THE EPISTEMOLOGY OF SCIENCE—A BIRD'S-EYE VIEW

Abstract

In this paper I outline my conception of the epistemology of science, by reference to my published papers, showing how the ideas presented there fit together. In particular I discuss the aim of science, scientific progress, the nature of scientific evidence, the failings of empiricism, inference to the best (or only) explanation, and Kuhnian psychology of discovery. Throughout, I emphasize the significance of the concept of scientific *knowledge*.

1 The aim of science

The aim of science is the generation of scientific knowledge. Science issues in propositional outputs that we seek to support with sufficient evidence that they are worthy of belief. So a principal product of science is scientific belief, and since the aim of belief in general is knowledge, the aim of scientific belief is scientific knowledge.

Let me clarify and defend these remarks. First, what is it for belief to have an aim? Mental capacities, such as the capacity for belief, have functions in much the same way as anatomical organs have functions (for functional views of epistemic capacities, see Kornblith 2002; Neta 2007; Bird 2009b). Arguably such functions may be explained in terms of their contribution to fitness. The aim of belief is to be identified with the function of the capacity to believe, which is, in my view, to generate knowledge and to make that knowledge available to other cognitive functions (Bird 2007b; see Williams 1973, Velleman 2000, and Wedgwood 2002 for alternative views of the aim of belief and Owens 2003 for the view that belief has no aim). Secondly, why concentrate on knowledge rather than truth? One reason is that having truth as an aim is in fact too easy. Take any pair of a proposition and its negation. One of these is true. So a subject who, for as many of such pairs as possible, believes one of the two propositions, chosen at random, will end up believing a very large number of true propositions, more than any scientist. Arguably it is possible to believe contradictory propositions, and so if a subject can arrange for that to occur for as many such pairs as possible, then he will be guaranteed to have many true beliefs. But such random or contradictory believing scarcely satisfies the requirements of science. Furthermore, the psychological difficulty of believing contradictory propositions or even propositions that are only weakly supported by the evidence, suggests that the function (and so aim) of belief is not satisfied by mere true belief. Some who wish nonetheless to continue to focus on true belief, will have to supplement the truth-seeking aim of belief with a falsity-avoidance aim (Papineau 1993). So belief would have a somewhat schizophrenic nature. If on the other hand, the aim of belief is knowledge, than random or contradictory believing will not achieve its aim, even if the resulting belief is true. A belief will achieve the aim of knowledge only if it produced in some reliable, truth-implying manner, for example, by being a judgment that is produced as a result of careful reasoning from evidence.

If the aim of science is knowledge, then scientific progress is, typically at least, the accumulation of scientific knowledge (Bird 2007d), on the grounds that if the aim of some activity is the production of X, that activity makes progress to the extent that it produces X. We may wish to add that progress might also be made when something is done that promotes the achievement of X, even if that is not doing X itself, as one might make progress in a building project by preparing the foundations. In science that preparatory work might include the production of some new piece of equipment, or the design and improvement of some experimental procedure. It might include the use of these things to produce evidence. It should be remembered though, that there will be knowledge at the core of these activities, such as knowledge how to measure the charge on an oil-drop.[1] As we shall see, the production of new evidence is itself an addition to scientific knowledge, precisely because evidence is knowledge.

The view that scientific progress is the accumulation of scientific knowledge is one that is often regarded as philosophically naive. Few philosophers have endorsed this view in recent times. Popper, for example, rejected such a view because he denied the possibility of positive scientific knowledge, i.e. knowledge of the truth of general hypotheses.[2] From a rather different perspective, others have rejected the accumulating knowledge view of scientific progress because they do not hold that an objective conception of truth is relevant to science. Kuhn (1970) and Laudan (1977) take progress to be a matter of the accumulation of *solutions to problems*—where what counts as a solution to the problem is determined *internally*, that is, by the standards of the prevailing science. The problem with this, of course, is that progress and the appearance of progress are not the same thing. In the two years following René Blondlot's alleged discovery of N-rays, many scientists reported finding these rays emitted from a variety of substances. Their findings were not a contribution to scientific progress, however much it may have seemed that way to leading scientists in France (and elsewhere), including Blondlot himself; that is simply because there are no N-rays.[3]

Nor does the accumulation of true scientific beliefs suffice for scientific progress, if those beliefs are only accidentally true. But 'accidentally true' I mean, for example, a claim that is true but which is believed on the basis of reasoning that is flawed.

