

PHILOSOPHY OF SCIENCE QUA EPISTEMOLOGY

Forthcoming in *Continuum Companion to the Philosophy of Science*

(eds Steven French and Juha Saatsi)

Alexander Bird

1 Philosophy of science between generalism and particularism

The relationship between philosophy of science and epistemology as practised at the heart of general philosophy has been variable. Philosophy of science is caught between potentially opposing forces. On the one hand philosophy of science needs to be true to the (at least apparently) distinctive and even arcane practices of actual scientists. This I call the *particularist* tendency, because it tends to emphasize the particular, special nature of science (and maybe even of the individual sciences). On the other hand philosophy of science needs to relate its account of scientific belief to the entirely general account of knowledge and justification provided by epistemology. This I call the *generalist* tendency, because it seeks to place the philosophy of science within a general epistemological framework.

Modern philosophy of science emerged in the middle of the nineteenth century. The most significant work of philosophy of science since Bacon's *Novum Organon*, William Whewell's *The Philosophy of the Inductive Sciences, Founded Upon Their History* (1840) exhibits the particularist tendency. As the title of his book indicates, Whewell's philosophy of science is built on his earlier work *The History of the Inductive Sciences* (1837). While he does construct his account within a general framework that bears a superficial resemblance to Kant's, Whewell's description of science and its processes of reasoning are clearly specific to science and are motivated in part by reflection on the details of particular episodes in the history of science. For example, he uses Kepler's discovery of the elliptical orbits of the planets as an example of discoverer's induction: the unfolding of fundamental ideas and conceptions permits the 'colligation' of the observed facts concerning the orbit of Mars to the conception of those facts as satisfying an elliptical orbit, which is then generalized for all planets. Whewell insisted that a philosophy of science must be inferable from its history. As a consequence Whewell's account is detailed, and permits multiple routes to discovery and knowledge. Discoverer's induction is itself a multi-stage process, and the colligation component may be achieved by a number of different inferences. Likewise the process of confirmation can involve three distinct components: prediction, consilience, and coherence. Whewell addresses questions that are still much discussed in philosophy of science, such as whether novel predictions have greater

value than the accommodation of old data—questions which have been scarcely addressed in general epistemology either then or since.

While Whewell came to philosophy of science as a polymath scientist (he invented the very term ‘scientist’), John Stuart Mill came from the other direction, that of a general philosopher. Mill’s *A System of Logic* (1843) demonstrated the generalist tendency, seeking a unified account of all reasoning within an avowedly empiricist framework. For Mill our inferential practices in science are essentially either enumerative or eliminative in nature (the latter being really deductive). Mill is generally sceptical regarding the methods Whewell describes, such as those that allow for inferences to the unobservable. All satisfactory science depends on enumerative induction, which is fallible, allied with elimination, which tends to reduce, if not remove, the fallibility of the induced hypotheses. For Mill, the inductions of science are extensions of the spontaneous inductions we naturally make. Science aims to improve on these spontaneous inductions, making inductions of a less fallible variety. But the pattern of inference is essentially the same. That being the case, Mill needed only a few, and simple examples, to illustrate his case. Whewell criticized Mill on precisely this point, that the paucity of his case studies suggests that Mill’s philosophy of science, unlike his own, could not be inferred from the study of its history.

Mill’s generalism was dominant in the first part of the twentieth century. The logical positivists’ conception of the relationship between scientific and everyday knowledge is best summarized in Einstein’s famous comment that “the whole of science is nothing more than a refinement of everyday thinking”. To refine something is to remove impurities; the basic substance is unchanged. The logical positivists took the view that there is no fundamental difference between epistemology as applied to science and epistemology as applied to ordinary instances of knowledge and belief. In his *General Theory of Knowledge*, Moritz Schlick asserts, “knowing in science and knowing in ordinary life are essentially the same” (1985: 9). Similar statements may be found in the work of other positivists, and beyond (from James Dewey and George Santayana for example). For Schlick, science and everyday knowing differ only in their subject matter. It is natural that empiricists such as Mill and the positivists should take such a view. After all, empiricism emphasizes the foundational role of individual experience. Furthermore it is sceptical (to some degree or other) of inferences that take us much beyond that experience. Consequently, they take the foundation and even the extent of all knowledge, scientific and everyday, to be fixed by something—individual sensory experience—that is clearly a central concern of general epistemology and not something abstruse and specific to science.

A corollary of the “science is refined everyday thinking” dictum is that everyday thinking is scientific thinking with impurities. Because science and everyday thinking are essentially the same, except that science is purer, an understanding of their common epistemological foundation is best gained by looking at science. Thus, for the positivists, epistemology just *is* philosophy of science.

As philosophy entered a post-positivist phase in the second half of the twentieth century, the particularist tendency returned to the ascendancy, with an emphasis on the significance of the history of science that resembles Whewell’s work rather than Mill’s. Thomas Kuhn gives an account of the development of science that is not deduced from some general epistemology but is inferred from the historical facts. Although not an epistemology of science, Kuhn’s picture is epistemologically significant. First, he rejects a simplistic, empiricist conception of observation. Perceptual experience cannot be the foundation of science, because it itself is theory-laden.

