Popper's Philosophy of Science: Looking Ahead

Peter Godfrey-Smith

Philosophy Department
Harvard University

2007

1. Popper's Peculiar Standing

2. Eliminative Inference

3. Skeptical Realism

4. Tests, Risks, and Pseudo-Contact

5. The Diachronic Perspective on Evidence

6. Conclusion

1. Popper's Peculiar Standing

Is Popper's philosophy alive or dead? If we make a judgment based on recent discussion in academic philosophy of science, he definitely seems to be fading. Popper is still seen as an important historical figure, a key part of the grand drama of 20th century thinking about science. He is associated with an outlook, a mindset, and a general picture of scientific work. His name has bequeathed us an adjective, "Popperian," that is well established. But the adjective is used for very general ideas that, according to most current philosophers, Popper did not develop convincingly. His detailed account is often seen as attractive on first impression, but full of holes that become bigger rather than smaller as discussion continues. The picture and the name remain, which is more than

most philosophers can hope for. But the name attaches more and more to a set of instincts and intuitions, less and less to views that are seeing ongoing philosophical development.

Inside science itself, Popper's standing is quite different. He continues to be just about the only philosopher who can seize the imagination and command the loyalty of successful professional scientists. And he is popular within science not only for purposes of general commentary and public relations. Popper's philosophy is a *resource* drawn on by scientists in their internal debates about scientific matters. This has been especially marked in some parts of biology (Hull 1999).

From the point of view of philosophers, this affection on the part of scientists may have an unflattering explanation. Popper offers a rather heroic view of the scientific character, featuring an appealing combination of creativity and hard-headedness. It is no surprise that scientists like to be described this way. It is no surprise that they prefer this picture to the one often (though inaccurately) associated with Hempel and Carnap – the scientist as a sort of logic-driven pattern-recognition machine. The same applies to the picture associated with Kuhn, who presents the normal scientist as a narrow-minded, indoctrinated member of a peculiar collective enterprise, an enterprise devoted to solving tiny puzzles most of the time but prone to occasional fits of crisis and chaos. Popper's picture is much more appealing.

Scientific admirers of Popper might respond with a similarly unflattering explanation of why Popper has faded from professional philosophy of science. Professional philosophy of science is, for many scientists, an irrelevant exercise that is prone to long excursions into fantasy-land, obsessed with pointless semantic disputes, and very far from the living enterprise of science. From this point of view it is no surprise that Popper, a philosopher who had the blood of science running through his veins, might come to be unwelcome in the dour conference rooms of the Philosophy of Science Association.

So unkind explanations for Popper's mixed reputation can be thrown in both directions. In this paper, my aim is to isolate and explore some parts of Popper's philosophy of science that seem to me to have continuing relevance. One of them, in fact, I see as a "sleeper" idea – something that Popper may have been well ahead of his time

on, and an idea that will eventually be developed in much more detail. The four ideas I will discuss are all entirely philosophical, rather than ideas at the technical or historical fringe of Popper's philosophy. I will discuss them in increasing order of complexity.

I do not suggest that the four ideas I will discuss are the only ideas of Popper's that have continuing philosophical importance. On the more technical side, Popper's views on probability continue to attract interest. Both the general idea of probabilities as propensities, and some of the mathematical details, are the subject of ongoing discussion (Hájek 2003a, 2003b). I will not discuss that work here. My focus is on some of the core philosophical – especially epistemological – themes in Popper. In addition, I will not discuss ideas for which Popper deserves some historical credit, but which have now become fully absorbed into the tradition. An example in this category would be Popper's anti-foundationalism, and his vivid metaphor of how we can successfully build scientific knowledge on a swamp, by driving our piles in just as far as they need to go at the present time (1959 p. 111). Popper was far from the only philosopher to develop such views in the 19th and early 20th century (Peirce, Dewey, and Neurath are examples). And the details of Popper's treatment had problems; his was a rather shaky anti-foundationalism. But the contribution was important.

A final preliminary point should be noted. I assume in this paper a particular interpretation of some of Popper's ideas about testing and evidence. Roughly speaking, this is a simple interpretation that emphasizes the skeptical side of Popper's work, especially his skepticism about induction and other mainstream ideas about the support of hypotheses by observations. One problem with Popper is his tendency to try to have things both ways, with respect to these issues. Confirmation as a relation of support between observation and theory is rejected with great fanfare, but corroboration is then ushered in, and it can be hard to tell the difference.[1] This uncertainty has been rhetorically useful for Popper. Many of his scientific backers do not realize how skeptical some of Popper's ideas really were. In this paper I assume a simple, strong interpretation of Popper's critique of mainstream ideas about evidence and testing. All the results of observations can do is refute hypotheses, never support them. Confirmation – a relation of support between evidence and theory that can rationally affect our level of confidence

3

in the theory – is an illusion. The most we can say about a theory that has survived our strenuous attempts to refute it is just that – it has survived our attempts to refute it so far. So the Popper of this paper is the more skeptical Popper, a Popper who is epistemologically something of a radical, despite his rather proper demeanor. Perhaps this is a simplistic interpretation, but if so, my aim is to work out what we can learn from this figure, the skeptical Popper.

2. Eliminative Inference

The first idea I will discuss is simple and can be handled quickly. This is the very basic Popperian idea that the rational mode of theory assessment proceeds by ruling *out* alternatives. For Popper, this ruling out of options is all we can do. There is no direct epistemic support that a theoretical hypothesis can gain via observation, because confirmation is a myth. And a theory cannot gain indirect support through the ruling out of other options, because the number of options is (in all scientifically significant cases) infinite.

