors the antirealist. Yet science is more than evaluation of theoretical hypotheses. It is also, and this is prejudicial, a method of discovery, a means of discovering things about which, and of which, we should know nothing were it not for science.

One of the key features of discovery in any field is that it cannot be done in a passive way. Discoveries of new lands, of mathematical theorems, and of new elementary particles are not simply presented to us. Yet the traditional point of departure between realists and empiricist antirealists, which is at the border of the observable, carries with it connotations of phenomena merely impinging on our sensory receptors. Let us call this observable-oriented tradition of empiricism *passive empiricism*. From Hume's sensory impressions through the logical positivists' sense data to more recent observation vocabularies, humans are taken to be merely important kinds of receiving devices. Here then is a definition of passive empiricism's anti-scientific-realism: *Among the entities that constitute the focus of scientific investigation, only the observables can justifiably be claimed to exist.*

Quite banal, seemingly innocuous, and it allows passive empiricists to prescribe what is for them the legitimate ontology of specific objects and, construed purely empirically, laws of nature. Yet it is, I think, an exceptionally limited characterization of scientific empiricism, for the methods of scientific investigation are not limited to observation of the world.[2] It is the availability of other methods of discovery which enables us to concede to the passive empiricist that the burden of proof is on us, as realists, to establish that unobservables exist, yet to demonstrate that his terms are not exhaustive. This will involve a defense of some version of direct realism, direct in the sense that linguistic intermediaries are not essentially involved in the argument to realism, but not, I hope, a version of naive realism. An aside on scientific discovery is needed here.

Scientific discovery was given a philosophical bad name when placed in a context opposed to the context of justification, and condemned as being of merely psychological interest.[3] The discovery process is thus regarded as a mere preliminary to the central task of the justification, acceptance, rejection, and so forth of statements. Not only is this terminology a peculiar one given the linguistic mode within which it is usually presented, for one invents hypotheses or has ideas, one does not discover them, but its emphasis on theories is an example of why it is dangerous to identify the discovery issue with truth or assertability conditions. For discoveries are not made exclusively by means of true hypotheses. Interesting cases of falsification occur when experiments and observations made on the assumption of the truth of a hypothesis produce phenomena which were not expected on the basis of that hypothesis, by means of apparatus and data collection methods constructed on the assumption that the false hypothesis is true. If interesting cases of falsification occur, it follows that genuine discoveries can be made without the availability of true theories describing them, in that the apparatus and methods are robust enough to be used independently of the truth of the tested hypothesis. Because the novel falsifiers will generally be observable, this point should be acceptable to empiricists as well as realists. What we need to show is that discoveries often involve phenomena that take us beyond the observable.

When Humphry Davy discovered potassium and sodium by means of electrolysis, the contemporary theories of electricity and of potash and soda (the alkalies from which the metals were separated) were radically wrong.[4] Moreover, since the observed samples of the metals liberated by electrolysis were formed by continuous accretions of metallic particles smaller than the limit of observability, and had properties quite different from the materials used to produce them, we have here an example of substances literally moving across the famous gap between the unobservable and the observable. Of course, once they are out in the open, as it were, the substances uncontroversially exist for both realists and antirealists, but we can discover regularities and causes as well as objects by means of experiments. Indeed, we can, without too much distortion in the empiricist tradition, view regularities as involved in all cases of experimental discovery, in part because of the repeatability condition required of all sound experiments, in part because to say that an object is observable implies that some potential or actual regularity is associated with it.[5] Thus we can construe many, perhaps all, scientific discoveries as involving the observation of a regularity that was not previously present in circumstances other than those in which the discovery is made. We then ask the question "If these regularities are real (and the empiricist does allow that, for they are observed), why do we require experimental methods to observe them?" An empiricist can, of course, refuse the request for an explanation. Such a refusal would seem to me profoundly unphilosophical, but it can always be made. If so, I cannot proceed. But as long as a need for an explanation is seen, a realist explanation can be given that is simple, in keeping with scientific practice, and methodologically fruitful. Those are virtues which instrumentally inclined empiricists should appreciate, but the explanation involves essential recourse to unobservables.