And observation, the process of generating scientific data, is not simply a matter of perception, “The operations and measurements that a scientist undertakes in the laboratory are not ‘the given’ of experience but rather ‘the collected with difficulty.’” (Kuhn 1970b: 126). Secondly, whereas Carnap (1950) took inductive logic to provide a generally applicable, rule-based approach to scientific reasoning, Kuhn argued that scientific cognition is driven by a scientist’s learned sense of similarity between the scientific problem she faces and some concrete exemplar—an earlier, exemplary piece of scientific problem-solving. Thirdly, in the positivists’ picture theory-neutral perceptual experience is supposed to provide a basis for the meaning of our theoretical terminology; but in the light of the first point, that basis is not theory-neutral at all. The languages of different scientific communities, working within different paradigms (i.e. employing different sets of exemplars), will not be immediately and simply inter-translatable—they will be incommensurable. Together these points indicate an epistemological relativism. We do not have an unproblematic, agreed basis for settling scientific disputes. The language disputants use may be incommensurable, the data may be laden with theories that are themselves subject of dispute, and our assessments of theories may be determined by differing exemplars. Such problems do not arise in an obvious way in the epistemology of everyday knowledge.

The debates between Popper (1959, 1970), Kuhn (1970a), and Lakatos (1970) illustrate the changing nature of the perceived relationship between the philosophy of science and the history of science. Although Sir Karl Popper was not a logical positivist, he shared many of the beliefs of logical empiricism. A key problem for Popper was a central issue in general epistemology—Hume’s problem of induction. It was also a central issue for others, such as Carnap (1950) and Reichenbach (1971), but whereas they had some hope of providing a probabilistic solution to the problem, Popper took it to be insoluble. This fact, of being an inductive sceptic, combined with the view that science is rational, is the defining feature of Popper’s philosophy. The latter can be seen as the attempt to reconcile inductive scepticism with scientific rationalism. While Popper’s work is scientifically well-informed, and does refer to episodes in the history of science, it is clear that these have only a minimal evidential role. Popper’s falsificationism—his rationally acceptable alternative to inductivism—is primarily normative in character. But since Popper holds that science largely meets the norms of rationality, his combined view is answerable to history, for the majority of scientists should behave in accordance with Popper’s falsificationist method. Kuhn’s complaint is that they do not. According to Popper, scientists ought to reject a theory as soon as they discover a piece of evidence inconsistent with it. But Kuhn’s historically well-attested account of normal science shows that they do not do this: a theory can accumulate anomalies, including data flatly inconsistent with the theory, without being rejected. Kuhn emphasizes that a theory only begins to be questioned when anomalies of a particularly serious sort accumulate that resist the attempts of the best scientists to resolve them. One aim of Imre Lakatos’s ‘methodology of scientific research programmes’ is to reconcile a falsificationist approach with the historical record. Lakatos (incorrectly, in my view) held that Kuhn’s account of scientific development would, if correct, imply that science is irrational. Thus while Popper’s view was founded in epistemology, with a relative disregard for historical detail, and Kuhn’s view was historically well-informed but disengaged from general epistemology (there is just one, rather off-hand reference to the problem of induction in *The Structure of Scientific Revolutions*), Lakatos seeks to combine history with a normative approach derived from Popper’s general

epistemology. Famously, Lakatos (1971: 91) adapted Kant’s words thus “Philosophy of science without history of science is empty; history of science without philosophy of science is blind.”

I suggested above that a thoroughgoing empiricist philosophy of science is likely to be close to a general empiricist epistemology, because of (i) the central place given to perceptual experience, and (ii) the disinclination of empiricists to go much beyond the evidence of our senses. Kuhn, and others such as N. R. Hanson (1958), emphasize the problematic nature of (i), in a manner that exacerbates the anti-realist tendency of logical empiricism. At the same time a return to scientific realism in the philosophy of science (e.g. Grover Maxwell 1962) questioned (ii). If science aims to get to know not just about perceptible items but also about the unobserved, then epistemological questions may need to be asked about scientific cognition that do not apply to a central, everyday instance of cognition. Pursuing this line of thought, it is open to the scientific realist to ask whether, since science has a subject matter that differs from the objects of everyday knowledge, the inference patterns of science are also distinctive.

2 Scientific reasoning

Even among those who do think that science may employ forms of reasoning that are not found in everyday cognition, the opposing generalist and particularist tendencies is apparent. A long-standing disposition among many philosophers of science (and even more scientists) is to hold that the successful development of science has the *scientific method* to thank. The scientific method, which, one might hold, emerged during the period of the scientific revolution, is a single approach to scientific investigation and inference that underlies all good science in all domains. Belief in the existence of a single scientific method exemplifies the generalist tendency. Opposed to such a view, a particularist might deny that there is a single scientific method; instead there are many methods, reflecting the diversity of scientific theories and subjects.