Like most philosophers, I reject many of Popper's claims about the impossibility of justifying theories. But this is compatible with a recognition of the great importance of the Popperian form of "eliminative inference" in science. And the crucial point here is that the role of this kind of inference was a neglected topic (among those outside Popper's school) in much of the 20th century, at least until very recently (Earman 1992, Kitcher 1993). The great obsession of mainstream philosophical theories of evidence in the 20th century, exemplified by Carnap and Hempel, was the direct positive support of generalizations by their instances. The formal problems that arose in the attempt to make sense of such support are notorious; some will be discussed below. What is important here, however, is the fact that during the development of the largely unsuccessful theories of "instance confirmation," little attention was paid to what seems in retrospect to be an obvious and central feature of the epistemology of science. This is the practice of seeking to support one theoretical hypothesis by ruling out others.

Perhaps it was often thought that only the ideal case would be an epistemologically significant one; the ideal case is where we are able to *decisively* rule out *all* options except one. Non-ideal cases depart from this one in two ways. First, there may be a less decisive ruling-out of alternatives; maybe we can only hope to show that all alternatives except one are unlikely. Second, there are cases where we might be able to rule out many or most, but not all, of the alternatives to a hypothesis.

Many philosophers have been discouraged by the thought that there will always be an infinite number of alternatives to any theoretical hypothesis. In scientific practice, the problem is made tractable by use of a notion of a *relevant* alternative (Goldman 1986, Kitcher 1993, Forber 2006). Only some options are seen as worth taking the time to exclude. This move alone does not solve the epistemological problem. What guides these judgments about relevance, and what rational basis could they have? It can be argued that scientists constantly tend towards over-confidence on this point (Stanford 2006). Scientists often think that they have ruled out (or rendered very unlikely) all the feasible alternatives to some theory. In hindsight we can see that in many cases they did not do so, as we now *believe* a theory they did not even *consider*. A focus on eliminative inference has the potential to illuminate both the successes and the failures found in scientific reasoning.

So an emphasis on eliminative inference and the introduction of a notion of relevant alternative does not solve the core epistemological problems here. It does, however, pose these epistemological problems in a better form than the form often imposed on them in mainstream 20th century discussion. It seems clear that much of the overt, day to day, practice of theory assessment in science proceeds by the explicit presentation of alternatives and the attempt to rule out as many as possible, by either deductive or probabilistic means. Some scientists have even sought to distinguish between fields that apply this procedure as an explicit *strategy*, from those that tend not to (Platt 1964). So whereas the traditional empiricist idea of the confirmation of generalizations by observing their instances is, at best, an ultra-idealized philosophical model of the epistemology of science, eliminative inference is a plain and central feature of scientific practice. The epistemological problems around this form of inference are far

from solved, but a focus on this phenomenon will surely be a large part of any good future account of evidence and theory choice in science.

3. Skeptical Realism

The second theme I will discuss relates to the debates about "scientific realism." Scientific realism is hard to characterize exactly. Part of the view seems to fall within metaphysics. Another part apparently has to do with the kind of "contact" with the world that scientific investigation makes possible. I think that in years to come it will appear that much 20th century discussion was focused on a rather awkward conglomerate doctrine, in much discussion under the "scientific realism" heading. In particular, there has often been the forging of a link between scientific realism and a kind of generalized epistemological optimism. The realist side of the debate is often associated with an overall confidence in current scientific theories as descriptions of what the world is like, or perhaps confidence in the trajectory or lineage within which current theories are embedded.

There is nothing wrong with asking about the levels of confidence in scientific theories that might be rational. The problem, as I see it, is the entangling of this family of questions with more basic questions about scientific realism. Roughly speaking, there are two questions (or kinds of questions) that should be considered in sequence. First, we can ask whether it makes sense to say that there is a real world existing independently of thought and theory that our scientific theorizing is *directed on*. Second, we can ask how confident we should be that our particular scientific theories are *succeeding* in representing how the world is.

These formulations of the two sets of questions are rough, and the treatment of them as questions to be addressed "in sequence" is obviously an idealization. But the separation is important. The first question is one about whether it is even *coherent* for us to take our scientific theories as directed on a world that exists independently of thought and theory. That question should be distinguished, as much as we can, from questions about how confident can we be that we are *succeeding* in the goal of representing what

this world is like. If the answer to the first question is "no," then the second question must be dropped or greatly transformed. But if the answer to the first question is yes, there are many options regarding the second. Some neglected options here include a skeptical position, and also a sort of "particularism" that I will discuss below.

First, though, a connection to Popper: it is part of the overall structure of Popper's philosophy that there is a good separation between these two kinds of questions. For Popper, there is no philosophical impediment to our regarding our theories as directed upon a mind-independent world. It is *possible* that we could devise a theory that is wholly accurate within its domain, where this includes the accurate description of unobservable things. Despite this being possible, the nature of theorising, evidence, and testing precludes us from ever having any confidence that we are *succeeding* in this task. Popper's view is, then, a moderate form of "skeptical realism" about science, and it is set up in a way that makes the possibility of skeptical realisms very clear.

A surprising amount of discussion of scientific realism in the 20th century paid little attention to skeptical realism as an option. Instead, "the question of scientific realism" was very often set up in a way that combined answers to both the questions above, and associated the label "scientific realism" with an optimistic, non-skeptical attitude towards current scientific theories, towards the scientific tradition as a whole, or both. Examples of this tendency include McMullen (1984), Boyd (1983), and Devitt (1997).