2. ANTI-HUME[6]

What is the picture that Hume leaves us with after his devastating attack on the existence of necessary connections in nature? It is one of a world of discrete events presented to us as passive observers, assessed for properties such as spatial and temporal contiguity and regular temporal succession.[7] Taken literally, this is a very restrictive account of causation. If one takes the world as it comes, there are very few regularities of this kind—the world as it presents itself to us simply is not regular in the required way. Events rarely come to us in such a pristine fashion that we can just read off the causal relationships, even by habit. So a passive Humean empiricist is faced with a real problem. He can either say that because such regular sequences are a necessary condition for the existence of a causal law, there are in fact very few such causal laws in the world (and hence few causes and causally produced phenomena), or he will have to account for why, although the laws are present, the observed sequences are not manifest.

Those who hold that subsumption under regularities (even where not causal) is a prerequisite for scientific explanation will also have serious problems with this absence of regularities. For most phenomena that occur naturally, there will be no explanation. This difficulty can be dissolved in a number of ways, but I want to suggest that adopting a realist attitude toward causal laws and causal influences allows a unified account of how certain kinds of scientific discoveries are represented and employed.

In his book *A Realist Theory of Science*,[8] Roy Bhaskar notes this scarcity of regularities in open systems, and contrasts this with the fact that such regularities are common in experimental contexts. An activist Humean can thus provide himself with a number of causal laws, and explain phenomena in closed systems. But if causal laws are a subset of the class of regularities between observed events, why do the laws appear in closed systems but not in most open systems? It surely cannot be that the experimenter created the law by creating the experimental context, and destroyed the law in opening up the closed system. A simple nonscientific example adapted from Mill will illustrate this. Consider two tug-of-war teams—realists and antirealists—pulling on opposite ends of an inelastic rope. Equally balanced, the teams produce no observable movement. I now go to the rope and cut it with a razor. Observable effects occur. Question 1: Did the razor cut produce the law that covered the subsequent accelerations of the teams? Question 2: Was my razor cut the sole cause of the acceleration? Answers: no and no. (One may agree that the cut triggered the acceleration, but it cannot have produced the acceleration by itself. Ask the team sweating on the other end of the rope.) Hence the law and the other causes must have existed before the acceleration was produced, even though the other causes produced no observed changes before the cut.

Another reason why this response is unappealing is that it creates an extreme anthropomorphism. Most laws would be dependent for their existence on the existence of human experimenters to create the regularities. No humans, few laws. In consequence, our own existence would have evolved in a world almost free of biological (and other) laws, making it a good deal more surprising than it already is. This position is possible to maintain, but unattractive.

A second option open to an empiricist is to claim that the regularity was indeed already there in the nonexperimental context, but was obscured by the presence of other factors. The regularity was, in effect, embedded in statistical noise much as Michelangelo's sculptures were already present in the Carrara marble, waiting only to be freed by his chisel. We do, on this picture, literally discover causal laws. To discover is to bare, to uncover or expose to view, a meaning largely usurped in the epistemological realm by the wider notion of obtaining knowledge of something previously unknown.

This approach is a common one, especially in scientific and philosophical methodologies that use mathematical and statistical surrogates for actual experimental methods. The essence of this approach consists in partitioning existing data so that a subset of the data reveals a regularity that is not immediately evident in the whole set. Such methods are the essence of induction by simple enumeration, gambling systems (including stock-market forecasting by trend analysis), and correlation methods (including statistical relevance methods and regression techniques). The reverse of this method, leaving a data set with no regularities, forms the basis of many relative-frequency interpretations of probability. It is a quintessentially passive empiricist device, and its deficiencies are well chronicled. Without restrictions on the partitioning methods, or on what counts as a regularity, it is trivial to find some regularity in the data.[9] More important from our perspective is the fact that this approach cannot, by itself, provide any explanation of the presence of a regularity in one context, and its absence in another. The standard devices of restricting the scope of a regularity, or of imposing ceteris paribus conditions are merely descriptive acknowledgments of this fact.