If there is a scientific method, it must be largely tacit, because there is little agreement on what exactly the scientific method is. Since it is tacit, it is the job of the philosopher of science to reconstruct it, showing on the one hand how it is in fact used in successful science and on the other hand that its success is explained by some general epistemological justification. Lakatos’s methodology of scientific research programmes provides one account of the scientific method in action, but not the only one; Carnap’s inductive logic may be considered another.

2.1 Bayesian epistemology

Of contemporary accounts of a single scientific method, the clear front runner is Bayesianism. Bayes’s theorem tells us how the conditional probability, $P(h|e)$ of a hypothesis h , given the evidence e , is related to the probability of the hypothesis independently of the evidence, $P(h)$, the probability of the evidence, $P(e)$, and the probability of the evidence given the hypothesis, $P(e|h)$:

$$(B) \quad P(h|e) = \frac{P(e|h)P(h)}{P(e)}$$

Traditional Bayesianism interprets these probabilities as credences—subjective degrees of belief. On learning that e is in fact the case, the Bayesian says that one is

rationally obliged to update one's credence in h so that it is equal to $P(h|e)$ as given in (B). This is known as *conditionalization*. Bayesians argue that Bayesian conditionalization can resolve a whole host of problems and issues in the epistemology of science, such as the ravens paradox, the problem of induction, and the Duhem problem (Howson and Urbach 2006).

What is unusual about Bayesianism is that it is an epistemological tool or approach that was developed principally within the philosophy of science, and within science itself, and only thereafter transformed into a general epistemology, applicable not only to scientific but also to more general epistemological problems (Bovens and Hartmann 2003). There are a number of reasons for this. First, orthodox, subjectivist Bayesianism concerns fairly minimal rational constraints on credences, subjective degrees of belief. It does not tell us what those degrees of belief ought to be, only what they ought to be given previous degrees of belief. Consequently Bayesianism does not provide any account of the relationship between belief and the world. But such relationships are at the centre of traditional epistemology; witness the centrality of concerns about scepticism for example. Thus Bayesian epistemology has no place for the concept of knowledge, which expresses precisely such a relation. Arguably it cannot account for the concept of justification. This point is more disputable, because unlike the concept of knowledge, the concept of justification provides no direct implication of a relationship between beliefs and the world. One might argue that a subject's degrees of belief are justified precisely when updated in accordance with Bayesian conditionalization. On the other hand, one might think that the notion of justification cannot be entirely divorced from the notion of truth: justification ought normally to be truth-tropic whereas Bayesian conditionalization is not except in a very weak sense. Updating in accordance with Bayesian conditionalization puts a constraint on possible credences that is much like the requirement of logical consistency in the case of full-on belief. Such a constraint helps avoid some falsehoods, but does not go beyond that in pointing towards the truth. If one had chosen, before the process of conditionalization began, a sufficiently way-out set of credences, then one can obey conditionalization impeccably over a large data set, and still end up with credences that are far from reflecting the truth. Bayesians may point to the phenomenon of the 'washing out of priors'. Subjects with differing priors who conditionalize on the same set of evidence will tend to converge in their credences. But the actual credences will only actually converge, i.e. reach the same value (within some small tolerance), in the long run, which may be a very long run indeed, if the initial priors differ enough. In the short and medium run, the credences of the subjects may still differ significantly. And the question of justification concerns one's *current* credences: is *this* subjective degree of belief justified?

Note that this case is not analogous to the possibility of justified false belief in traditional epistemology. A standard case is that of the subject who, looking at a convincing barn façade, as one might have found in a Potemkin village, believes there is a real barn ahead. In this case the subject's failure to believe correctly is due to factors outside the subject's control. On the fact of it, the subject did exactly what she ought to do: form a belief on the basis of the way things appear to be, there being no reason to suppose that things are unusual. That seems to be a good, if informal, explanation of why she is justified in believing that there is a barn in front of her. Consider, on the other hand, the subject whose starts with way-out initial priors and consequently now has credences that are far from the truth (i.e. he has high degrees of belief in falsehoods and low degrees of belief in truths). This subject is in that position not through any bad luck or fault with the evidence, but only because he

started out with bizarre initial priors. It is debatable whether the well-behaved conditionalizer with way-out priors is justified in his credences that radically misalign with the truth. The extent to which one is willing to ascribe justification in such cases will depend on how strong a link between belief and world one thinks is signalled by the concept of justification. As we shall see, epistemological internalists will be more sympathetic to the Bayesian case than epistemological externalists: internalists insist on the epistemological primacy of internal relations among beliefs whereas externalists prefer to focus on the belief–world relation.

2.2 Inference to the best explanation

In many respects *inference to the best explanation* (IBE) provides an important contrast to Bayesianism. Unlike Bayesianism, IBE is concerned with the objective truth of hypotheses rather than their subjective probability. Thus IBE connects more directly to issues in traditional epistemology than does Bayesianism. On the other hand, supporters of IBE as a model for scientific inference typically do not regard every inference in science as conforming to the model. For example, it is not clear how one might construe many statistical inferences in science as IBE; likewise inferences based on specific test procedures. One might attempt to force some of these into the IBE mould. For example, a litmus test may license the inference that a particular substance is acidic. This might be construed as an inference that the best explanation of the test result is that the substance is acidic. But many such Procrustean moves seem implausible, psychologically at least.

