On Scientific Method, Induction, Statistics, and Skepticism

Abraham D. Stone (abestone@uchicago.edu)
University of Chicago

31 July, 2003

Abstract. This paper is part of an attempt to explain the distinctive status of modern science within human intellectual life, and in particular the way in which it seems to undermine the legitimacy of other, traditional modes of thought. I argue that the key feature of modern scientific methodology is its testing, not of our judgments, but of the concepts employed in them. Statistical techniques represent a systematic way of doing this. I show, however, that something similar happens on a larger scale in so-called scientific revolutions. (Hence I argue that there is no discontinuity between the methods of "normal" and "revolutionary" science.)

My aim in this paper is to explain how universal statements, as they occur in scientific theories, are actually tested by observational evidence, and to draw certain conclusions, on that basis, about the way in which scientific theories are tested in general. [1] But I am pursuing that aim, ambitious enough in and of itself, in the service of even more ambitious projects, and in the first place: (a) to say what is distinctive about modern science, and especially modern physical science, as a human intellectual activity; and (b) to show how this distinctiveness explains the unique status of modern science in human intellectual life. So I will begin by saying a few words about that larger project.

One might doubt, first, whether that project is legitimate. Although everything is different from everything else, the question "What is distinctive about X?" is not necessarily well put, because X may not be anything—that is, anything distinct. And there are indeed many philosophers, otherwise of the most diverse intellectual backgrounds and tendencies, who would deny, in various ways, that modern science is a distinct thing, or that its status is unique. It seems to me that they deny something terrifyingly obvious—a fact which confronts us far more urgently than the fact that, say, ravens are black. I will put off further remarks about this until the end of the paper, however, because the discussion of my more limited present aim will focus precisely on the ways to tell when a question is well put, and whether there is such a (distinct) thing as X. That one's methodological problems are also the subject of one's investigation is a sign that that investigation is philosophical, although (or because) also a threat to its coherence.

Second, some points about methodology. I am not a sociologist, or even a historian, nor (unlike some recent philosophers of science) will I pretend to be

$Id: scimethod.tex,v 5.6 2003/10/17 15:11:53 abestone Exp $

one. A reliable and philosophically neutral history and sociology of science is very much to be desired, if indeed such things are possible in principle. But my interest is in (a) and (b) as logical or epistemological points, rather than historical or sociological ones. To claim that (a), in the way I am taking it, is connected with (b), is thus to claim that modern science has a unique status not, or not only, de facto, but also de jure. I have not yet said anything about what that unique status is. But it should be obvious that it is at least in certain ways a *dominant* one: that modern science claims and exercises the rights of a lord. The project of showing that those rights are not violently taken but deserved is therefore similar to the projects of various philosophers who have tried to justify or explain the rationality of scientific practice. I am not embarrassed by such company. I do, however, want to distinguish my approach from certain of theirs with which it might be confused.

The classic issue in this area was to justify the principle of induction—i.e., to explain why we are entitled to conclude from observed to unobserved instances of a regularity. Some have asserted that this problem was always silly on the face of it, and that it's time we moved on to other things. With that point of view I am not much in sympathy. But others have suggested various modifications of the problem, and I do agree with many of these. For one thing, it is the method, rather than the results, of science that should be justified: physical science in its present state does not offer any guarantee that its results are true. This justification, furthermore, cannot be absolute, but only relative to other possible (human) methods. And finally, we ought not to expect an explanation of how science chooses the best theory from among all possible theories, since we cannot even formulate, let alone seriously entertain, all possible theories. What we can expect is to understand how science tests for acceptability some statements or theories that are actually proposed.

These modifications have been widely accepted. A further, more controversial, modification is represented by the view of Karl Popper: that "induction" as described above is neither particularly justifiable nor an actual feature of practice—not, that is, insofar as that practice is really scientific.[2] With this, too, I am in agreement—with the first part, completely, and with the second, insofar as modern scientific practice is concerned. As to the nature of the proper and (relatively) justifiable method, however, I do not agree with him, and the body of this paper will put forward my alternative suggestion.

My questions, moreover, are specifically about modern science. Both the "modern" and the "science" bear emphasis. The methodological features we will be discussing cannot be traced back to Anaximander, and, while there are other modern practices (most crucially of textual interpretation) which exhibit similar features, this is not generally true of disciplines outside of science. The most important example, for us, of such a discipline is philosophy. It is true that philosophy from the beginning sets up a certain standard (of criticality), and that modern science above all meets that standard, as philosophy itself does not. And it is also true, for that very reason, that modern philosophers have repeatedly proclaimed the merger, or rather re-merger, of philosophy and science. Still, modern science and philosophy remain stubbornly different. I hope, indeed, to say something useful about what that difference, and that stubbornness, are.

Finally, the term "justification," applied to my project here, might suggest, misleadingly, that I am out to prove the rationality of modern science to someone who does not accept it. I do not say that that is impossible. On the contrary: as modern intellectual history amply demonstrates, such persuasion is, at least in favorable circumstances, pretty much inevitable, and proceeds whether or not philosophy deigns to argue on science's behalf. What philosophy does need to justify is that persuasiveness itself: the seeming inevitability with which we accept modern science as relatively rational, or, in other words, with which other, traditional modes of thought lose their legitimacy in the face of it. This is the unique status of modern science to which I alluded under heading (b). In relation to the results of science, it concerns, not their truth, but their irresistible and irreversible character: that they cannot, that is, be resisted or reversed in favor of tradition (for they *can* be resisted or reversed on science's own terms, or in the type of pseudo-traditional reaction which we loosely call "fundamentalism").

Irresistibility and irreversibility are things which I have learned to associate with modernity in general—and hence with modern science, which is paradigmatically modern—from Stanley Cavell. From the same source, I have learned to associate them with skepticism. To mention them (and him) here at the beginning is thus to suggest certain limitations on the kind of justification we should want in this area, or certain qualifications on the way we should feel about such justification. In particular, it is to suggest that the success of a project such as mine should not be greeted with triumphalist joy. Philosophical epistemology operates with two constraints: a theoretical and a practical one. The theoretical constraint is that knowledge, whatever it is, must be something *justified*, and it is to that constraint that one responds in explaining the uniqueness of modern (scientific) rationality. But the practical constraint is: that I must always know enough to do my duty. An irresolvable conflict between these two is both terrible to contemplate and unfortunately all too plausible. But there would be little point in saying more about such large and obscure questions yet, while we still have much detailed work ahead of us.

1. "Induction": Two Examples

No doubt it is a stylistic error to enclose one's section heading in scare quotes. In the present instance it seems to be necessary, however. For the topic of this section will be the way in which modern science confirms (or disconfirms) universal statements such as

(i) All ravens are black.

Such statements are traditionally supposed to be verified by induction, and induction is taken to be a process in which, after observing some state of affairs over and over (in this case: after observing many black ravens, and none that are not black), one concludes that it occurs universally. That, I will argue, is an inaccurate description of modern scientific methodology, and one which leaves out precisely its distinctive feature. Thus far, as promised, I agree with Popper. But should one say that science does not proceed by induction, or rather that traditional accounts of induction are incorrect? Unlike Popper, I do not regard such terminological questions as trivial or unimportant. Hence the scare quotes.