This feature of the debate has been due to influential arguments from both sides. Those in the "realist" camp have often been attracted to the idea that general arguments can be given from the predictive success of scientific theories to their likely truth. These success-based arguments can be given about particular scientific theories, or about the enterprise of science as a whole. On the other side, Larry Laudan (1981) argued that the overall pattern of change in the history of science gives us reason to expect that our present theories are *not* true.[2] So it has been common for philosophers to look for arguments, on one side or another, that would tell us what our overall level of confidence in science ought to be. And optimistic answers to this question have often been attached to the label "scientific realism."

As noted above, there are two possibilities that become marginalized by this way of setting things up. One is skeptical realism, of which Popper is an example. The other possibility is that the answer to the epistemological question that people associate with realism is complicated and field-specific – not summarizable or sloganizable in the manner of the standard debate. According to this "particularist" option, different scientific fields apply different representational strategies, attended by different kinds of risk and difficulty. Some fields (including parts of physics) are dominated by highly abstract mathematical formalisms, where it is unclear what sorts of entities and facts are being posited at all. In other fields (evolutionary theory, ecology, economics), we know what sorts of entities to believe in, but might wonder about how to treat the highly idealized models that have become the currency of much theoretical discussion. And yet other fields seem able to engage in straightforward mechanistic description of how systems are composed (neuroscience, molecular biology), of a kind that make many traditional philosophical anxieties seem quite misplaced. Different kinds of epistemological optimism will be appropriate in these different areas, and it seems unwise to attempt a blanket summary asserting that "most posits of well-established theories are real," "mature and predictively successful theories are approximately true," or anything like that. (See Psillos 1999 and Devitt 2005 for examples of expressions of optimism of this kind.)

So it is not Popperian skeptical realism that I see as likely to inherit the field in the scientific realism debate, but a "particularist" position that skeptical realism helps make visible.

4. Tests, Risks, and Pseudo-Contact

My third idea is a familiar and central one in Popper, but I will discuss it with the aid of a specific set of contrasts and connections. Popper claimed that a good scientific theory should take risks, should "stick its neck out." This is central to the whole idea of falsificationism, and Popper's attempt to give a "demarcation" of science from non-science. For Popper, a hypothesis that cannot be falsified by any possible observation

might be perfectly meaningful and even quite important, but it is not science. A genuinely scientific hypothesis must take risks, in the form of exposure to potential falsification via observational evidence.

The details of Popper's account of this risk-taking have well-known problems. Here I focus on a simple form of this idea, on the underlying intuition rather than the details. The idea is not just intrinsically useful, but it can be used to cast light on key problems with mainstream empiricist epistemology.

The mainstream empiricist tradition claims that experience is the only genuine source of knowledge. Some empiricists add, perhaps metaphorically, that empiricism is the idea that we acquire knowledge of the world by being brought into *contact* with it, via experience. Maybe "contact" is a questionable term to use, but let us accept it for now. Experience can certainly be described in a low-key way as bringing us into contact with what is going on in the world.

Given this, it seems that the details of an empiricist epistemology will be largely concerned with what this "contact" is and how it works. From this point of view, an important theme in Popper is the idea that there is something we might call *pseudo-contact* between theory and observation. There can be a genuine and intimate relationship between a theoretical idea and a piece of observational evidence, that has no epistemic value.

In Popper's philosophy this idea appears in various forms. It is central to Popper's "demarcation" criterion, as noted above. It also appears in a stronger form in the idea that successful prediction of surprising facts does not confirm a theory, if "confirmation" is seen as something that motivates an increased confidence that the theory is true. But one can accept the importance of the Popperian idea of pseudo-contact within a more moderate view. The key idea is that there are lots of ways for an observation to "conform" with a theory without that relationship having a genuine evidential *bearing* on the theory.

From some standpoints, this is a very obvious fact. If we develop an epistemology as Popper did, by thinking specifically about science and testing, then it will probably seem obvious. We can appreciate a simple version of the point just by noting that there

are lots of ways for data to be cooked up (deliberately or accidentally) in such a way that they *fit* a theory without *telling* us anything significant. More broadly, as Popper said, if you only look for confirmation then you will probably find it everywhere. So from some points of view, the importance of pseudo-contact seems obvious and straightforward. But apparently this message is not clear, obvious, and easy to take on board when one approaches the topic within other frameworks.

The fact that this message is not always easy to take on board is seen in an important feature of the literature on confirmation. Here I have in mind the view, especially associated with Hempel and Carnap but taken seriously by many others, that all positive instances of a generalization provide some support for the generalization. Formally, if we are considering a generalization "All F's are G," then all instances of Fs that are also G provide some degree of support for the generalization. (Hempel called this "Nicod's criterion.") Many people think there are problems with the idea that all positive instances do confirm, but a very large number of philosophers seem to be strongly *attracted* to this view; I have often seen in action the desire to salvage as much of the idea as is humanly possible.

From the Popper-informed standpoint, however, there is no reason whatsoever to believe it. Merely being a positive instance is not nearly enough for there to be confirmation or support. The evidential relation between a hypothesis and an observed instance will depend on much more than that bare logical fact. Roughly speaking, support of this kind can be expected to depend on whether the observation was the product of something like a *genuine test*, a procedure that had the possibility to tell either for or against the hypothesis.