Instead, the IBE supporter should hold that only a particular subset of the inferences of science, albeit a particularly important one, exemplifies IBE. According to IBE, the truth of a hypothesis is licensed when that hypothesis provides a better potential explanation of the evidence than all its rivals (Harman 1965; Lipton 2004). That is, we select as being the *actual* explanation, an explanation from a pool of *potential* explanations that is (a) clearly better than all its rivals at explaining the evidence, and (b) is in itself a good enough explanation. IBE faces a number of challenges and demands for clarification. What counts as a ‘good’ explanation? Isn’t goodness, whatever it is, too subjective to be a measure of truth? How can we be sure that we have considered the actual explanation among the pool of potential explanation?

The most obvious and important of the objections to IBE is that which Lipton calls ‘Voltaire’s objection’. Why should explanatory goodness be correlated with the truth? For IBE to be an inference procedure that leads to the truth, the actual world must be the best of all possible worlds, explanation-wise. Why should we think that is the case? On the face of it, it doesn’t seem to be a necessary truth that the actual world is explanatorily better (e.g. simpler, more unified, more elegant, ‘lovelier’) than all other possible worlds. Consequently, if IBE is a reasonably reliable form of inference, that fact is a contingent fact. Furthermore, it does not look to be knowable *a priori*.

In the light of this one might ask, is it really the case that IBE can lead to knowledge, or even to justified belief? In using IBE, it appears that scientists are using a method that they appear to have no reason to think will lead to the truth. The case seems analogous to sceptical questions surrounding perception. On the (admittedly doubtful) assumption that we make inferences concerning our immediate physical surroundings from our sensory experiences, the sceptic points out that it is neither necessary nor *a priori* that there should be a match between the two. Hence the

inference is made upon a basis of a supposed match that the subject has no reason to believe is the case. How then can the inference lead to knowledge or even justification? Note again the contrast with Bayesianism. Bayes's theorem itself is clearly *a priori* and may be inferred with ease from the intuitively compelling axioms of probability. Conditionalization is more contentious. Even so, Bayesians can quite plausibly claim that conditionalization is *a priori*, and they bring forward arguments of an *a priori* kind to defend conditionalization. For example, Bayesians appeal to diachronic Dutch book arguments (e.g. Teller 1976). Such arguments aim to show that a subject who updates credences in a way that does not obey Bayesian conditionalization is susceptible to a Dutch book: accepting a set of bets that can be shown to lead to a loss, whatever the outcome. This shows, it is suggested, that failure to obey conditionalization is a kind violation of rationality.

3 Internalist and externalist epistemology

The contrast between Bayesian epistemology and IBE exemplifies a deep divide within epistemology, between epistemological internalists and epistemological externalists. There is more than one way of characterizing the difference, and not all ways are equivalent, but most ways will have something in common with the following. The internalist requires that if a subject uses some method, such as an inference pattern, to come to knowledge or justified belief, then the subject must have some reason to believe that the method in question is likely to lead to the truth. The externalist will deny this. Another characterization is to say that the internalist holds that the state of epistemic justification supervenes on the subject's intrinsic states—such as their sensations, perceptual experiences, beliefs, and credences (ignoring issues of content externalism). Denying this, the externalist will hold that in some cases the difference between a justified believer and an unjustified believer may be some further feature that is extrinsic to the subjects, such as some relation between the subject and the world.

We can see why there is a connection between Bayesianism and internalism on the one hand, and IBE and externalism on the other. The internalist thinks that justification is a matter of one's internal states, in particular an internal ability to discern the truth-conduciveness or rationality of one's methods. The Bayesian says that the truth of Bayes's theorem and the rationality of Bayesian conditionalization are *a priori*, that is discernible by internal reflection rather than empirical investigation. The externalist will say that in some cases the justification provided by a belief-forming method will depend on factors of which the subject may not be aware. For example, according to the reliabilist version of externalism (Goldman 1979), a justified belief is one that is brought about by a method that is reliable. A reliable method is one that produces, or tends to produce, true beliefs. A method can be reliable without anyone being aware of its reliability. Consequently a subject can have a justified belief without knowing that the belief was produced by a method that is reliable, so long as it is in fact reliable. For example, one might think that perception can justify one's beliefs or give knowledge without anyone having considered the reliability of their perceptual capacities. The point is most clear with regard to small children or animals. It would seem obvious that they have perceptual knowledge, although they are not sophisticated enough to consider, let alone evaluate, the reliability of their senses. What is relevant in deciding whether they do gain perceptual knowledge is the actual reliability of their senses, not their knowledge of that reliability. Returning

to IBE, perhaps the fact that one cannot know by any kind of reflection that explanatory goodness and truth are correlated does not matter—it may be sufficient for one to gain knowledge from IBE that there is in fact such a (non-accidental) correlation, as a result of which the process of inferring from the fact that a hypothesis is the best explanation of the evidence to the truth of the hypothesis, is a reliable form of inference.