The statement (i) is a traditional example, and there is always good reason to use traditional examples, whenever possible. But though it *will* prove possible discuss a hypothetical scientific project of verifying (i), it will not be easy. The problem is that (i) is not, in the current state of our practice, part of any scientific theory, but at best part of the data for which some theory (perhaps a theory of evolutionary or of molecular biology) might have to account. It is therefore not easy to see how scientists would go about testing the truth of (i), if there were or had ever been such a research program.

Let me make clear what I do and do not mean by this complaint. Empiricists used to claim or hope for a distinction between theoretical statements and those which are intrinsically prior to all theory ("observational"). That claim or hope has, thanks to the efforts of Duhem, Popper, Neurath, Putnam and others, fallen into disfavor, and rightly so.[3] But though any statement could in principle become the object of scientific testing, there is in practice a large and crucial difference between the things that a theory proposes and those which it need only explain or, at the least, not contradict. Now, it may well be that we or our ancestors have carried out some crude or unsystematic or unconscious version of the scientific testing of (i), and that we owe our (alleged) knowledge of its (partial) truth to such a process. But certainly (i) is not and has never been the object of explicit testing in modern scientific research. It is difficult even to imagine its being so. Yet we will have to imagine just that, if we are to see what kind of evidence would then be brought to bear, and how.

Let me clarify here even further. Perhaps someone will now think (if only because of certain recent fashions in the philosophy of science) that my complaint concerns the use of an imaginary example, rather than a real one. That is very far from my intention. The problem is not, per se, that we need to imagine a fictional case: fiction, rather, as we should know from Aristotle (and, yes, even science fiction), is more philosophical than history. If there can be such a thing as philosophy of science at all, then there is more we can learn from a single good fiction than from the whole actual history of science combined. But nothing good comes cheaply: imagination is difficult. Let us therefore *start* with something a bit easier—with a statement which, more or less, actually does form a part of a modern scientific theory. I have chosen the following, in part because it sounds very similar to (i):

(ii) All post-main-sequence stars are red.

Like (i), (ii) is not precisely correct, both in that there are exceptions (some ravens are albino or painted) and in that it is not completely or absolutely true even in the normal cases (a raven is not black all over, and even the "black" parts of its surface are not absolutely black—just very dark gray). But, like (i), (ii) is relatively true in almost all cases. Or so we think. In the case of (ii), however, unlike in the case of (i), we think so because it has been tested in the process of modern scientific research.

What is the empirical evidence for (ii)? In a way, that is a trick question. The phrase "empirical evidence" makes one imagine that someone went out and made some observations—observed a lot of post-main-sequence stars, say, and noted that they were all red—and that we believe (ii) because of those observations. The truth of the matter—and this should be a familiar point—is nothing like that simple. The statement (ii) is part of a complicated scientific theory, the theory of stellar structure and evolution, which we believe to be a (more or less) correct theory for any number of reasons. In a sense, there is no empirical evidence whatsoever which is evidence simply for (ii) as such. But recall: we are interested in how proposed statements are scientifically tested, rather than in how they emerge out of all possible interpretations of the data. And there *is* a kind of empirical evidence which constitutes a *test* of (ii). Consider the (imaginary and highly simplified) data plotted in fig. 1, which is known as a Hertzsprung-Russell diagram. The vertical axis represents the absolute visual magnitude of a star—a measure, roughly speaking, of its luminosity. The horizontal axis represents a star's color, measured by taking the difference between its visual magnitude as seen through two different colored filters. The stars in the narrow band running from the upper left to the lower right are (mostly) on the main sequence; the more diffuse group of stars on the upper right are post-main-sequence.

From the diagram one can see part of the reason for saying that (ii) is only relatively true. Many of the post-main-sequence stars are not as red as many of the main-sequence stars; some, in fact, are not as red as our sun, a star which we would hardly call "red" in absolute terms. What is (mostly) true, however, is that post-main-sequence stars are redder than main-sequence stars of the same luminosity—or, to put it differently, that they are brighter than main-sequence stars of the same color. If we assume that stars are more or less spherical and more or less uniform black-body emitters, then their color is associated with their surface temperature and hence with their luminosity per unit surface area, and two stars of the same color which differ in luminosity

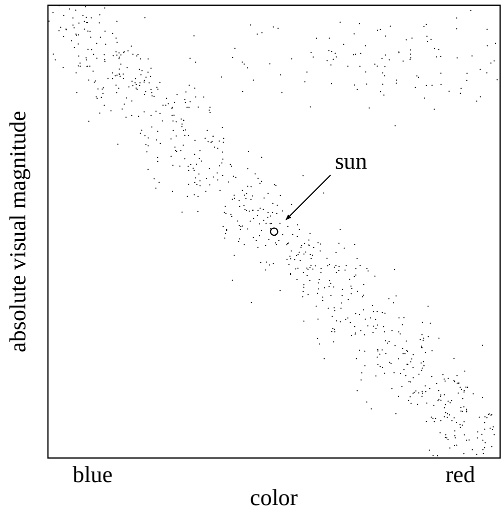

Figure 1. Hertzsprung-Russell diagram (confirming).

must therefore differ in total surface area, i.e. in size. Hence it would be perhaps less misleading, in place of (ii), to say that all post-main-sequence stars are bright, or, equivalently, that all post-main-sequence stars are large. (In fact these stars are normally said to belong the "luminosity classes" of "giant" and "supergiant.") Still, (ii) is more or less accurate, properly understood, and has in its favor also its similarity to the traditional (i). Let us therefore continue with (ii).

There is a classic debate between Popper, on the one hand, and others (especially Putnam and Kuhn) on the other: Popper maintains that no instance can ever serve to confirm (even partially) a universal statement, whereas his opponents maintain that no instance can ever serve to falsify it. We should note right away that, at least as far as (ii) is concerned, and in relation to the data shown in fig. 1, both sides are correct. No one of the single observations recorded in fig. 1 can be said to confirm (ii), contribute to its verification, or make it more probable. No such activity is represented here as examining many stars, finding that all the post-main-sequence ones are red, and then generalizing to the universal statement (ii). For the same reason, moreover, there is no single observation we could add to fig. 1 which would serve to falsify or disconfirm (ii). That is: fig. 1 also does not record a process of examining many stars in search of a post-main-sequence one which is *not* red. On the contrary: we can only know (at least, from the data presented here) that a given star is post-main-sequence by the very fact that it is (relatively) red.