The place where this issue becomes most vivid is the huge literature around the "ravens problem" (Hempel 1965). Famously, the ravens problem arises from three innocent looking assumptions. One is the idea that any observation of an F which is also G supports the generalization "all Fs are G." The second is the idea that any evidence that confirms a hypothesis H also confirms any hypothesis H* that is logically equivalent to H. The third is the idea that "All F's are G" is logically equivalent to "All non-G things are not-F's," which looks odd initially but is accepted in the most basic and

uncontroversial kinds of modern logic.[3] So the hypothesis "all ravens are black" is logically equivalent to "all non-black things are not ravens." But then we note that this last hypothesis is confirmed by the observation of a white shoe, as the shoe is a positive instance. Given the logical equivalence of the two hypotheses, and the fact that anything that confirms one confirms the other, the observation of a white shoe also confirms the hypothesis that all ravens are black. This looks absurd.

The relevance of Popperian ideas to the debate was noted rapidly by John Watkins (1964; see also Musgrave 1974). The problem only arises if one thinks that all observed positive instances confirm a generalization, regardless of whether those observations were made as part of a genuine test. If we insist that confirmation depends on more than the bare logical relation between instance and generalization, then the problem is easily handled.

Suppose that we know antecedently that an object is black, and we then inspect it to see whether it is a raven. We find that it is. But because there was no possible outcome of this observational process that would tell *against* the hypothesis, we should deny that the hypothesis gains any support from this observation of a black raven. And now suppose that we have an object antecedently known to be non-white in color, and we inspect to see whether it is a raven or not. We find it to be a shoe. This observation *was* part of a genuine test of the black-ravens hypothesis. It could have come out in a way compatible *or* incompatible with the hypothesis, and in fact it came out as the hypothesis claimed it would. Here, there is support for the hypothesis, because the observation was made as part of a genuine test.

The basic point about the role of "order of observation" in the ravens problem was also noted by some people outside the Popperian context (Hempel 1965). But the connection to Popperian themes is clear, and it is surely no accident that well-focused versions of the idea were developed by those around Popper's circle.

Since the original discussions, this idea has developed in more detail by various others. Horwich (1982) embeds it within a Bayesian approach. Giere (1970) and I embed it within an approach inspired more by classical statistical ideas (Godfrey-Smith 2003). But what is striking here is that the literature has not, in general, simply accepted and

absorbed the point. The more common response to the situation has been to try to salvage something closer to Hempel's original picture, according to which all instances support a generalization as a matter of quasi-logical fact. The tendency has often been to treat the considerations discussed by Watkins and others as extraneous details that concern special cases and do not get to the heart of the issue. From the Popper-informed point of view, though, these considerations are absolutely the heart of the issue. Favorable observations or positive instances may or may not have any real evidential significance in relation to a hypothesis we are considering. Whether they do or not depends on much more than what the hypothesis says and what was observed. It is also essential to consider the context in which the observation was made, for that is essential to determining whether or not the observation was part of a genuine test.

I call this point "Popper-informed," even though in one way it is anti-Popperian. The aim of this move is to resurrect a distinction between cases where observed instances do, and do not, provide epistemic support for a generalization. For Popper, such concepts of non-deductive support are misconceived (putting to one side, again, some complexities involving corroboration). I take the point to be importantly linked to Popper because of its connection to the ideas of risk-taking and possible disconfirmation.

I do not want to overstate the extent to which mainstream discussion rejected the Popper-informed point. Modern Bayesian treatments of the ravens hypothesis respect, more or less, the key constraint. For a Bayesian, evidential support is a contrastive matter. An observation only supports H if there are other hypotheses on the table that treat the observation as more unlikely that H does.[4] Those outside the Bayesian framework seem often to want to hang onto the idea that all instances confirm, however, and even some Bayesians seem to think this is a desirable output for a theory of evidence. To me, there is nothing whatsoever to be said for the idea that all instances, *qua* instances, confirm a generalization. Even the hedged versions of the view (the hypothesis is not already falsified, and there is no special background knowledge of the kind that Good (1964) discussed) should have no appeal. We learn from Popper that there is such a thing as the *illusion of support*, in many non-deductive contexts, and this is a very important illusion.

12

Why has this idea been so unobvious to those thinking within the mainstream empiricist tradition? I conjecture that this is because of two distinct factors.

One, which is fairly obvious, is the great influence that formal logic had on 20th century empiricism. If logic is our tool and our exemplar, then it will be natural to treat confirmation as a logical or quasi-logical relation. That means it will probably be a relation between sentences themselves – between the sentences describing observations and those expressing a theoretical hypothesis. Once we are committed to that framework, it seems attractive to do as much as possible with the relationship between universally quantified generalizations and statements reporting the satisfaction of generalizations by particular cases. If our goal is a logical theory of non-deductive support, this seems as simple and clear as things could get.

The second reason has to do with the other great influence on mainstream empiricism, a tradition of general theorizing about mind and language. The influence of psychologistic assumptions on empiricism was more muted in the 20th century than the 19th, but I think it still played a role. If we approach epistemology within the kind of psychological picture associated with traditional empiricism, then it might seem very natural to insist that that all positive instances confirm a generalization. The mind is seen as acquiring its epistemologically-sanctioned *contents* via the conduit of experience. One *sees* each positive instance – where seeing is a local and particular matter – and the world thereby impresses itself upon the mind. Each episode of this kind will and should increase the confidence the agent has in the generalization. And if these episodes do not suffice, nothing else is available to do the job.

So within the mindset of both 20th century logic-based epistemology and more psychologistic forms of empiricism, it can be quite hard to move away from the intuition that all positive instances confirm. But as the Popperian tradition, and some later work, have taught us, if we insist on this we fail to realize the importance of a simple but crucial kind of illusion of support, involving cases where a positive instance has been observed in the wrong context to have any epistemological significance. From the viewpoint of Popper's epistemology, which is free from the psychologistic assumptions of traditional

empiricism and less committed to the primacy of formal logic, it is easy to appreciate the importance of this fact.