Externalist epistemology as applied to science faces the following problems. Reliabilism, a leading form of externalism, seems most plausible when (a) regarded as a theory about knowledge, and (b) is applied to relatively basic (e.g. innate) forms of belief production; reliabilism looks good as a theory about perceptual knowledge. Correspondingly, reliabilism may seem rather less plausible when considered as a theory of justification, as applied to the sophisticated methods developed by science. Could someone be justified in believing in neutrinos on the basis of Pauli's arguments from the details of beta decay and the the conservation of mass, without understanding why those arguments justify the belief? It would seem that in science some degree of reflective awareness of the justificatory power of evidence and argument is required in order for the subject to be justified in holding a belief (when caused by that evidence and argument). Thus internalism would appear to be on stronger ground when we think about scientific ways of producing beliefs than when the topic is simple perceptual belief-forming methods. This may be one reason why the internalism–externalism debate has been more intense among general epistemologists than philosophers of science.

This evidence in favour of internalism should not be exaggerated, however. First, not that although simple reliabilism does not require reflective awareness of the belief-forming method employed, reliabilism is only one externalist theory. Externalism itself says that it is possible in some cases to know or to have justified belief without such awareness; it need not say that such awareness is never required. Secondly, it should be recognized that the individually accessible justification that scientists have in their beliefs is very limited. For example, a chemist justifies her belief in a favoured hypothesis by referring to the fact that electrodes in a certain solution produced a current of 5.2 mA. But were she to be asked what justifies her faith in that fact, she may be able to do no better than to tell us that this is the reading that the ammeter showed, and that the ammeter was made by a reputable manufacturer. In another case, a physicist may be able to justify belief in Einstein's general theory of relativity by referring to the details of the precession of the perihelion of Mercury. But he may be able to tell us nothing about how the astronomers calculated that precession nor about the observations they employed. In these cases the justification of a belief seems to depend ultimately on the reliability of some instrument or the testimony of other scientists. (The social nature of scientific evidence and the importance of trust is emphasized by John Hardwig 1991.) By the internalist's standards, this is insufficient to provide justification. One possible response would be to abandon internalism as a thesis about individual justification, but retain it as a thesis about group beliefs. The community as a whole is justified in the relevant beliefs because as a whole it can provide an explanation of the functioning of the ammeter or the basis for the belief about the perihelion of Mercury.

Thirdly, internalism does not say that the subject should be aware of *something* relevant to the justification of their beliefs, but *everything*. The justification of a belief should depend solely on what is available to the subject. This is important because we may wish to make a distinction between the reasons given that are specific to some particular case and the general form of the justification. Let us imagine (as

seems plausible) that Pauli's reasoning concerning beta decay is an instance of IBE. Certainly Pauli and anyone else who uses this reasoning to gain a justified belief in the existence of neutrinos ought not simply to be able to point to the relevant evidential facts (the observed loss of mass-energy during beta decay, the laws of conservation of mass). But they should also be able to articulate the idea that the existence of the neutrino provides a simple but powerful explanation of the observed loss of mass; and they ought also to be able to articulate, to some extent, why this is a better explanation than some obvious alternatives, such as the proposal that the law of conservation of mass is not strictly true. However, we would not expect the scientist to be able to articulate, on a more general level why the fact that the neutrino hypothesis possesses the property of being the best explanation we've come up with gives us reason to believe that hypothesis. So science differs from perception insofar as it does require for justified belief that the scientist be able to articulate reasons for her belief and to be able to provide some explanation how they support the belief. But that fact does not fully vindicate internalism, for we do not expect the scientist to be able to articulate why the general pattern exemplified by this case is a good one to use; why this pattern is a pattern that makes it that case that the specific reasons proffered by the scientist really are reasons for belief.

This line of thought then raises the following questions for the externalist thinking about scientific reasoning. In order to have justified beliefs, the subject is required to be able to justify herself—to give her reasons—to a fair degree, but not beyond. Where does the cut-off point come? Why does it come where it does? And why is sophisticated thinking different in this respect from more direct, mundane or innate ways of forming beliefs which do not seem to require self-justification? While these questions are significant for the philosophy of science in particular, they are ones for general epistemology to answer.

4 Naturalized epistemology and the philosophy of science

Naturalized epistemology is one manifestation of epistemological externalism (Quine 1969; Kornblith 2002). Cartesian epistemology takes it to be the task of the epistemologist to show how we can come to know things from the vantage point of the subject whose basic resources are limited to his or her reflective capacities plus the sensory impressions with which he or she is presented. An optimistic Cartesian epistemologist, such as Descartes himself, might hope to be able to deduce the veracity of his sensory impressions and thus get to know about the external world. But such a project, whether in a foundationalist or coherentist guise, faces severe obstacles; it is susceptible to scepticism and threatens to collapse back to its solipsistic roots. The naturalistic epistemologist, on the other hand, does not start from the subjective experiences and reasoning capacities of the isolated subject, but sees the subject as an organism in an environment, an organism whose cognitive capacities have evolved in certain environmental conditions (which may include social conditions) and who uses those capacities to navigate and interact with that environment.