And yet fig. 1 does provide a striking confirmation of (ii) and, by way of (ii), of the theory of stellar evolution of which (ii) is a part. And there could well be data like those of fig. 1 which would serve to undermine (ii) and

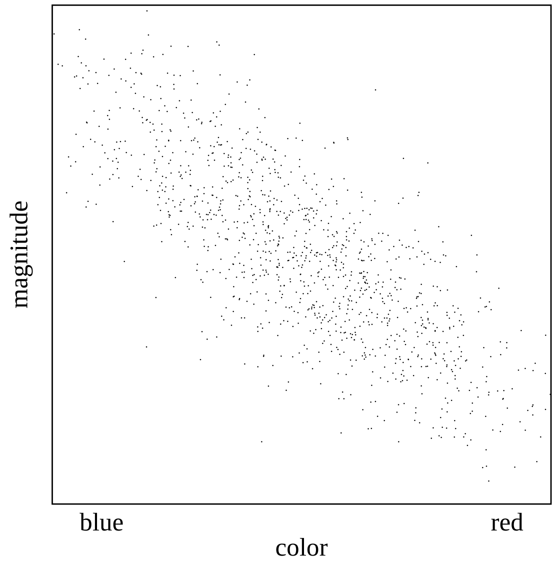

Figure 2. Hertzsprung-Russell diagram (undermining).

thereby to disconfirm that theory. Not that the theory could tolerate no departures whatsoever from the pattern shown. Observation of particular samples of stars (for example, from very young or very old clusters) would yield quite different patterns, and the theory is able to account for that. But if, in the long run, nothing at all similar to fig. 1 were obtained—if, for example, the data always tended to come out like those in fig. 2—then it is safe to say that the theory would be in hot water. As for (ii), it is perhaps not quite right to say that it would be "falsified" or "disconfirmed"—that is why I instead chose to use the vaguer term "undermined." What would emerge, in such a case, is not that (ii) is false, in the sense that, say, some post-main-sequence stars are blue; it is rather that (ii) would be seen to rest on a classification of stars, into main-sequence and post-main-sequence, which has no basis in the data. To put it slightly differently: it would emerge that (ii) implies a question ("What color are post-main-sequence stars?") which has—on the basis of the available data—no objective answer.

Note that the point here is not about essence versus accident, or about what is true by definition. It is unclear exactly how we should today apply the (quasi-)Aristotelian terms "essence" and "definition." But however we might conclude in that respect, we probably would *not* want to say that (relative) redness is an essential attribute of post-main-sequence stars, or that they are red by definition. If there is an essence of post-main-sequencehood, it consists rather in a star's having exhausted the hydrogen in its core and begun to shine due to some new energy source. It is not even that far-fetched, let alone literally inconceivable, that we should someday find that some large class of such stars is not (relatively) red. And even as things stand it is not the case that *all* such stars are red: the big empty spaces in fig. 1 are mostly due, not to the fact that no star ever occupies such positions, but to the fact that no star remains there for more than a small fraction of its lifetime. Hence, far from being a necessary truth, (ii) is not even universally true, only approximately so. The point here is not logical or metaphysical, but epistemological: the evidence which, in our current state of practice, serves to confirm the (approximate) truth of (ii), is not separable from the evidence by which we establish that the classification main-sequence/post-main-sequence is a good one. Evidence which served to undermine the former would thus also have to be such as to strike a blow against the latter.

Note also that the point here is not about whether sentences face the tribunal of experience alone or as a corporate body. It is true enough, as I have already pointed out, that (ii) could not, or could not easily, stand or fall alone. And it is also true that under some circumstances in which a statement like (ii) is faced with data like those in fig. 2 we might be inclined (for whatever reason) to say that—in terms of this example—post-main-sequence stars are in fact blue, or that all stars are in fact main-sequence, or even that all post-main-sequence stars are surrounded by dark clouds of dust and hence cannot be detected at visible wavelengths. But although all these things are true and important, they are not the point of my example. The point is about how, under circumstances such that we *would* allow the data to confirm or (roughly speaking) disconfirm (ii), we would go about allowing them to do that. And on that topic it should be clear that, at least in this example, it is the overall pattern of observations, rather than any particular observation as such, which contributes to the confirmation or rejection of (ii). One might say that the data, as judges in the tribunal of experience, sit not individually but en banc.

Note also, finally, that the point here is not about the so-called "theory-ladenness of observation." Observation *is* always laden with theory,[4] and I will have something to say, below, about precisely why that must be the case. But, once again, although that is true, it is not the point of this example. In fact, what we can see here is almost the opposite: the way in which the observational data themselves can, in the proper circumstances, throw off a piece of theory with which they are laden. If we approach fig. 2 on the theory that there are two kinds of stars, main-sequence and post-main-sequence, which differ, among other ways, in color, then we will be frustrated: the data shown there do not allow themselves to be interpreted in that way. That is precisely why they serve to undermine the theory-laden statement (ii). And, by the same token, the data of fig. 1 serve to *confirm* (ii) precisely because they validate the classification main-sequence/post-main-sequence in terms of which it is stated. The data, in their overall pattern, do not, so to speak, determine whether (ii) is the right or the wrong answer, but whether it is an answer to the right or the wrong question. It remains to be established how this works and whether and how it can be generalized beyond this one case.

Let us, first of all, return to (i), the statement that all ravens are black. Although it is formally similar to (ii), it looks at first blush as if the process of confirming it is or would be different. Wouldn't we simply, as discussions of this example normally assume, go out and examine many ravens to see if they were black? By finding many black ones, while at the same time not finding any (or very many) non-black ones, we could gradually build up evidence for the truth of (i).

Once again, I would not deny that we do often go through a process like this. But I agree with Popper, first of all, that such a process is not particularly scientific or even rational. Later I will argue that this kind of "induction" is in itself nothing more than superstition. And I would point out, secondly, that the imagined process is only possible at all because the question of what ravens are—of how ravens are to be distinguished from other things—is regarded by us as unproblematic. This is exactly why (i), in our current state of practice, is merely a fact with which scientific theories must be consistent, rather than itself a statement of scientific theory. That we can simply "go out and examine many ravens" means that we *present* our scientific theories (e.g. of evolutionary or molecular biology) with a certain thing, and say, in effect, "This is a raven. Note that it is black. Explain." And our theories had better be able to explain that, or at least to explain why it is possible. But the discovery of a white raven would not, as such, be any problem for those theories. That would simply be a new challenge: explain why *this* raven is white. Of course, if the white ravens turned out to be identical (genetically, chemically, ecologically) to the black ones, then we might have a problem on our hands. But that would be because something which, according to our theories, ought to be black, is actually white (or vice versa); it would have nothing to do with the whiteness of ravens as such.

If, therefore, (i) were to be an actual statement of scientific theory, which could be confirmed or undermined by a process of scientific research, it would have to be the case that the theory itself has a stake in which things are ravens, as well as in the blackness of those things. This "having a stake" is in reality a somewhat vague matter which admits of degrees—that, in part, is the moral of Quine's point about sentences facing the tribunal of experience. But simply by talking about "theories," about statements included in or entailed by them, and so forth, we are already engaged in an extreme idealization. And, continuing in the spirit of that idealization, we can understand how it is that a theory comes to have a stake in the nature of some (kind of) thing by means of an idealized account of the way in which one might introduce a name for that (kind of) thing as a new, theoretical term.