5. The Diachronic Perspective on Evidence

The last theme I will discuss is more complicated and controversial. I will look at the possibility that Popper was seeing something important in some of the most-criticized parts of his work. Here, again, I have in mind his work on testing, and this time the focus will be on Popper's rejection of what look like moderate and reasonable concepts of evidential support. Without endorsing Popper's actual claims here, I want to raise the possibility that Popper was successfully seeing past some standard ways of thinking, and glimpsing significant new options. My discussion of these ideas will be cautious and qualified.

In this section I will use a distinction between *synchronic* and *diachronic* perspectives on evidence. A synchronic theory would describe relations of support within a belief system at a time. A diachronic theory would describe changes over time. It seems reasonable to want to have both kinds of theory. In the case of deductive relationships, formal logic gives us a detailed synchronic account, that can also be the basis of diachronic descriptions of valid reasoning. In the non-deductive case, epistemology in the 20th century tended to suppose we could have both kinds of theory, but often with primacy given to the synchronic side. The more novel possibility, which I will discuss in this section, is the primacy of the diachronic side, once we leave the deductive domain.

This idea of "primacy" gestures towards a family of ideas and options. A very moderate version of such a view appeared in the previous section. There I looked at the idea that an observation only has the capacity to support a hypothesis in the context of a test or procedure, where a "test or procedure" is something that extends over time. Musgrave (1974) extended these ideas, taking them to motivate a "party historical" or "logico-historical" approach to confirmation. There are also more radical views in the same family; perhaps there is no substantive non-deductive synchronic theory of

evidence possible *at all.* Or perhaps the only synchronic theory that can be given is much weaker and less rich than people have supposed.

A diachronic view of this kind would describe rational or justified change, or movement, in belief systems. The assessment is of motions rather than locations. Such a theory might enable us to recapture some, but not all, of what people wanted from the traditional synchronic account.

In this section I suppose that we do not, at present, have the right framework for developing such a view. But we can trace a tradition of sketches, inklings, and glimpses of such a view in a minority tradition within late 19th and 20th century epistemology. The main figures I have in mind here are Peirce (1878), Reichenbach (1938), and Popper. This feature of Popper's view is visible especially in a context where he gets into apparent trouble. This is the question of the epistemic status of well-tested scientific theories that have survived many attempts to refute them. Philosophers usually want to say, in these cases, that the theory has not been proven, but it has been shown to have some other desirable epistemic property. The theory has been confirmed; it is well-supported; we would be justified in having a reasonably high degree of confidence in its truth.

In situations like this, Popper always seemed to be saying something inadequate. For Popper, we cannot regard the theory as confirmed or justified. It has survived testing to date, but it remains provisional. The right thing to do is test it further.

So when Popper is asked a question about the present snapshot, about where we are now, he answers in terms of how we *got* to our present location and how we should *move on* from there in the future. The only thing Popper will say about the snapshot is that our present theoretical conjectures are not *inconsistent* with some accepted piece of data. That is saying something, but it is very weak. So in Popper we have a weak synchronic constraint, and a richer and more specific theory of movements. What we can say about our current conjecture is that it is embedded in a good process.

In some ways, this development in the shape of epistemological theory is not as alien as it might initially look. Something like this moral, in a moderate form, is implicit in standard versions of Bayesianism, the dominant view in philosophy of science about testing and evidence at present (Howson and Urbach 1993). (The next page or so is more

complicated than the rest of the paper and can be skipped by those unfamiliar with Bayesianism. See also note 4.)

Bayesianism is often seen as completely transforming the issues that people like Popper, Hempel and Carnap were concerned with. Further, it is often seen as coming down on the side of the optimists about traditional philosophical notions of confirmation and induction. But this standard story is not entirely accurate. Bayesianism of the standard kind treats belief systems both synchronically and diachronically. Constraints are placed on belief profiles at a time, and also on change over time as new evidence comes in. But Bayesianism imposes *weak* synchronic constraints, and *stronger* diachronic ones.

The synchronic constraint is often called "coherence." (This includes Popper's constraint of deductive consistency). Degrees of belief must obey the probability calculus. This is a weak constraint; all sorts of very unreasonable-looking belief profiles meet it. You can coherently believe, for example, that the coin in my hand is almost perfectly symmetrical and about to be tossed high with plenty of spin by a normal human, and that it will almost certainly come up tails. Substantive principles that link or coordinate subjective probabilities with objective chances or their physical bases are often discussed, but they are controversial and not intrinsic to the basic Bayesian picture.

The standard Bayesian framework imposes much richer constraints on how a rational agent's belief system changes over time. The updating of subjective probabilities must be done via conditionalization. When a Bayesian agent learns some new piece of evidence e, the agent's new unconditional probability for any hypothesis P'(h) must be set equal to the agent's old conditional probability P(h|e), which is related to various other old subjective probabilities via Bayes' theorem. This diachronic constraint is much richer than the coherence constraint that applies to "synchronic" assessment of belief systems.

Some (though not all) advocates of Bayesianism urge the importance of a larger-scale diachronic perspective here. Standard Bayesianism allows the initial "prior" probabilities of hypotheses to be "freely" set, so long as updating is done properly. The results of updating in the short term then depend on these assignments of prior probability. But Bayesians often argue that in the long term, after many rounds of

updating, the initial settings of priors "wash out." Two agents with very different distributions of prior probability for a set of hypotheses, who also satisfy some other constraints, will come to have degrees of belief that "converge," once enough evidence has come in. The details and significance of these "convergence" arguments are complicated, and not all Bayesians put much stock in them. But for some, an essential part of the Bayesian story concerns these larger-scale dynamic patterns in how bodies of incoming data are handled.