The realist answer to this problem was, curiously, provided by J. S. Mill. In *A System of Logic*, Mill argued that the distinguishing characteristic of a causal regularity was its unconditionalness, this characteristic being exactly what is meant by a cause necessitating its effect: "That which is necessary, that which must be, means that which will be, whatever supposition we may make in regard to all other things."[10] Although Mill does not give an explicit argument for the unconditionalness requirement, the argument would seem to be the following. Suppose A is asserted to be a sufficient cause of B, yet there are circumstances X within which A is not followed by B. Then in those circumstances in which A is sufficient for B, it is not A *simpliciter* that is sufficient for (causes) B but A together with the absence of X or the absence of certain elements of X. Hence the original assertion regarding the causal generalization that A is a sufficient cause of B was incorrect. Note that this is not just an argument about how we speak of causal factors, but a condition on the very way in which such factors cause their effects. This doctrine of the unconditionalness of causal influences led Mill into serious difficulties when trying to reconcile empiricism with the unconditionalness requirement. Indeed, we shall see that the unconditionalness of causation is indispensable, and forces us to adopt a realist position with regard to causes.

Although Mill was a determinist, the unconditionalness requirement can be applied to causal analyses which do not require that a cause be sufficient for its effect, as long as those analyses have causal generalizations as their focus.[11] For if a causal factor must be necessary (not merely in the circumstances, but generally) for its effect, then because its absence is sufficient for the absence of the effect, a parallel argument to the above can be given. For probabilistic causation, the situation is more complex, for it requires us to adopt a realist attitude toward propensities, which would be begging the question here. However, a similar stricture of unconditionalness has to be imposed if one holds that probabilistic contributing causes are those which raise the propensity of an effect.

A number of writers, including Mill himself, realized that although the regularity requirement could be satisfied by creating experimental contexts, the lack of regularities in the natural world makes the unconditionalness requirement extremely difficult to satisfy. In particular, causal connections can be operating in the everyday world, but observable regular sequences are not identical with them. The natural inference to draw is that the causal sequences are associated with something unobservable, that there are causal influences "beneath" the observable phenomena.[12]

It is not enough to say this is the natural move to make. We need an argument. Here is one. Consider an experimental situation S in which a regularity R has been isolated, that is, one in which a single observed factor A is uniformly associated with a second observed factor E, that is, E regularly appears whenever A is present.[13] Then introduce into S a third factor B which, in S, in the absence of A, is uniformly associated with a fourth factor F. Now suppose that instead of insisting on the unconditionalness requirement, we claimed that a straightforward Humean regularity was sufficient, in the simple situation we have described (together with certain additional features such as temporal succession – what these are does not matter here), to identify A as a cause of E and B as a cause of F. Suppose further that neither E nor F is observed when both A and B are present, and that the situation is completely deterministic. Now ask what has happened to E. Why is it not present when B appears together with A? Now, as I mentioned earlier, it is possible for someone to deny that an explanation of this fact is called for. For him, there are three brute facts: situations with only A also have E present; situations with only B have F present; and situations with both A and B have neither E nor F. I assume in contrast that the burden of proof is always on those who deny that an explanation exists for a given fact. And the case we have in mind should be taken to be the most routine, everyday kind of situation, with no exotic quantum effects. The rope example without the razor will do.

If E was initially present, but disappeared when B was introduced into S, something must have been responsible. Let us call this "the feature that prevents A from causing E in S." We must establish that this feature, although existent, goes beyond what an empiricist would allow as real. What is this a feature of? If the feature that prevents A from causing E in S is a permanent feature[14] of B, then B would, when alone in S, simultaneously possess both the feature of causing F in S and the feature of preventing A from causing E in S. This second feature, given that it is permanent, cannot be (observably) instantiated, because it is false to claim that something has the (observably) instantiated feature of preventing A from causing E in S if A, by virtue of being absent from S at that time, does not have the feature of causing E in S. This permanent, but not always observably instantiated feature gives us the required realist entity. If, on the other hand, the feature of preventing A from causing E in S is a transient feature of B, then consider the possible cases when both A and B are present. Either A permanently possesses the feature of causing E in S, or it transiently does. If it permanently possesses it, we again have a realist feature, for it is not observably instantiated when B is present in S. It cannot transiently possess it when B transiently possesses the feature of preventing A from causing E, because in that situation A's feature is not observably instantiated and would hence be permanent, because A possesses it when alone in S and when accompanied by B. If A transiently possesses its feature when B's transient feature of preventing A from causing E is absent, then we have no explanation of why A's feature does not result in E. It is impossible for A transiently to lack its feature and for B transiently to possess its feature, for B cannot possess the feature of preventing A from causing E when A does not have the feature of causing E, any more than I can prevent a demagnetized tape from recording a message. Finally, if both A and B transiently lack their properties, then we again have no explanation for why A causes E when B is absent, but does not cause it when B is present, because the fact that A lacks the property of causing E is the fact that needs to be explained, and there is nothing in B or S to explain that fact.