The particular idealized account which I have in mind here is due to David Lewis (1970).[5] Briefly and roughly put, Lewis imagines that a theory says something (as it might be, in one long conjunctive sentence) using one or more new terms—new in the sense that we plan to treat them as if everything we know about their meaning is contained in what the theory says. Or almost everything, rather. For such a term-introducing theory, according to Lewis, is always accompanied by an implicit claim that there is a unique way of assigning denotations to those new terms such that the theory comes out true. If this existence-and-uniqueness claim turns out to be false, then the theory will be false, as well, and its new terms denotationless. But if, conversely, the theory is true, then the existence-and-uniqueness claim must be true, and this supplies the crucial missing piece which will make our terms fully interpreted, and thus, among other things, allow us to find out more about the entities which they denote. To put the same point somewhat differently: the explicit theory specifies the entities named by its terms by giving something like an (empirical) concept under which they are supposed to fall, whereas the accompanying implicit claim attaches that concept to an object (i.e., provides the *tertium quid* which makes empirical synthesis possible) by adding to it the transcendental predicates of existence and uniqueness.

If our theory is supposed to introduce "raven" as a theoretical term, therefore, it must say something about ravens, and we must be prepared to treat that as (almost) everything we know about the meaning of the term. Now, it could be that what the theory says about ravens includes or strictly entails (i)—i.e. that (i) is "part of" the theory in a strictly formal sense. For example, we might have:

(T) ... & a raven is a black bird with a short neck & ...

In that case (i) would turn out to be true by definition: if there is anything at all which is properly called a raven, it would have to be black. A more normal case, however, would be that (i) is "part of" the theory in the looser sense that, given suitable approximations and simplifications, it follows from that theory, together with reasonable background assumptions, that (i) is (mostly, relatively) true. Imagine, for example, that we have an evolutionary theory which (unlike our actual ones) allows us to say for certain what ecological niches there are, and to predict with certainty whether they will be filled. Then the theory might introduce the term "raven" as follows:

(T′) ... & the ecological niche Δ is filled by the species of ravens & ...

And it might follow from *that* (together with some reasonable background assumptions, under certain approximations, etc.), that ravens are (mostly, relatively) black and short-necked, and that they are birds, or at least flying animals (i.e., animals which mostly fly pretty often and relatively well).

Given the intrinsic unreliability of data, the difference between (T) and (T′) is perhaps less important than it seems. But, in any case, they both have in common what we need for our purposes: they both require the truth of (i) in a way that our current theories do not. Neither can be allowed to explain with equanimity why some ravens (a significant number of them) turn out to be white, after all. For they both, as I put it above, have a stake in which things are ravens and which are not. In the case of (T) this works directly: if (i) is not true then it can only be because there are no such things as ravens, in which case the implicit claim which goes along with (T) is false.[6] In the case of (T′), the situation is more complicated, since it is strictly speaking possible for (T′) to be true while (i) is not (even mostly and relatively). It is possible, for example, for there to be such things as ravens, according to the (T′)-definition, and yet for those things to be mostly white. This is similar to the case of post-main-sequence stars: it is not inconceivable that many or even all of them might be blue. But that could only happen if some of our background assumptions, approximations, etc., were incorrect (e.g. if many stars turned out to have an initial helium content much different from what we were led to expect, or if our approximate treatments of convection turned out to give seriously wrong answers), and the same holds in the case of (T′). Since, however, our background assumptions, etc., are by definition things we are inclined to take as basically right, whatever kind of evidence tends to rule out (i.e., undermine) (i) will be prima facie evidence that there are no such things as ravens, and vice versa.

But what kind of evidence would that be? Note that, although we have so far been blithely assuming that the implicit claims made by (T) and (T′) might be true, in fact that is far from obvious. Let's assume, for example, that (T) says nothing about ravens other than what I've written out. Then the claim is that there is a unique property ϕ (the property of being a raven) such that the things which possess that property are black, short-necked birds. But there is a serious problem, because, although (T) makes (i) true by definition, the definition in question is nominal, not real: (T) doesn't claim that ravens are black, short-necked birds by essential necessity. ϕ can therefore be any property such that, as a matter of contingent fact, all ϕ things are black, short-necked birds. And we can easily cook up infinitely many such properties. This is easy to see in Lewisian metaphysics, where a property is class of possibilia. We can start, for example, with the class of actual black, short-necked birds, and then stick on any class whatsoever of non-actual things to get a property which satisfies (T). Even, moreover, if (T) said something more complicated about ϕ (and, in particular, even if it said something modal) this problem would recur as long as the theory remained empirical, as opposed to metaphysical—i.e., roughly speaking, as long as its claims continued to be about the things which happen to be actual, rather than about possible things in general. If, therefore, there is to be any chance of (T)'s coming out uniquely realized, we must add the stipulation that the property of being a raven is, as Lewis would later put it, a (more or less) *natural* one (strictly speaking: a natural property which is instantiated at our world).

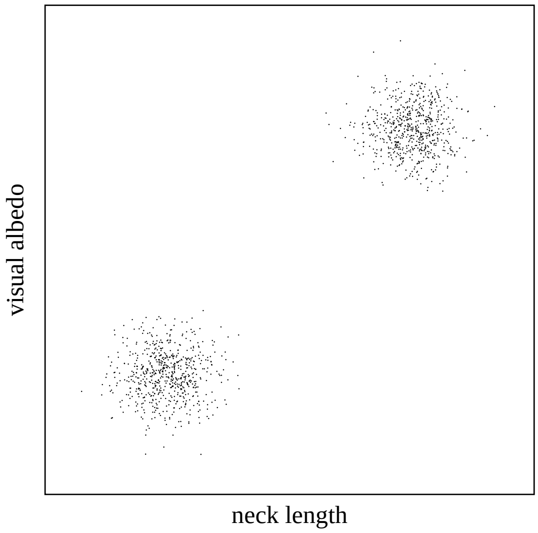

Figure 3. Ravens and swans (confirming).

I don't know if Lewis ever mentions this stipulation explicitly, even after he has already bought into the idea of "natural" properties (which at the time he wrote "How to Define Theoretical Terms" were still not a part of his metaphysical apparatus). But I think he would nevertheless have agreed that he needed such a stipulation in our kind of case. He does think, at any rate, that, "in putting forward as comprehensive theories that recognize only a limited range of natural properties, physics proposes inventories of the natural properties instantiated in our world" (1983a, 38).[7]

Now, I do not claim to know what, metaphysically speaking, makes a property natural, or even, as Lewis does, to know what some of the main competing explanations might be. Nor do I know how we could, starting with a list of all possible properties (including, for example, the property of being grue), proceed to pick out just the most natural ones, or to rank them all in order of naturalness. Nor do I claim to know how we ever come to treat any property as more natural than any other (i.e., I do not know how to fix "global descriptivism"). But I think I do know how we go about testing the proposal that a certain group of things have a single (more or less) natural property in common. Consider, for example, the data presented in fig. 3 (in which each data point represents a bird: ravens on the lower left, swans on the upper right). It is this kind of data, I assert, that we would take as evidence that the black, short-necked birds all have something—that is, something more or less natural—in common.[8]