Leaving Bayesianism now, the general possibility on the table is the idea that many traditional epistemological questions might be recast in a way that treats diachronic features of evidence as primary. The discussion below will be more informal than a Bayesian treatment. The discussion will also be idealized in several ways, to bring it close to the general picture of scientific theorizing that Popper, Hempel and others tended to assume in the 20th century. This is a picture that I would in many ways reject, but it is common in philosophy. According to this picture, what scientists aims to develop are "theories," often in the form of generalizations. These theories are assessed for empirical adequacy, and when a theory does well under testing scientist can hope the theory might be true. Theories are discarded or modified when their predictions fail.

This may look like a fairly neutral account of scientific work, but it is very much a creature of philosophy itself. A more accurate view would include ideas such as these: a lot of theoretical science is concerned with *models* rather than "theories." Modeling involves the description and investigation of deliberately idealized, hypothetical structures that can have a range of different resemblance relations to real-world systems (Giere 1988). Scientists often retain multiple models when dealing with a single domain, including deliberately oversimplified ones that are known to be inaccurate in many ways. Traditional philosophical questions about truth and reference are not straightforwardly applicable to scientific work of this kind, and questions about evidence and testing look different as well. But having registered these qualifications, in most of the discussion below I operate within the confines of the standard picture of scientific activity that Popper and many other philosophers assume. This is, in a way, a piece of model-building of its own.

Let us now look at some ways in which some problems in epistemology might be transformed by adopting a strongly diachronic view of evidence and testing. I will discuss three issues in turn: conservativism, simplicity, and the underdetermination of theory by evidence.

Conservativism: In many discussions in philosophy of science, "conservativism" is seen as an epistemic virtue. We are told that it is reasonable not to alter or discard our theories unless there is a positive reason to do so. When we are induced to make changes to our theories, we should not change more than we have to. Quine called this principle the "maxim of minimum mutilation" (1990 p. 14).

Why should conservatism be a virtue? If our goal is to believe theories that are true, or that are well supported by the evidence, then it is hard to see why conservativism should be a good thing. If we take a snapshot of our current theory and its relation to current evidence, and we then note that some other theory does equally well with this body of evidence, why should the "incumbent" theory get an advantage?

One reason why philosophers are attracted to principles of conservativism is the fact that they seem to fit to some extent with scientific practice. It is also true that the principle can be justified in part on pragmatic grounds; it will generally be inconvenient to change from one theoretical framework to another. Kuhn emphasized this pragmatic side, when he said that "retooling is an extravagance reserved for the occasion that demands it" (1970 p. 76). But the kind of pragmatic role being envisaged here seems to be one that is at odds with, rather than being independent of, epistemic considerations.

From the point of view of a diachronic view of evidence, the role of conservativism looks different. It comes to have something closer to an epistemic justification. Or perhaps it would be fairer to say that from a diachronic point of view, there is a role for conservativism that is in some ways pragmatic but is positively tied to the epistemic rather than being at odds with it.

Suppose we have a view according to which the epistemic credentials of a theory derive from its embedding in an ongoing process. We see science as committed to a particular process of the development of theory in response to observation; our present

theory is just where this process has brought us for the time being. To drop that theory and adopt some other theory would be to depart from what we take to be the best process.

Perhaps this move does not really resolve the problem of the status of conservativism. Someone might object: why is it not *just as good a process* to arbitrarily switch to a different theory once it is shown to be compatible with all available evidence? In general, the best process is to take a theory and modify it as evidence comes in, but in the special case where we note the existence of a rival theory that has all the same relations to evidence, then we can freely switch.

There will certainly be good practical reasons not to follow such a procedure. The result will be needless retooling, and disruption of the development of key theoretical concepts. So perhaps what the shift to a diachronic perspective does in this first case is bring the more "practical" and more purely "epistemic" considerations into some kind of concord, rather than having them pull in different directions.

Simplicity: My second example is the problem of simplicity preferences, or "Occam's razor." Scientists are usually seen as preferring simple to complex theories whenever possible, and they are usually seen as justified in this preference. This is a more famous, more important, and more vexed issue than conservativism.

Occamism has been very hard to justify on epistemological grounds. Why should we think that the a simpler theory is more likely to be true? Once again there can be an appeal to pragmatic considerations, but again they seem very unhelpful with the epistemological questions.

This problem has connections with the problem of conservativism discussed above. Sometimes the preference for existing theories is itself described as an Occamist preference, but this can be misleading. We might have a situation where an incumbent theory is more complex than a newcomer, and both deal with the evidence equally well. Then conservativism tells us not to shift, and simplicity pulls us the other way.

The right eventual philosophical response to the problem of Occamism will surely contain a mixture of elements. First, there is an important range of cases in science where substantive assumptions about the objects and processes being studied can justify a

limited and field-specific preference for simplicity (Sober 1988). Secondly, this is one of the areas where philosophical discussion is partially out of step with scientific practice. In some areas in science, the response to the development of a simpler theory that handles the same observations as an existing theory would be to keep *both* theories on the table. This is one of the areas (noted earlier) where a philosophers' obsession with theory *choice,* rather than the development of a range of useful models, can lead us astray. But let us bracket those aspects of the problem, and see whether things look different when we switch to a diachronic perspective on evidence and testing.