Thus, even without presupposing the unconditionalness requirement, we must either allow the existence of permanent features that are sometimes not observably instantiated, or claim that there is no explanation for why the factor E is present in closed experimental systems and disappears when the system is opened, even minimally. The first option is, to my mind, preferable to the second.

I want to emphasize that this kind of argument is not antiempiricist in nature. It simply maintains that empiricism is not to be identified with the passive observation of events as they present themselves to us. We must explain what happens when the activist intervention in the natural course of events that experimentation allows produces a regularity that otherwise would not have occurred, and what happens to those regularities when other causal influences are added. The argument is applicable only in those cases in which the other causal factors do not permanently alter the causal properties of the original cause. For example, by grounding an electrically charged sphere, I remove the causal disposition to move the leaves of a gold-leaf electroscope apart. But we can identify which situations are of this kind by isolating the causal factor before and after its interaction with the other factors. If the same regularity reappears, we have a case to which the argument applies. The difference between the two types of situations is commonly acknowledged in ordinary discourse, as when the physician is enjoined to treat the disease rather than the symptoms. In the case where the original cause retains its dispositional properties, it is its effect or symptoms that are counteracted; in the case where the disposition is destroyed, it is the cause itself, or the disease, that is counteracted.

3. CONSEQUENCES

This account of scientific realism and scientific empiricism in terms of the discovery of causal factors via experimental isolation requires us to say something about how it affects other philosophical issues that involve realism and causation.

At the end of the last section, we established the existence of unobserved causal features. What the argument establishes is a general claim: that there must be causal influences operating in the world that are not to be identified with empirical regularities. It does not specify what they are, or even claim that our current theories correctly describe them. The issue remains of how much further we ought to go along the realist road. That, of course, depends upon the nature of the causal features, but the arguments given in this paper do not commit us to more than a minimalist realism.

Is it a defect that this approach enables us initially to say nothing about the properties of the causal features other than that they must exist unobserved in certain circumstances? I do not think so, because much of scientific progress consists of filling in detailed causal pictures about entities of which we initially know nothing except that they are responsible for certain observed effects. This construction of causal detail can be seen in a number of areas of scientific activity. It is seen in model building. Models being a simplification of reality, they frequently concern themselves initially with only a single causal component of a system, such as the models of a hydrogen atom. Perturbation theory then successively adds causal contributions to the previous stage of modeling. Even when interaction terms are required, the causal contributions of the existing components are generally left in place. If a causal component is removed, then that is simply an acknowledgment that the previous model was incorrect. The information about successive causal contributions may be acquired either experimentally or theoretically. The former is preferable, for one of the principal purposes of experimentation is to isolate the effects of single causal factors. However, in contexts where experimentation is impossible or impractical, resort to conceptual modeling may be required. We must, however, regard models in a Campbellian sense, within which causal influences are real, not representations. Semantic accounts of scientific theories are no better at representing causal realism than are syntactic approaches. Model theoretic structures are suitable for mathematical physics (and other sciences), not for experimental physics.[15]

Accounts of scientific explanation that closely associate prediction and explanation, that require true law statements under which the explanandum sentence falls, and that use a humean regularity account of causation have serious problems in explaining phenomena that are the result of multiple causal influences. This is primarily due to their inability to separate true explanations from complete explanations, because the lack of a comprehensive regularity covering the joint set of causes results in no explanation at all being available. Science, however, almost always proceeds in explaining phenomena by specifying a partial list of causes that contribute to and counteract the observed phenomena. Two previous papers detailed how this kind of explanation works in probabilistic cases.[16] For present purposes, the important fact is that in order to provide such explanations, we do not need a complete law under which the explanandum (not the explanandum sentence) falls; we require instead a set of a true laws covering the separate causal influences on the explanandum, and it is these which experimentation will give us.