The data of fig. 3 are theory-laden. They are irremediably theory-laden in that they presuppose, for example, that black is a more natural property

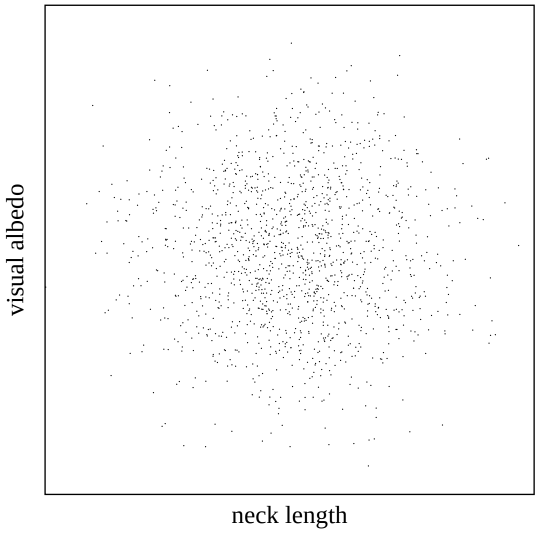

Figure 4. Ravens and swans (undermining).

than grue (or, as it might be, blite). But their ladenness with a theory of ravens, with (T) or (T′), for example, is not similarly irremediable. Rather, the legitimacy of such a theory, its objective right to load itself onto our data, is tested, and confirmed, here—confirmed, not by any particular data point, but by their whole distribution. There is, of course, no logical (deductive) relationship between that distribution and the truth of (T) or (T′). And who expects to find such a relationship between our data and our empirical theories? But it is nevertheless only the kind of data presented here which can give us objective warrant for deploying the concept of ravenhood, for classifying birds into ravens and non-ravens, just as it is only the kind of data shown in fig. 1 which warrants our classification of stars into main-sequence and post-main-sequence. The metaphysical concept of natural property thus functions as a regulative ideal for the empirical employment of the understanding: the ontological requirement on things as such, that they have the transcendental predicates of unity and uniqueness, is represented (schematized) by the empirical requirement that they produce a recognizable regularity in our data. And it is in this way, by legitimating the concept *raven* (by showing that "being a raven" has a denotation) that the data of fig. 3 serve to confirm (i).

Just as (ii), moreover, could not be falsified by finding a single post-main-sequence star which is not red, (i) could not, in the imagined context of (T) or (T′), be falsified simply by finding a white raven. Under (T), of course, "finding a white raven" is literally inconceivable. Under (T′), it is not, but, given our faith in our background assumptions, etc., we would normally take the whiteness of a bird as a sign that its species very likely does not fill niche

Δ—i.e., that it is not a raven at all. And under those assumptions, etc., the kind of data we see in fig. 3 *ought*, therefore, to be exactly the kind we need to distinguish between ravens and other things. In that fictional situation, a whole cluster of points on the upper left would be taken, not as a discovery that some ravens are white, but as showing that there is some other kind of bird that is white and has a short neck. Yet this ladenness of the data with (T′), and with it their confirmation of (i), could be shaken. The data shown in fig. 4, would not, it is true, logically compel us to give up on our concept of ravenhood—we could simply draw a circle in the lower left corner and call everything inside it a raven, if we so desired. But our objective warrant for using that concept would be gone.[9] Under (T), it would be gone absolutely; under (T′), which is the more realistic case, gone *unless* we decide to give up something else.

This is all an extreme idealization, as I said to begin with. Our resolve to treat our terms in a certain way is never absolute. Faced with the results of fig. 4 we might in reality end up saying that ravens are actually gray, that a raven is just a funny kind of swan, or what have you. But it likewise remains true that if any kind of observational data could get us to give up (i), they would be, under the circumstances we are now imagining, something like those shown in fig. 4. And under such circumstances it doesn't much matter whether we decide to say that (i) is false or meaningless or even true: the concept deployed in (i), the reason for saying it, the question it is meant to answer, are shown not to be usable. In other words, (i) has, as I put it above, been "undermined."

2. Science and Superstition

It emerges, therefore, that (i), regarded as a statement of scientific theory, would actually be tested in the same way as (ii). Perhaps that is not so surprising. I chose originally to investigate (ii), after all, precisely because of its apparent similarity to (i). Not every statement of scientific theory is of the form "All A's are ϕ," though, so we shouldn't expect every process of scientific testing to be precisely similar to the ones we have been discussing. We can nevertheless, I think, generalize the results with respect to (i) and (ii) to cover all empirical testing of scientific claims.

To see what is wrong with the usual image of "induction," and how it needs to be corrected in line with what we have learned from examples (i) and (ii), let me introduce a third example, this one taken from an early paper by Charles Sanders Peirce:

(iii) Bleeding tends to cure cholera.

This is an example, not of a simple observed regularity, but of an alleged causal law. As we will see, however, what constitutes (or would constitute) scientific rationality in the testing of this claim is the same feature we have already encountered above in the cases of (i) and (ii). And we will see, furthermore, that that feature would in this case be supplied by the use of statistics.

I begin with Peirce, who gives the following example of an argument supporting the conclusion (iii):

> A certain man had the Asiatic cholera. He was in a state of collapse, livid, quite cold, and without perceptible pulse. He was bled copiously. During the process he came out of collapse, and the next morning was well enough to be about. Therefore, bleeding tends to cure cholera. (Peirce, 1868, 43–4)

Peirce uses this as an example of a *good* inference, rather than of a bad one. A possibly good inference, that is—although not, of course, a possibly apodictic one. Its validity, if it had any, would depend on our lack of knowledge about other cases. As he explains:

> This is a fair probable inference, provided that the premises represent our whole knowledge of the matter. But if we knew, for example, that recoveries from cholera were apt to be sudden, and that the physician who had reported this case had known of a hundred other trials of the remedy without communicating the result, then the inference would lose all its validity. (Ibid., 44)

Now, I do not at all intend to pick out Peirce's analysis as particularly mistaken. In fact, I think his analysis (which, incidentally, forms only a small part of an extremely interesting paper) is unusually elegant and sophisticated. The (Hegelian) idea that what makes inductive reasoning possible is our *lack* of knowledge is appealing. He hints, moreover, at the importance of experimental control, to which we will soon return. Still, Peirce's version of these things, according to which the reasoning stated would under *some* circumstances be scientific, cannot be correct. It is true enough that we often make inferences like the one he outlines, and it is even true, in all likelihood, that they often lead to correct conclusions, and therefore also to hypotheses which, when properly tested, end up surviving the test. But the rationality of a method is not to be established by truth of its conclusions. And the rationality of modern science—in this case, in particular, of modern medicine—rests precisely on avoiding such inferences wherever possible. Regarded as a clinical trial, the procedure Peirce describes is deficient in at least two key respects: (1) the sample size is too small; (2) there is no control.[10] Of these, it is (2) that is the more fundamental. If there were a control, we would need a large sample size in order to carry out appropriate statistical tests; since there is none, there are no such tests to be made, and hence no possibility whatsoever of scientifically confirming or disconfirming (iii).