For the naturalized epistemologist it is no surprise that our cognitive capacities have their limitations. They are evolved to cope with environments of a certain kind. Consequently it cannot be regarded as a failure of those capacities if they would be unable to distinguish some common state of affairs from some hypothetical state of

affairs that would obtain only in some reasonably distant possible world. Thus for our perceptual systems to be operating properly it is *not* required that they should be able to tell (in a normal circumstance) that the subject is not a brain in a vat or that she is not being deceived by an evil demon. If we further associate knowledge with what is produced by a properly functioning cognitive system in a propitious environment, then we have a rather more liberal account of when the subject knows than is provided by the Cartesian epistemologist.

Naturalized epistemology provides a role for science in epistemology. Under the Cartesian approach, it is rational (internal) reflection that assesses the ability of some method to produce knowledge—whether the reliability of the method is immune from rational doubt. The naturalized epistemologist, on the other hand, regards the knowledge-generating capacity of a belief-forming process to be an *a posteriori* matter whose assessment requires scientific investigation. Thus naturalized epistemology can adopt a methodological pluralism that exemplifies the particularist tendency referred to above. There is no single scientific method—and certainly not an *a priori* one. Rather, as science develops it refines and adds to its methods. For example, early clinical trials involved no form of blinding or placebo control. Techniques such as double blinding, randomization, and placebo controls were developed to improve methods of assessing therapies as it became clear that outcomes could be biased by the both patients' and doctors' knowledge that of who is being treated and who not.

Naturalized epistemology gives a context for and support to externalism. It also thereby provides some degree of support for scientific realism. Externalist epistemology is less susceptible to sceptical worries than internalism. The externalist will regard it as potentially sufficient for being knowledge that a subject's belief stands in an appropriate relation to the world, for example by being produced by a reliable belief-forming method, whereas the internalist will require in addition that this connection itself be known to the subject. The empiricist displays this kind of internalism and is reticent about making inferences from sense-experiences to the nature of an independent physical world. Consequently the empiricist inclines either towards epistemological anti-realism—our theories do not tell us about the way the world is, they aim only to provide accurate predictions of future sense-experiences (Duhem 1914, c.f. van Fraassen 1980). Or they incline towards metaphysical anti-realism—our theories do tell us about the world, but only because the 'world' is constituted by our sense experiences (Carnap 1961). Externalism, bolstered by naturalized epistemology, says that the reliability of a microscope, MRI scanner, radio telescope, and so forth can suffice to give us knowledge of independent entities we cannot experience directly.

On the other hand naturalized epistemology does present a *prima facie* puzzle. How is it that the cognitive capacities of *Homo sapiens*, which evolved to give us knowledge of our immediate physical and social environment, are also able to give us knowledge, as the realist claims, of facts of a very different kind, including mathematically sophisticated laws of nature, the age and structure of the universe, and the constitution of atoms? One (part of) a possible answer is the idea that our cognitive capacities continue to evolve: our innate capacities are supplemented by those that develop thanks to advances in science and in the social organization of science. Let us consider Inference to the Best Explanation. As noted, IBE requires externalism—it is not *a priori* that the best explanations we devise are likely to be true. How exactly does externalism help IBE? Let us (the epistemologists) assume that the world is indeed one where the best explanations are likely to be true. Then a subject who ar-

gues in accordance with IBE will be using a reliable method and will thereby acquire justified beliefs. While that justification may require the subject to have reasons for thinking that her inferred hypothesis is a better explanation than its competitors, it does not require than she have reasons for thinking that IBE is a reliable method or that the world is one where the best explanations are likely to be true. Do we, as epistemologists have such reasons? An externalist epistemologist might regard herself within her rights to decline to answer this question. This minimal response says that the importance of the externalist perspective is its response to scepticism: it can show that justification and knowledge are possible. Whether they are actual depends on the way the world actually is, and such contingent questions are not within the realm of philosophy.

The naturalized epistemologist, on the other hand, might feel emboldened to go further, regarding it as legitimate to use empirical results in answering questions originating in epistemology. Showing that IBE is reliable is a matter of showing that the world is one where the best explanations are often true. And this looks like a tall order. After all, that would seem to make the world rather special, and why should that be the case? A possible strategy is to argue as follows. There can be different standards of explanatory goodness. What is needed is that the world is such that explanations that are the best by *our* standards are likely to be true. Furthermore, it might be the case that standards are discipline-relative. So in order that IBE as used in, say, palaeontology to be reliable, it needs to be the case that the world is one such that the explanations that are the best by the standards used by palaeontologists are likely to be true. Now imagine that the standards in question are not fixed for all time, but evolve in response to the developments of science. In particular it might be that the explanatory standards exemplified by a particularly successful hypothesis are more likely to be adopted than those exemplified by an unsuccessful theory. If so, it would be the case that the explanatory standards employed by a field of science are ones that are correlated with the truth and were adopted for that reason. Thus we would have an explanation of why IBE as employed in advanced sciences is reliable. So far, what has been presented is a hypothesis that stands in need of evidence, evidence that can only be gathered empirically. Arguably that evidence is exactly what is supplied by Kuhn's historically informed account of normal science, according to which the evaluative standards, including explanatory standards, employed by scientists working in a particular field are embodied in an especially successful theory (or practice) that stands as a model or exemplar for subsequent science in that field (Kuhn 1970b). Kuhn's history can be supplemented by the work of psychologists and cognitive scientists, who show that scientific thinking is heavily informed by thinking in terms of analogy with understood scientific problems, models, exemplars, cases, and so forth (Dunbar 1996). (For this approach to work, it has to be argued that there is significant continuity of explanatory standards across scientific revolutions—as Kuhn allows.) This examples illustrates how the naturalized epistemologist may be able to use the results of empirical investigations—in this case from history and psychology—to argue for an epistemological proposition—in this case concerning the reliability of IBE in advanced sciences.