[1] Knowledge *how* is still knowledge. We could add, but do not need to, that this may be construed as further propositional knowledge: Stanley and Williamson 2001.

[2] One could reconcile Popperian falsificationism with the accumulation view of scientific knowledge, by thinking of progress as accumulating knowledge that theories are *false*. But strictly speaking, even that view is not available to Popper. That is because he accepts the theory-ladenness of experimental observation. Hypotheses are not tested and falsified by having theory-free perceptions of the world. Rather tests are carried out using instruments, procedures, and patterns of inference that are all informed by auxiliary hypotheses that have to be taken as given to produce a result that can be used to falsify a hypothesis. For example, the hypothesis that the Earth is only 6,000 years old is refuted by evidence of the age of rocks. But evidence of the age of rocks depends on the employment of hypotheses concerning radioactive decay. So the successful refutation of the young-Earth hypothesis depends on the truth of a theory in physics. But, according to Popper's own epistemology, if one cannot depend on the truth of the auxiliary hypothesis (the theory of radioactive decay) one cannot trust the resulting datum (the age of the rocks) used to refute the hypothesis (that the Earth is young). And according to Popper one cannot trust the truth of a general hypothesis. So one cannot know that one has indeed refuted the hypothesis after all. So, strictly, on Popper's epistemology one cannot even have progressively accumulating negative knowledge.

[3] See Rowbottom (2008) and Bird (2008) for further discussion.

Examples of this in actual science are rare. If a scientist uses flawed reasoning, the outcome is likely to be false belief. Nonetheless, for sake of argument, imagine that we now discover rays that possess some of the key theoretical features ascribed by Blondlot to his N-rays; but we also know that his techniques could not possibly have produced those rays. So his theoretical beliefs were not supported by his experiments. Those beliefs were true, but only accidentally so. We would not regard the accidentally true theoretical claims as contributing to progress, since they do not have the right evidential connection to the truth.

Many philosophers, from Popper (1963) to Psillos (1999), have tended to see progress in terms of increased verisimilitude rather than the accumulation of truth. That difference is immaterial to the issue I am discussing. First, progress as increasing verisimilitude suffers from the same objections. An unjustified, accidental increase in verisimilitude is not progress, for the reasons given in the preceding paragraph. Secondly, an increase in verisimilitude is always accompanied by in an increase in (full, exact) truth for related propositions. If '*p*' is approximately true, then 'approximately, *p*' is exactly true. 'H_0=74.2 (km/s)/Mpc' is nearly true; but 'H_0=74.2 $\pm$ 3.6(km/s)/Mpc' can be exactly true. Consequently, it is no objection to the accumulating knowledge account that some important propositions are not exactly true and so are not knowable, for a relevant related proposition will be exactly true and knowable.

So neither accumulating truth nor increasing verisimilitude can suffice for scientific progress, for they may lack justification. My conclusion is that only knowledge, which does have the right justifying connection to the truth, typically through reliable reasoning and good evidence, suffices for scientific progress.

2 Evidence

The view that science aims at knowledge gives us a handle on the notion of evidence. Crudely speaking, a new scientific claim will be arrived at by a process of reasoning from propositions that are presented as evidence. Not any old proposition constitutes evidence. Imagine that someone seeks to persuade you of some claim using impeccable reasoning from a set of propositions claimed as evidence. It then turns out that he had no reason to believe those propositions, but chose them simply for their being persuasively useful in this argument. We would not regard those propositions as genuine evidence. So which propositions are those that are genuine evidence?

Let us assume, as argued for in Section 1, that the aim of belief, and scientific belief in particular, is knowledge. We can think of evidence in functional terms: the function of evidence is to provide the propositional input into a process of reasoning that can achieve this aim. More precisely, then:

> (E) The proposition *p* is in S's evidence if and only S can gain knowledge by (non-redundant) inference from *p*.