To see this even more clearly, and to see why I use the strong term "superstition" here, consider the following statement, which is similar in form to (iii):

(iv) Having a black cat cross one's path tends to cause bad luck.

Suppose my evidence for this is that a black cat once crossed my path, and that I subsequently had bad luck. Or suppose even more: suppose that a black cat crosses my path every day, and that every day, without exception, I have bad luck. On Peirce's analysis, or indeed on any of the usual analyses of "induction," I would then have very good reason to believe (iv)—even Popper would agree that (iv) would then have passed a severe test and hence be very well corroborated.[11] But although, once again, we often do reason in that way, the conclusions which we reach when we do so are not scientific theories, but simply superstitious beliefs. Scientists are not, as a matter of fact, now engaged in testing the truth of (iv) (or, for that matter, of (iii)). But if they were, then the scientific testing of those statements would begin only with the examination of cases in which the alleged remedy was *not* applied and of days on which a black cat did *not* cross my path—that is, with the establishment of a control.

How is all of this related to the points I made above about (i) and (ii)? If there were a control, then, as I mentioned, its purpose would be to allow the application of a statistical test. Such a test would serve to rule out the so-called null hypothesis: that the difference, if any, between the control and experimental groups is due to chance alone. But this so-called "hypothesis" has of course no positive content; it simply means that the classification of cases into those two groups is not relevant to the outcome. Not that this is the only way that empirical data could undermine (iii) or (iv). It could be found that one is significantly *less* likely to have bad luck after a black cat crosses one's path, or (somewhat more plausibly) that copious bleeding lessens the chances that a cholera patient will recover. But even the simple, superstitious procedure for testing (iii) and (iv) leaves itself open, in principle, to that kind of disconfirmation. What is new in the scientific approach is that it allows (iii) or (iv) to fail, not because the difference between the control and experimental groups goes the wrong way, but because there is no (significant) difference to be observed. Not, in other words, because (iii) or (iv) is the wrong answer to the question we have asked about the data, but because that question itself is wrong.

In the cases of (iii) and (iv), the key element which allows scientific testing is the use of statistics. And, though it is perhaps not quite so evident, that is the key element in the cases of (i) and (ii), as well. We tacitly assumed, in discussing figs. 1 and 3, that the blank spaces in them were not due to selection effects—assumed, that is, that we had looked sufficiently hard that, if there were stars or birds of the type in question, we would have seen them. This corresponds to the establishment of a control. Given that assumption, the data shown are sufficiently clear that, simply by inspecting them in their graphical presentation, one can already see that the differences which interest us are significant. In certain areas of science, given the right experimental design, one can expect to be able to do that relatively often. But this method (sometimes known as “chi by eye”) really just amounts to doing by rough estimation what a statistical test would do more precisely. If things in figs. 1 and 3 were less clear—if there were not so many data points, or if the different regions of the plots were not so clearly distinct—then only the more precise statistical technique would do. Before going on, let me therefore say something about the implications of this central role of statistics.

First, the use of statistics implies the recognition that one’s “data” may not be entirely reliable. As the famous astrophysicist Fritz Zwicky is supposed to have said: “No theory should fit all the facts, because some of the facts are wrong.” To this one could add that the fitting theory is probably wrong as well, in the sense that one is rarely in a position to derive exact predictions—one must make do with approximations and simplifying assumptions. In philosophical discussions of the scientific method, these issues of the unreliability of data and predictions are, where they are not entirely ignored, frequently treated as if they were secondary complications. But in fact they go to the heart of the matter. Every measurement or calculation yields results which in principle are indefinitely precise, which is to say: every measurement or calculation can be used, in principle, to answer whatever questions we may put. In Peirce’s terms one could say that every question can be answered by a valid probable inference on the basis of induction on a single case. But the beginning of scientific wisdom is the fear that some of these “results” are *insignificant*, and the methods of modern science, including first and foremost the methods of experimental design and statistical analysis, are aimed at addressing that fear.[12]

I would not, however, want to overemphasize the importance of formal statistical methods, or claim that all scientifically rational testing must lend itself to their explicit use. If we were to discuss statements at a higher level than (i)–(iv), such as for example

> (v) A body in motion remains in motion unless acted upon by an outside force.

or

> (vi) $\mathbf{F}_{\mathrm{grav}} = -(Gm_1m_2/r_{12}^3)\mathbf{r}_{12}$

then we would not be able to point to any specific body of data which furnishes their statistical support. And yet I must certainly discuss such statements as (v) and (vi), since it is this kind of statement that is adopted in a so-called scientific revolution, and since, as Kuhn points out, it is precisely there that only a sociological, rather than a logical or epistemological, account of the distinctiveness and dominant status of modern science might seem possible.

Let us begin by noting that it is the adoption of (v), rather than of (vi), that constitutes the Newtonian revolution, properly speaking. Given (v), the problems of celestial mechanics, of the trajectory of falling bodies, and of the tides, already have something in common: in each case, bodies are accelerated without being (obviously) pushed. Once that is noted, the step to (vi) (and to the concomitant final acceptance of a heliocentric system) is incredibly difficult and brilliant, certainly, but does not as such involve anything like incommensurable competing theories. And within that framework the data relevant to (vi)—the data, that is, by means of which it is tested—will be similar to those recorded in figs. 1 and 3, even though they cannot be conveniently represented in a single figure. The evidence by which they confirm or undermine (vi) is statistical—not in the sense that we might actually apply some gigantic statistical test to all the relevant data at once (and after all such tests hadn't even been developed yet in Newton's time), but in the sense that the function of the data in providing evidence here is like the function we have seen them fulfill in statistical testing. That means, first of all, that we do not treat as entirely reliable either our data or our predictions based on (vi): we know that observers make errors, and that we can derive predictions from (vi) only based on approximations and on simplifying assumptions which are, at best, not known to be true and, at worst, actually known to be false. It also means, secondly, that no single observation can serve either to confirm or to disconfirm (vi). Note well: the point is not (what should be obvious) that no such single observation serves as a *conclusive* verification or refutation. The point is that none of them is by itself worth anything, any more than is a single dot on fig. 1, or a single case in which copious bleeding allegedly cures cholera. The statement (vi) was confirmed, as long as it continued to be confirmed, by an entire pattern of data—by the fact that (vi) explained that pattern better than "the null hypothesis." And it was undermined, when it finally was, by a disruption of that whole pattern.