5 Knowledge and evidence

To conclude I shall address a neglected topic in the epistemology of science, but one which in my view (Bird 2010) offers considerable scope for philosophy of science to

benefit from general epistemology. Despite the fact that we talk about ‘the growth of scientific knowledge’, very little epistemology in the philosophy of science is concerned with *knowledge*. There are a number of related reasons for this. Bayesianism simply eschews the notion of knowledge altogether—its sole epistemological notion is *credence*. (Bayesians do use a notion of evidence, but the latter is typically treated as matter of having credence equal to unity.) Even non-Bayesian accounts of confirmation tend to focus on incremental confirmation—which is when a piece of evidence adds to the likely truth of a hypothesis. This sheds little light on the question, when does one have enough evidence to know that the hypothesis is true. Although motivated by scepticism, Popper’s fallibilism has proved a popular view, especially among scientists. According to this view, our best theories are merely tentative, always open to the possibility refutation. This is difficult to reconcile with the claim that we know some scientific theories to be true. A moderately sceptical attitude is common even among realist philosophers of science. While they reject outright scepticism, they are sensitive to arguments such as the pessimistic meta-induction (Poincaré 1943). The latter argues that past theories that have been replaced by current theories were very often successful even though ultimately refuted. This shows that success is no good indicator of truth, and that we ought not hold that our current theories are different in this respect from their predecessors. A common realist response is to accept that our theories are indeed always strictly false while maintaining that they are nonetheless gaining in verisimilitude over time—later theories are getting closer to the truth, even if they do not reach it (Psillos 1999). But being strictly false, there is no possibility of knowing these theories to be the truth.

Science is sometimes held up as humankind’s most remarkable epistemic achievement. If so, it would be odd if it failed to reach the standard set by a very common epistemic standard, knowledge. Furthermore, according to one important epistemological approach it would be impossible for science to reach any significant level of achievement without there being knowledge in science. Timothy Williamson’s ‘knowledge first’ epistemology emphasizes the centrality of knowledge (Williamson 2000). For example, Williamson argues that all and only what we know is evidence. So the denial of knowledge in science would be tantamount to the denial of evidence in science. In the absence of evidence, little or nothing in science would even be justified. Williamson also argues that many of the problems facing Bayesian conditionalization are ameliorated if what is conditionalized upon on knowledge. The equation between evidence and knowledge also puts pressure on empiricism, which assumes that all our evidence is perceptual. This equation does not limit evidence to any particular kind of proposition. It thus supports scientists’ practice of citing heavily theoretically laden propositions as evidence (such as the anomalous precession of the perihelion of Mercury as evidence for General Relativity). The focus on knowledge also allows for a natural account of progress in science as the accumulation of scientific knowledge (Bird 2007).

6 Conclusion

Williamson’s knowledge first approach was developed with general epistemology in mind. The fact that it has many important implications for the philosophy of science shows that the relationship between general epistemology and the epistemological philosophy of science is still important, at least to the latter. Particularists will argue that such general approaches, while significant, can only provide a framework

and partial answers to question in scientific epistemology. There are still epistemological questions that are specific to science and which require specific answers: for example, how is it that the explanatory power of a theory is indicative of its truth, or can science be rational if it is guided by Kuhnian paradigms. Answers to such questions might be provided in part by empirical naturalistic investigation, drawing on history of science and cognitive psychology. The particularist–generalist tension is significant one for philosophy of science. On the one hand, philosophy of science must be true to science itself; it cannot deal only with an idealization that bears only a superficial resemblance to anything with which real scientists are engaged. On the other hand, the parts of science do not exist in epistemic isolation, nor does science as a whole. Even if the particularists are right, and that the details of the particular patterns of inference found in science are not derivable from a single epistemological rule, it remains a significant task for epistemology in science to show how these details cohere with a general approach to epistemology. It is notable that few of the recent developments in general epistemology—virtue epistemology, anti-luck epistemology, contextualism, and knowledge first, for example, have had much impact in philosophy of science. The preceding section gave a brief indication of the potential significance of the last of these. An important task for epistemologists and philosophers of science is to generalise this approach and to see how other developments in epistemology relate to science, on the one hand to test the fruitfulness of such approaches in the epistemically crucial case of science, while on the other hand seeing whether insights in general epistemology can help us advance our understanding of epistemological questions arising from the study of scientific discovery and inference.