This cumulative causal picture requires us to adopt a causal atomism rather than a logical atomism. I have deliberately refrained from construing the issue of scientific realism in terms of the truth of theories or the successful reference of theoretical terms. This semantic descent must be accompanied by a willingness to allow that the extension of the term 'observable' (I prefer 'detectable') is not fixed a priori, but is a temporally dependent function of scientific and technological advances. The problem with Russell's famous dictum "Wherever possible logical constructions are to be substituted for inferred entities," which underlies such enterprises as Ramsey's second-order eliminative approach, the first-order commitment criterion of Quine, and Craig's reaxiomatization method, is that it has become too tied to passive empiricism. (The insistence by Putnam that causation must be treated linguistically[17] seems to me to be a residue of this kind of approach.) We may, in initial phases of investigation, use a referential apparatus that employs a causal descriptive operator applied to a Ramseyfied sentence without the existential quantifier, so that we are talking of "the thing which is responsible for observed phenomenon O," but the moving boundary of the observable will frequently produce a much more direct presentation of the phenomena. Many advances in observation techniques involve procedures very similar to experimental techniques—the removal of causal factors that disturb the pure display of a cause. But we have to do the removing, and that is why I prefer another of Russell's aphorisms, to the effect that if you want to find out what exists, you must *go to the zoo*. To deny that this stripping away of confounding influences results in something observed is to deny the reality of entities such as viruses, atoms, synapses, neutrinos, status structures, and the cancer cells that caused the unfortunate death of David Hume.

The isolation of causal factors by means of experiment does not require us to adopt a manipulability account of causation, because as we have seen, we can treat the resulting causal display in the closed system purely in regularity terms. We do, however, require that the other potential causal influences be controllable, in that they can be held constant or removed entirely. If one does wish to view causation in manipulability terms, the argument sheds light on how we ought to view the proper role of such accounts. In order to impose experimental controls, we must already have a certain amount of causal knowledge. It is often claimed that manipulability accounts of causation are circular, because definitions of the causal relation in terms of manipulations require reference to causal properties in the definiens as well as the definiendum. That charge is correct if we are concerned to give an explicit definition of the general term 'causes'. For the present purposes of realism, however, that is not what is wanted—indeed, such an approach runs the risk of elevating causation to the status of a universal, with the consequent risk of having to account for what kind of thing a general causal necessity is. It is easy to make this error when employing a formal calculus of a causal necessity operator. The operator for logical necessity is usually treated as being the same for all propositions. But in the case of natural necessities, there is, if any such things exist, no reason to suppose that there is one monolithic necessity that brings about things; instead there is a variegated set of causal influences of very particular sorts. For the realist, why should there be something that pervades all actual causal connections between objects of radically different types? We already have names for these different kinds of influence: electromagnetic force, sexual attraction, economic demand, status expectations, and so forth. These may or may not be reducible to a single kind of causal factor, but that is a matter for science to discover, and if it were true, the basic influence would not be the second-level influence that natural-necessity advocates wish to establish.

In a similar vein, one should be wary of too quick an inference to the existence of 'causal powers'. If one holds that every disposition has a permanent or semipermanent basis, and that the causal features in the situations I discussed are dispositions or tendencies, then we may infer the existence of unobserved structures in such cases of multiple causal influences. The permanency of the structural underpinnings to the manifest regularities, together with the fact that such structures interact in different ways with different external variables, is enough to account for the dispositional properties possessed by such structures. Nothing in the realist positions requires that there exist an infinite regress of explanations in terms of finer-grained levels of reality, and which could only be stopped by reference to mysterious causal powers. There can be fundamental structures and processes that are *sui generis*, which are not caused by anything, which have no further explanation, and which operate unobserved in conjunction with other basic processes and structures to produce observed regularities. To that extent, the empiricist is correct in denying causal powers along with modal realism, but taking unobserved structures seriously is something which is enormously helpful even if one is primarily interested only in prediction of observed phenomena. This is because empiricist theories based only on observed phenomena will generally lose their predictive power if the internal structure of the system changes. Armed with a true account of that structure, a realist will be able to predict what effects those changes will have on the observed regularities.[18]