From a diachronic point of view, simplicity preferences take on a quite different role. Simplicity does not give us reason to believe a theory is true, but a simplicity preference is part of a good *rule of motion*. Our rule is to start simple and expect to get pushed elsewhere. Suppose instead we began with a more complex theory. It is no less likely to be true than the simple one, but the process of being pushed from old to new views by incoming data is less straightforward. Simple theories are good places from which to initiate the dynamic process that is characteristic of theory development in science. Occasionally a very simple theory might actually be true; that is merely a bonus.

This feature of simplicity preferences has been noted, in more specific contexts, before. Popper himself (1959, Chapter 7) argued that the importance of simplicity lies in the fact that simple statements are more easily falsified than complex ones, and he used the example of different mathematical relations between variables (linear, quadratic, etc) to make the point. More recently, Kevin Kelly (2004) has argued, on the basis of formal models, that a simplicity preference will be part of a procedure that reliably approaches the truth via the fewest dramatic changes of opinion *en route*.

Underdetermination: My third example is more subtle, and even more general than the problem of simplicity. This is the problem of the "underdetermination of theory by evidence." In simplest terms, this is argument that for any body of evidence, there will always be more than one theory that can, in principle, accommodate it. As observational evidence is all we have to go on, this seems to show that our preference for any specific theoretical view must always be based to some extent on non-evidential factors like

aesthetic considerations or convenience (see Psillos 1999, Chapter 8, for a review). This, in turn, is often taken to be a problem for scientific realism, particularly scientific realism of the "mixed" kind discussed in Section 3 above.

There are many versions of the underdetermination thesis, some stronger than others. Some versions are both very strong and very general; it is argued that for any theory T_1 we might come to hold, there will be another incompatible theory T_2 that that we cannot hope to empirically distinguish from T_1 via any conceivable evidence. But as Stanford (2006) argues, these very strong versions of the view tend to rely on extreme skeptical hypotheses (perhaps of the Cartesian demon kind), or on small manipulations of T_1 that produce a variant that is not scientifically interesting. There are worked-out illustrations of underdetermination for some physical theories, usually involving space, time, and motion, that are neither ultra-skeptical nor trivial, but certainly not for all theories.

The statement of an underdetermination problem that I will focus on is more moderate, but still important. My formulation is modified from one used by Psillos (1999, p. 164).

U: For any particular body of evidence we might have, there will always be more than one scientific theory that can, in principle, accommodate it.

In my discussion here I will bracket some questions about the role of probability and confirmation. To say that more than one theory can "accommodate" the data is not saying much, as two theories may both permit various observations but assign very different probabilities to them. But I won't worry about that complication here.

Suppose U is true. How worrying is it? My argument will be that its importance is sometimes over-stated because of philosophers' routine assumption of a particular point of view. The usual situation imagined is one in which we assume we have a body of data and a theory T_1 on the table. Principle U then appears in the form of a kind of *barrier* to successful theorising. But so far at least, U is compatible with another principle that might apply to the situation.

D: For any particular comparison of theories we want to make, there is some possible body of data that will discriminate the two.

That is, many of the usual underdetermination anxieties are compatible with a kind of symmetry in the situation: for any comparison of theories, we can hope to find discriminating data; for any data, there will be rival theories that are not discriminated.

Of course, D might be false. Once we bring in traditional skeptical possibilities, it seems that it may well be false. But most philosophy of science discussion is not supposed to be concerned with those extreme possibilities. Perhaps in the case of some specific hypotheses or scientific domains, D is again a vain hope. But that, again, is not the usual focus or thrust of the discussion. Something like U alone is often seen as sufficient to endanger realism.

Suppose U and D are both true, or have similar standing. Then we have a "glass half full" and "glass half empty" situation. When we look at U, the glass looks half empty. When we look at D, it seems half full. What must be resisted, or done more cautiously, is the drawing of conclusions solely from the "glass half empty" side. And here the diachronic point of view is relevant. The glass looks half empty when we think about the problem synchronically in a particular way. If we take a snapshot of the data present at a time, and ask which theoretical possibilities it can distinguish, then we will note the data's limited power. If, however, we think diachronically about what data can do, the situation looks different.

I cannot claim too tight a connection between the diachronic point of view and a glass-half-full impression of the situation. We could, in principle, think diachronically about the introduction of new theoretical possibilities. If at any time we have succeeded in using our data to discriminate between T_1 and T_2, we can expect someone to put T_3, a new theoretical possibility, on the table at the next time-step, that our data cannot rule out. But it seems easier or more natural for people to think diachronically about a flow of data, rather than a flow of theoretical options being put on the table. And most of the time

people seem to think simply in terms of a snapshot where we hold our data fixed and lament its limited powers of discrimination.

I do not say all this in order to urge that people look only at the bright side, the glass-half-full side. But clearly, we should look at the situation from that side as well. Principles like U and D should be assessed as pairs, when we look for general philosophical morals. The question of what we *should* conclude from the pair of D and U, if both are true, I leave for another occasion.

This concludes my discussion of three epistemological issues that may look different from a diachronic point of view. I emphasize that a *purely* diachronic view in this area seems to lead to bad consequences. For example, consider our present epistemic situation regarding evolution and creationism. When describing this case, it seems misleading and incomplete to say merely that the evolutionary view of life on earth was arrived at by a good process and that we expect to refine it rationally via more good "motions". In cases like that, it seems undeniable that we have good reason to believe that one theory, seen as a snapshot, is highly likely to be true, at least with respect to the basic features. We do not always expect to move on.

As has often been noted, the Popperian view of evidence is strongly informed by particular parts of science – the collapse of Newtonianism, the heroic conjectures of Einstein, the permanent sense in 20th century physics that more surprises may be just round the corner. Not all science is like this. No one who started their epistemology by thinking about 20th century biology would be led to a picture with this overall shape to it. So a future view must somehow strike a balance here.