References

- Bird, A. 2007. What is scientific progress? *Noûs* 41: 64–89.
- Bird, A. 2010. The epistemology of science—a bird’s-eye view. *Synthese*. DOI: 10.1007/s11229-010-9740-4.
- Bovens, L. and S. Hartmann 2003. *Bayesian Epistemology*. Oxford: Clarendon Press.
- Carnap, R. 1950. *Logical Foundations of Probability*. Chicago, IL: University of Chicago Press.
- Carnap, R. 1961. *Der logische Aufbau der Welt* (2nd ed.). Hamburg: Meiner. First edition 1928.
- Duhem, P. 1914. *La Theorie physique, son object et sa structure* (2nd enlarged ed. ed.). Riviere. (Reissued, Vrin, 1981) Translated P. Wiener *The Aim and Structure of Physical Theory* 1954 Princeton NJ: Princeton University Press.
- Dunbar, K. 1996. How scientists really reason. In R. Sternberg and J. Davidson (Eds.), *The Nature of Insight*, pp. 365–95. Cambridge, MA: MIT Press.
- Goldman, A. 1979. What is justified belief? In G. Pappas (Ed.), *Justification and Knowledge*, pp. 1–23. Dordrecht: Reidel.
- Hanson, N. 1958. *Patterns of Discovery: An Inquiry into the Conceptual Foundations of Sciences*. Cambridge: Cambridge University Press.

- Hardwig, J. 1991. The role of trust in knowledge. *Journal of Philosophy* **88**: 693–708.
- Harman, G. H. 1965. The inference to the best explanation. *Philosophical Review* **74**: 88–95.
- Howson, C. and P. Urbach 2006. *Scientific Reasoning: The Bayesian Approach* (3rd ed.). Chicago: Open Court.
- Kornblith, H. 2002. *Knowledge and its Place in Nature*. New York: Oxford University Press.
- Kuhn, T. 1970a. Reflections on my critics. In I. Lakatos and A. Musgrave (Eds.), *Criticism and the Growth of Knowledge*, pp. 231–78. Cambridge: Cambridge University Press.
- Kuhn, T. S. 1970b. *The Structure of Scientific Revolutions* (2nd ed.). Chicago, IL: University of Chicago Press.
- Lakatos, I. 1970. Falsification and the methodology of scientific research programmes. In I. Lakatos and A. Musgrave (Eds.), *Criticism and the Growth of Knowledge*, pp. 91–195. Cambridge: Cambridge University Press.
- Lakatos, I. 1971. History of science and its rational reconstructions. In R. C. Buck and R. S. Cohen (Eds.), *PSA 1970. Boston Studies in the Philosophy of Science VIII*, pp. 91–108. Dordrecht: D. Reidel.
- Lipton, P. 2004. *Inference to the Best Explanation* (2nd ed.). London: Routledge.
- Maxwell, G. 1962. The ontological status of theoretical entities. In H. Feigl and G. Maxwell (Eds.), *Minnesota Studies in the Philosophy of Science, vol. III; Scientific Explanation, Space, and Time*. Minneapolis, MN: University of Minnesota Press.
- Mill, J. S. 1843. *A System of Logic, Ratiocinative and Inductive*. References to the 1974 edition, J. M. Robson ed.; Toronto: University of Toronto Press.
- Poincaré, H. 1943. *La Science et l'Hypothèse*. Paris: Edition Flammarion. First published 1902; English language edition Dover, New York 1952.
- Popper, K. 1959. *The Logic of Scientific Discovery*. London: Hutchinson.
- Popper, K. 1970. Normal science and its dangers. In I. Lakatos and A. Musgrave (Eds.), *Criticism and the Growth of Knowledge*, pp. 51–8. Cambridge: Cambridge University Press.
- Psillos, S. 1999. *Scientific Realism: How Science Tracks Truth*. London: Routledge.
- Quine, W. V. 1969. Epistemology naturalized. In *Ontological Relativity and Other Essays*, pp. 69–90. New York, NY: Columbia University Press.
- Reichenbach, H. 1971. *The Theory of Probability*. Berkeley, CA: University of California Press. A translation by Ernest R. Hutton and Maria Reichenbach of *Wahrscheinlichkeitslehre. Eine Untersuchung über die logischen und mathematischen Grundlagen der Wahrscheinlichkeitsrechnung*, Leiden, 1935; revised by the author.

- Schlick, M. 1985. *General Theory of Knowledge*. La Salle, Ill.: Open Court Publishing. Originally published as *Allgemeine Erkenntnislehre* (1925); this edition trans. A. Blumberg.
- Teller, P. 1976. Conditionalization, observation, and change of preference. In W. Harper and C. A. Hooker (Eds.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*. Dordrecht: D. Reidel;.
- van Fraassen, B. 1980. *The Scientific Image*. Oxford: Oxford University Press.
- Whewell, W. 1837. *History of the Inductive Sciences*. London: Parker.
- Whewell, W. 1840. *The Philosophy of the Inductive Sciences, Founded upon Their History*. Cambridge: J. and J. J. Deighton.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.