Finally, by not putting the issue in terms of theories, we avoid being placed in a position where the 'scientific' in 'scientific realism' requires us to provide a demarcation criterion between scientific theories and other kinds of theories. Given the difficulty of disposing of theoretical holism, and the endless debates about scientific rationality that surround the demarcation criterion, that is welcome. Nor need any mention be made of the aims or goals of science, or of limit science. It is often objected, for example, that arguments for scientific realism that are based on scientific explanation are unconvincing to many antirealists because they simply deny that explanation is a goal of science. It is true that one's methods are usually determined by one's goals. However, I take experimentation to be a given fact of the scientific enterprise. It is not that we philosophers suggest to scientists that they ought to experiment (as, for example, we might suggest they do some explaining instead of just predicting), it is a given feature of at least many of the natural sciences, and surrogates for it have been painstakingly constructed in many of the social sciences. Furthermore, although limit science is a proper concern of those interested in theoretical truth, and of those who are bothered by underdetermination issues, limit science is also an exhausted science, and of no interest to those for whom discoveries constitute the evidence for realism. The minimal realism advocated here does not *require* that we say anything true of the unobserved entities, except that they exist. But if we do liberalize the notion of observability, we are in a position successively to say much that is true of discoveries, and that at least allows us to be practicing realists.[19]

4. CONCLUSION

The preceding arguments have established a minimal kind of scientific realism. They do not, to be sure, give us the kind of conclusions that arguments based on scientific theories aimed to provide. But they do, I think, provide us with a specifically scientific ontology, a class of entities of which we should not know were it not for science. And I am not convinced that realism requires much knowledge. When Leif Eriksson discovered North America, he knew virtually nothing of it upon leaving, a fact that is all too evident from *The Vinland Sagas*. But its existence had been established, and the course of scientific discovery is not entirely dissimilar, I think, to geographical discovery.[20]

Notes

1. B. van Fraassen, *The Scientific Image* (Oxford, 1980), 1.

2. An obvious but necessary point: it may well be that at the most fundamental epistemological level, observation is all we ever have. But philosophy of science is not fundamental in that sense—it has to take for granted the existence of at least some objects of scientific knowledge, the peculiarity of scientific method, and the addition of special techniques to supplement observation, in order to distinguish science from other kinds of investigation. These techniques, and those methods, may ultimately be telling us something very different from what scientists think they are telling us, but that there exists a scientific enterprise engaging in more than observation has to be an uncontroversial claim.

3. As in H. Reichenbach, *Experience and Prediction* (Chicago, 1938), 5–7; K. Popper, *The Logic of Scientific Discovery* (London, 1959), 31.

4. The article on Davy in *The Dictionary of Scientific Biography*, edited by C. C. Gillespie, claims that Michael Faraday listed a dozen incompatible theories of electrolysis, all of which he claimed could be derived from Davy's views. For Davy's work, see H. Hartley, *Studies in the History of Chemistry* (Oxford, 1971), Chapter 4. Further examples showing that discoveries can be made without access to true theories are given in Part B of Ian Hacking's *Representing and Intervening* (Cambridge, 1983). One curious feature of this piquant book is its fascination with traditional issues involving observations.

5. In everything that follows, "experiment" is to be taken in the sense of controlled laboratory experimentation, in which all but one of the factors are removed or held constant. The use of randomized experiments, statistical surrogates for experimental controls, and quasi-experimentation all introduce additional complexities into the interpretation of causal laws, even though certain kinds of regularities may be discovered by means of them.

6. The arguments in this section have drawn heavily on ideas and arguments contained in Roy Bhaskar's *A Realist Theory of Science* (Atlantic Highlands, 1975). It was Bhaskar's book which convinced me that the correct way to argue for scientific realism must be by assessing the role played by experiment in causal discoveries. Although I depart from his views about the social basis of experimentation and the nature of tendencies, I believe that his book remains the best written on this topic.

7. Hume does use the term "experiment" in Book II, Part II, section II of the *Treatise*, and also in scattered places in Book I, Part III (especially section VIII). In all these references, however, he seems to take the term to be synonymous with "observation."