That undermining, when it came, was itself part of a "revolution": the statements by which (vi) was ultimately replaced are incommensurable with it: $\mathbf{F}_{\text{grav}}$ no longer occurs there. Rather than analyze that abstruse development, however, it will be easier to return to the genuinely revolutionary (v). How could (v) be tested, or by what kind of data could it be confirmed? It has often (though, alas, not always) been noticed that there is no way at all to check it by "induction." For of course there is only one way we can tell that an outside force is acting, namely to see that some body does not (simply) remain in motion (i.e., is accelerated). Hence (v) may seem tautological, or rather may seem to serve, together with the law that $\mathbf{F} = m\mathbf{a}$, merely to define the term "force." Or it may, on the other hand, seem to have a trans-empirical, metaphysical content. To see what is really going on, however, we have to remember that (v) is a revolutionary statement, and therefore examine its relationship to the doctrine which it replaced.

What is called "Aristotelian physics" is not the kind of theory which could be represented by statements like (ii) or (vi). Within the Aristotelian tradition, it was a matter of interpreting Aristotle. For us, who can no longer (or can longer entirely) belong to that tradition, it thus becomes a matter of knowing how Aristotle was actually interpreted—interpreted, that is, by those who regarded him as an authority. There is no single consistent doctrine which answers to that description. But since the assertion of (v) was (a part of) the birth of modern science, it in effect served to set up its predecessor as a modern scientific theory for the purposes of criticism. And we can think of that imaginary theory as asserting something like this:

> (vii) An inanimate sublunar body moves to its natural place and remains there, to the extent that it is not violently prevented from doing so.[13]

This Aristotelian law of motion is not problematic in the way that (v) is; it is as empirical as (ii) or (vi). That is to say: under certain simplifying assumptions, and with the aid, perhaps, of some approximations, it predicts a certain correlation—in this case, between the velocity of a body and certain other of its circumstances (its material composition, its location relative to the center, the resistivity of the medium which surrounds it). To confirm (vii) would be to show that the data are correlated in this way to a significant degree (i.e., better than can be explained "by chance alone"). And what (v) says is that (vii) fails that test—fails not because it implies the *wrong* correlation, but because there is no such (significant) correlation at all. That is why (v), and the new statements such as (vi) that follow in its wake, are revolutionary, incommensurable with (vii): they oppose it not as giving the wrong answer (making the wrong judgments) but as answering the wrong question (applying the wrong concept). And yet for all that the method followed in choosing (v) over (vii) is no different than—is merely the negative side of—the "normal" scientific method of testing (ii).

Without going into the details, I would argue that the episodes which Kuhn calls scientific revolutions generally have this character, and that that is why *incommensurability* is so characteristic of them. They result from the discovery that what the old theory claimed to explain on its principles can in reality be explained just as well by "the null hypothesis." Their fundamental revolutionary insights, as opposed to the new detailed theories that grow up in their shadows, are thus negative rather than positive. If stated as positive explanatory principles—the kind of things that could be confirmed or disconfirmed in "induction"—they appear to be empirically empty.[14] And in fact such insights do not in themselves explain anything. But they do bring about the need for radical new explanations, by overthrowing the old idea of what there was to be explained.

3. Conclusion: Science, Skepticism, and Philosophy

That is as much as I have to say here about the (relatively) limited aims of this paper. It seems worthwhile, however, to give a few indications, in closing, as to how these conclusions are related to the larger project which I mentioned to begin with.

First, though I promise something about skepticism in the title, that connection is going to have to remain mostly implicit. Briefly, however: Cavell understands the philosophical skeptic, and I think correctly, as in effect asking "How do you know there is an X here?" in a case where there is no problem of correctly identifying or recognizing or describing X (for example, looking at the table, "How do you know there's a table here?"), and in that way making that question into a question about the existence of the world in general. If science proceeds by asking, not whether we've applied our concept correctly in this case—not whether we've performed a correct act of recognition or identification or description—but whether that concept is legitimate at all, whether there is anything at all for it to apply to—then it's perhaps not so surprising that scientific revolutions, like skeptical questions, are irresistible and irreversible. Of course, modern science isn't simply to be identified with skepticism about the external world, so there's plenty more to said about this topic. Presumably the use of mathematics in modern science, a subject upon which I've hardly touched, will be an important part of the story (because mathematics, after all, is about *succession*, i.e. about how to generalize a concept *step by step*).

But what, secondly, of the practical problem or contradiction at which I hinted towards the end of my introduction? It must be emphasized that the existence of modern science as a standard and even, to some extent, as an actual technique with respect to our social world, is not *only* a problem. If it were, then the kind of tendency to which I alluded under the term "fundamentalism" would be an attractive one. But the truth is that to apply incorrect concepts (concepts without objective basis in the data) in one's practical life is to be technically incompetent and politically unjust. Hence resistance to modern science, far from being justified, is both foolish and morally culpable.

There *is* a problem, however, as can be seen from the details of my own investigation. For it is now time to return to a question which I have put off until the last possible moment. What licenses me to assume that there is something distinctive about modern science? Or, in other words: what gives me the right to deploy that concept, to assume that the phrase "modern science" denotes some unique natural kind of activity? If there were an unproblematic modern science of sociology or of history, then by what right could I avoid submitting to its conclusions? By what right could I challenge its empirical results on the basis of mere science fiction? The actual difficulties confronting the establishment of such disciplines are of course well known, and as I indicated, moreover, there may also be difficulties in principle. But although that lessens the risk of embarrassment on my part, it only serves to make the epistemological situation worse. Just as I cannot expect refutation from that quarter, I cannot expect much support, either—whereas the demand for such support, once raised by the existence of any modern science whatsoever, remains in force.

It seems that philosophy can have no right to deploy such concepts. And yet: one must always know enough to do one's duty (to philosophy and to the city). Philosophical self-understanding stands and has always stood under the threat of practical contradiction, of moral incoherence. It has always stood under that threat; I do not have a way, do not believe that there could be a way, to remove it. But I do hope that philosophy has something or everything to say about such matters. And I would claim one thing (the only thing that Socrates will claim): that the existence of such a paradox is no excuse for abstaining from the work and duty of philosophical investigation.

Notes

[1] I am currently working on an entirely new version of this paper, to be titled "On Scientific Method as a Method for Testing the Objectivity of Concepts." The introductory section of the new version is available in PDF at http://home.uchicago.edu/~abestone/papers/newsmintr.pdf.

[2] Popper does not deny (and neither will I) that human beings, and even scientists, may at times follow such a procedure; but he holds that to do so would be a mistake. (See 1935, 22 [=1959, 52–3].)

[3] I say this despite Ian Hacking's sustained but, I think, unconvincing attack on their collective work (for details, see 1983, ch. 10).

[4] For example, the observation, "My typewriter is on the table," is laden with the (relatively false) theory of absolute simultaneity and with the (even falser) theory that there are typewriters (i.e. that "typewriter" names some natural kind of things), not to mention with the theory of private property (if there is such a theory).

[5] Lewis's treatment is closely based, as he makes clear, on ideas which he has taken from Ramsey and Carnap. It is also closely related (as he makes somewhat less clear) to some ideas of Hempel's (see especially 1952, 81 n. 26). It can be seen, furthermore, as following up on suggestions of Putnam's (1962, 216, 219). But it is Lewis's own treatment, rather than theirs, which is most useful for my purposes.