6. Conclusion

I have discussed four Popperian themes: the importance of eliminative inference, skeptical realism and neighboring possibilities, the link between risk-taking and evidence, and (more speculatively) the general viability of a more diachronic perspective on problems in epistemology. The first two of these ideas are straightforward. Philosophical

neglect of the first, in particular, looks very strange in retrospect. The second two ideas are more subtle, and their future importance is more uncertain. Especially in these latter cases, my emphasis has not been on the details of Popper's arguments, but on broad possibilities, philosophical priorities, and directions of analysis that he seemed to see when others did not.

<p style="text-align:center">*   *   *</p>

References

Boyd, R. (1982). "On the Current Status of the Issue of Scientific Realism." *Erkenntnis* 19: 45-90.

Devitt, M. (1997). *Realism and Truth.* (2nd ed.) Princeton: Princeton University Press.

Devitt, M. (2005). "Scientific Realism." In F. Jackson and M. Smith, (eds.), *The Oxford Handbook of Contemporary Philosophy*, Oxford: Oxford University Press, pp. 767-91.

Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.

Forber, P. (2006). "The Traces of Change: Evidence in Evolutionary Biology." PhD Dissertation, Department of Philosophy, Stanford University.

Giere, R. (1970). "An Orthodox Statistical Resolution of the Paradox of Confirmation." *Philosophy of Science* 37: 354-362.

Giere, R. (1988). *Explaining Science: A Cognitive Approach.* Chicago: University of Chicago Press.

Godfrey-Smith, P. (2003). *Theory and Reality: An Introduction to the Philosophy of Science*. Chicago: University of Chicago Press.

Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

Good, I. J. (1967). "The White Shoe Is a Red Herring." *British Journal for the Philosophy of Science* 17:322.

Hájek, A. (2003a) "Interpretations of Probability." *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/probability-interpret/>

Hájek, A. (2003b). "What Conditional Probability Could Not Be." *Synthese* 137: 273-323.

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

Horwich, P. (1982). *Probability and Evidence*. Cambridge: Cambridge University Press.

Howson, C. and P. Urbach. (1993). *Scientific Reasoning: The Bayesian Approach*. Second edition. Chicago: Open Court.

Hull, D. (1999). "The Use and Abuse of Sir Karl Popper." *Biology and Philosophy* 14: 481-504.

Kelly, K. (2004). "Justification as Truth-Finding Efficiency: How Ockham's Razor Works." *Minds and Machines* 14: 485-505.

Kitcher, P. S. (1993). *The Advancement of Science.* Oxford: Oxford University Press.

Kuhn, T. S. (1970). *The Structure of Scientific Revolutions*. (2nd ed.) Chicago: University of Chicago Press.

Laudan, L. (1981). "A Confutation of Convergent Realism." *Philosophy of Science* 48: 19-49

McMullin, E. (1984). "A Case for Scientific Realism." In J. Leplin, (ed.) *Scientific Realism*. Berkeley: University of California Press.

Musgrave, A. (1974). "Logical Versus Historical Theories of Confirmation." *British Journal for the Philosophy of Science* 25: 1-23.

Newton-Smith W. (1981). *The Rationality of Science.* Boston: Routledge & Kegan Paul.

Peirce, C. S. (1878). "How to Make Our Ideas Clear." *Popular Science Monthly* 12: 286-302.

Platt, J. (1964). "Strong Inference." *Science* 146: 347-353.

Popper, K. R. (1959). *The Logic of Scientific Discovery*. New York: Basic Books.

Psillos, S. (1999). *Scientific Realism: How Science Tracks Truth*. London: Routledge.

Quine, W.V.O. (1990). *Pursuit of Truth*. Cambridge, MA: Harvard University Press.

Reichenbach, H. (1938). *Experience and Prediction: An Analysis of the Foundations and the Structure of Knowledge*. Chicago*:* University of Chicago Press.

Salmon, W. (1981). "Rational Prediction." *British Journal for the Philosophy of Science* 32: 115-125.

Sober, E. (1988). *Reconstructing the Past: Parsimony, Evolution, and Inference*. Cambridge, MA: MIT Press.

Stanford, K. (2006). *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.

Van Fraassen, B. (1980). *The Scientific Image.* Oxford: Oxford University Press.

Watkins, J. (1964). "Confirmation, Paradox and Positivism." In M. Bunge (ed.), *The Critical Approach to Science and Philosophy*. New York: The Free Press.

Weisberg, M. (forthcoming). "Who is a Modeler?" *British Journal for the Philosophy of Science*.

Notes

[1] See Salmon (1981), Newton-Smith (1981), Godfrey-Smith (2003).

[2] Van Fraassen (1980) is a special case. Van Fraassen argued that the properly scientific attitude to have to a successful theory is to "accept it," where acceptance does not include belief in the claims the theory makes about the unobservable domain. Realism is characterized as a view about the proper *aim* of science, not (explicitly at least) as a view about the chances of success. See Godfrey-Smith (2003) for further discussion.

[3] This will not be the case if law-like generalizations in science are seen as subjunctive conditionals or some other quirky form of conditional.

[4] The Bayesian model of rational belief change characterizes agents as holding sets of *degrees of belief* in various hypotheses. At any time, a rational or "coherent" agent's degrees of belief must be related in ways that conform to the axioms of probability theory. As new evidence comes in, the agent updates via "conditionalization." Roughly speaking, if the agent observes *e*, the agent's new degree of belief in *H* is set equal to the agent's old degree of belief in *H given e*. See Howson and Urbach (1993) for a good introduction.