8. See note 6.

9. The trivialization was highlighted by Russell in his discussion of functional determinism in "On the Notion of Cause," reprinted in his *Mysticism and Logic* (London, 1917).

10. J. S. Mill, *A System of Logic*, 8th edition (London, 1881) III, v, 1–3, 6.

11. These arguments are inapplicable to singular causal claims that are not derived from causal regularities. By building "A caused B, in the circumstances" into the claim, or by relativizing the claim to a background field, the conditionality of the causal claim is thereby introduced. Excluding such cases

will not affect our argument for realism, in part because it is a requirement on experiments that they be freely repeatable.

12. Mill was caught in a dilemma between remaining faithful to an empiricist regularity account of causation and adopting a position involving tendencies to produce effects. See Mill, *A System of Logic*, III, v, 1–3, 6, for the first view; III, x, 5, and III, v, 5, for the second view. This last section is not in the first edition, and contains Mill's somewhat obscure views on dispositions: He says that 'capacities' are not real, but are merely names for inferential habits. Whether this humean line applies to tendencies, which are active capacities, is not stated. Mill's dilemma is discussed by Geach in *Three Philosophers*, by G. E. M. Anscombe and P. Geach (Ithaca, 1961), 101–4; chapter 4 of Alan Ryan's *John Stuart Mill* (New York, 1970); Roy Bhaskar, *A Realist Theory*; and Essay 3 of Nancy Cartwright's *How the Laws of Physics Lie* (Oxford, 1983).

13. Take 'factor' to represent whatever you consider the causal relata to be: event types or tokens, changes in variables, states, conditions, etc.

14. The term 'feature' is used rather than 'property' to avoid unnecessary worries about the status of properties. Even if one does not want to allow properties as real, there are many features of the world that uncontroversially exist for empiricists.

15. In his "Galilean Idealization" paper (*Studies in the History and Philosophy of Science* 16 [1985]: 247–73), Ernan McMullin has traced the development of such modeling methods. In his "Structural Explanations (*American Philosophical Quarterly* 15 [1978]: 139–47) and "A Case for Scientific Realism" (in *Scientific Realism*, edited by J. Leplin [Berkeley, 1984], 8–40), he had already argued on philosophico-historical grounds for adopting a position on scientific realism of the kind I adopt in the present paper. I hope the arguments here complement those in his papers. Further examples of the relation between causal modeling and experiment can be found in the papers by Blalock, Duncan, Glymour and Scheines, and by Cook and Campbell in *Synthese* 68 (1986): 1–188; and from a different perspective, in P. Humphreys and J. Berger, "Theoretical Consequences of the Status Characteristics Formulation," *American Journal of Sociology* 86 (1981): 953–83.

16. See my "Aleatory Explanations," *Synthese* 48 (1981): 225–32, and "Aleatory Explanations Expanded," in *PSA 1982*, Vol. 2, edited by P. Asquith and T. Nickles (East Lansing, 1983), 208–223.

17. "Why There Isn't A Ready-Made World" and "Introduction," both in H. Putnam, *Realism and Reason, Philosophical Papers*, Vol. 3 (Cambridge, 1983); "Is the Causal Structure of the Physical Itself Something Physical?" *Midwest Studies in Philosophy* 9 (1984): 3–16.

18. A detailed argument for this claim, with illustrative examples, is given in P. Humphreys, "Quantitative Probabilistic Causality and Structural Scientific Realism" in *PSA 1984*, Vol. 2, edited by P. Asquith and P. Kitcher (East Lansing, 1986).

19. A quite different set of arguments for minimal realism has been offered by Robert Almeder in his "Blind Realism," *Erkenntnis* (Dec. 1986): 1–45. There he concludes that although some of our warranted beliefs about unobservable entities must be true, we cannot pick out which, among our current set of beliefs, those are. His discussion of Pierce's experimental demonstration that laws of nature are mind-independent brings out very clearly a key aspect of experimental, as opposed to theoretical, realism—that there are features of the world over which we have no mental control.

20. A very preliminary version of this paper was read at the University of Pittsburgh in May 1985. I thank Ernan McMullin for his supportive criticisms at that time; also Nicholas Rescher and Robert Almeder for many lengthy discussions on this issue. The work on which this paper is based was supported by NSF grant SES 8410898.