[6] It doesn't much matter whether we count (i) as true, false, or meaningless in the case that "raven" fails to denote, and in fact Dana Scott's logic of denotationless terms, which Lewis proposes to use, is compatible with at least the first two of these alternative (see Lewis [1970, 81]). The main point is that if there *are* things properly called "ravens," then (i) must be true of them.

[7] I note parenthetically, to avoid confusion, that the problem and the suggested solution here are not the same as "Putnam's paradox" and Lewis's suggested solution to it, respectively. "Putnam's paradox," as Lewis understands it, is a general, model-theoretic refutation of "global descriptivism," and Lewis's solution is to reject global descriptivism in favor of constraints on reference which are not "just more theory"—i.e., not stipulations which we (implicitly or explicitly) make (see Lewis [1984, 58–64]). The present problem, however,

concerns "local descriptivism." It seems to me that there can be no general model-theoretic refutation of local descriptivism, for there is no metalogical reason that more than one person, say, must have committed the crimes attributed to Jack the Ripper. (See, however, van Fraasen [1997]. In private correspondence on this issue van Fraassen and I have unfortunately been unable to come to complete agreement.) In any case, the argument as presented is not a refutation of local descriptivism at all, since the suggested solution is not to reject local descriptivism, but to add, as "just more theory," an (at least implicit) stipulation: physicists (and other scientists) must be understood as (at least implicitly) *claiming* that they discover natural properties. (Cf. Lewis [1984, 67]: "physics professes to discover the elite properties.")

[8] I take it that this is in the same spirit as Lewis's observation that "Of course, the discovery of natural properties is inseparable from the discovery of laws. For an excellent reason to think that some hitherto unsuspected natural properties are instantiated . . . is that without them, no satisfactory system of laws can be found" (1983a, 38). Certainly the situation in fig. 3 is the kind which a satisfactory system of laws is traditionally supposed to explain. Lewis, however, (a) seems to have in mind some other reasons—I can't imagine what—for thinking that there are natural properties, and (b) neglects to mention what kind of discovery can lead us to reject as natural a property which we previously accepted. This last omission leaves him open to the objection that, since the concepts with which science begins are historically and biologically conditioned, the methods of science are inherently unsuited to a search for natural properties: see Elgin (1995, 293–5).

I take it, too, that this state of affairs is the same kind Hempel has in mind when he remarks that the "rational core" of the distinction between natural and artificial classifications is that "in so-called natural classifications the determining characteristics are associated, universally or in a high percentage of all cases, with other characteristics, of which they are logically independent" (1952, 53). Missing here, however, is the idea that we need such evidence of naturalness to validate even our introduction of classificatory terms: without it we lack, not just a useful or fruitful or simple theory, but a ground for taking our theoretical terms as meaningful. Term-introducing statements, in other words, must have"experiential import," not only because, as Hempel explains (e.g., ibid., 18–20; 27–8), such statements may lead to contradictions if certain empirical laws do not hold, but also and more importantly because their implicit claim of unique denotation can only be justified where we have evidence for such laws. The advantage of Lewis's treatment over Hempel's, Ramsey's, and Carnap's is that, in identifying and separating out the implicit claim in question, it makes that dependence clear.

[9] For a related point see Kuhn (1974, 811–12).

[10] There is also the issue of placebo effects, though these are perhaps not worth worrying about when the patient had no perceptible pulse before the procedure began. In any case, I will ignore that complication.

[11] Strictly speaking: that it would be relatively very well corroborated—relative, that is, to its maximum possible degree of corroborated (which in the case of such a particular statement is low).

[12] To *recognize* this unreliability of theory and observations is not the same as simply knowing that sensible measurement is imperfect: it means knowing what to do with that fact, how to use it to modify one's conclusions. This has important bearing on the status of so-called ancient science, but I cannot discuss that matter further here.

[13] This principle could be extended to animate and/or celestial bodies, but only at the cost of settling various issues which are highly controversial within Aristotelianism.

[14] Cf. Popper's complaint against Darwinian evolutionary theory that its principle of natural selection is tautologous (e.g. [1979, 241–2]). See also (Schlick, 1936). (Schlick explains the

content of the law of inertia and the thermodynamic principle of energy conservation in a way which fits in very well with my discussion, although his general diagnosis is different.)

This negativity of fundamental theoretical insights is related to the fact that many of them (for example, (v)) can be stated as symmetry principles (i.e., as principles of invariance under some family of transformations). Such a symmetry principle is a way of saying exactly what it is that a theory proposes not to take as potentially significant.

References

Elgin, C. Z.: 1995, 'Unnatural Science'. *Journal of Philosophy* **92**, 289–302.

Hacking, I.: 1983, *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. New York: Cambridge.

Hempel, C. G.: 1952, *Fundamentals of Concept Formation in Empirical Science*, Vol. 2, no.7 of International Encylopedia of Unified Science. Chicago: University of Chicago Press.

Kuhn, T. S.: 1974, 'Logic of Discovery or Psychology of Research?'. In: P. A. Schilpp (ed.): *The Philosophy of Karl Popper*. La Salle, IL: Open Court, pp. 798–819.

Lewis, D. K.: 1970, 'How to Define Theoretical Terms'. *Journal of Philosophy* **67**. Reprinted in Lewis 1983b, 78–95.

Lewis, D. K.: 1983a, 'New Work for a Theory of Universals'. *Australasian Journal of Philosophy* **61**. Reprinted in Lewis 1999, 8–55.

Lewis, D. K.: 1983b, *Philosophical Papers*, Vol. 1. New York: Oxford University Press.

Lewis, D. K.: 1984, 'Putnam's Paradox'. *Australasian Journal of Philosophy* **84**. Reprinted in Lewis 1999, 56–77.

Lewis, D. K.: 1999, *Papers in Metaphysics and Epistemology*. New York: Cambridge University Press.

Peirce, C. S.: 1868, 'Some Consequences of Four Incapacities'. *Journal of Speculative Philosophy* **2**, 39–72.

Popper, K. R.: 1935, *Logik der Forschung: zur Erkenntnistheorie der modernen Naturwissenschaften*. Vienna: Springer.

Popper, K. R.: 1959, *The Logic of Scientific Discovery*. London: Hutchinson; reprinted New York: Routledge, 1992.

Popper, K. R.: 1979, 'Of Clouds and Clocks'. In: *Objective Knowledge*. New York: Oxford University Press, pp. 206–55.

Putnam, H.: 1962, 'What Theories Are Not'. In: E. Nagel et al. (eds.): *Logic, Methodology and Philosophy of Science*. Stanford, CA: Stanford University Press. Reprinted in: *Mathematics, Matter and Method*, Vol. 1 of Philosophical Papers, second edition (New York: Cambridge University Press, 1979), 215–27.

Schlick, M.: 1936, 'Sind die Naturgesetze Konventionen?'. In: *Induction et probabilité*. Paris: Sorbonne, pp. 8–17.

van Fraasen, B.: 1997, 'Putnam's Paradox: Metaphysical Realism Revamped and Evaded'. *Philosophical Perspectives* **11**, 17–42.