Whether the aim of knowledge is achieved will depend on the quality both of the propositional input and of the reasoning process. We will come to the reasoning process later. For the purposes of (E) we may assume an impeccable reasoning process; what constraints then does (E) put upon evidence? We can see that anything less than knowledge will not do as evidence. If conclusion *c* depends inferentially on evidence *e*, then if *e* is not known then *c* will not be known either. At the same time, any piece of knowledge *k* can be evidence, since there will always conclusions drawn from *k* that can in principle thereby come to be known. So (E) supports Williamson's (2000) claim that a subject's evidence is precisely what that subject knows, (E=K).

Thus the distinction between 'evidence' and 'theory' is not an absolute one but is contextual. In the context of the debate between Millikan and Ehrenhaft over whether there is a single basic unit of electrical charge, the unique charge on an electron, Millikan's experimental data are the *evidence* that helps establish the truth of the *theory* that electrons all have the same charge, with a value of 4.8×10^{-10} esu. In the context of assessing evidence relevant to the standard model of particle physics, these latter propositions are considered as *evidence* not as hypotheses.[4]

3 Empiricism

The consequences of (E) and (E=K) for philosophy of science are profound but have not been explored in depth. They do not constrain our evidence to any particular kind or source of knowledge (or belief). In particular they do not identify evidence with our observational, let alone perceptual knowledge. They are, at very least, in *prima facie* tension with the epistemological arm of empiricism. That in itself does not refute emprcism, but it can undermine arguments for empiricism. For example, consider the following argument:

> (P1) All our evidence concerns the observable.
> (P2) From evidence that concerns only the observable, only conclusions concerning the observable follow.
> *therefore*
> (C) Only propositions concerning the observable can be known by inference from our evidence.

If that argument were sound it would limit scientific knowledge (which consists of propositions inferred from the evidence) to the observable, as positivists and empiricists like to claim. Note, however, that in the light of (E=K), (P1) must be read as asserting that all our knowledge concerns the observable. But that is precisely the conclusion, (C). So this argument for empiricism assumes what it sets out to prove; given (E=K), one could only be motivated to assert (P1) if one already accepts the conclusion that all knowledge is observational.[5] In the light of (E=K), arguments for empiricism of the kind given above, that depend on a premise to the effect that our evidence is limited in kind will be question-begging.

Empiricists will therefore want to reject (E=K). The challenge to the empiricists is to provide a defensible and principled account of what evidence is, that is more restrictive than (E=K). The empiricist might propose that function of evidence is to provide the *basic*, i.e. non-inferred propositions from which knowledge can be inferred. Thus evidence would be is non-inferential knowledge, (E=NIK)—compare Maher's (1996) view that evidence is what is known directly by experience. The immediate problem with (E=NIK) is that it faces problems with forgotten (non-inferential) evidence. We typically do not remember the perceptual knowledge from which our inferred knowledge comes. So that knowledge would seem to be without evidence, and so also would any supposed knowledge inferred from it. Similar

[4] Furthermore, when we *present* a proposition as evidence we are thereby claiming not only that the proposition is true, but also that we know it, consequently is must meet an additional epistemic standard.
[5] See Bird 2007c for further details.

problems may arise because our evidence shifts for rational reasons—evidence may be undermined but replaced by fresh evidence, but without the derived knowledge being inferred anew from the updated evidence set. Again, the derived knowledge would then be inferred from no current evidence. [6]

The 'unlimited' conception evidence, according to which all knowable propositions could be evidence, not just some observational or perceptual subset, fits with our normal usage. Part of the evidence for Einstein's general theory of relativity is the precession of the perihelion of Mercury. But that precession is only observable in a very extended sense of 'observable'. The precession is inferred from many individual observations of Mercury and the inference involves considerable mathematical work as well as substantive auxiliary hypotheses. As exemplified above, the fact that evidence is not restricted to any particular kind of knowledge undermines certain empiricist images of science. In particular I reject a view of science in which perceptual experience necessarily plays a central part. Both propositions concerning the perceptible world and those concerning deeply unobservable or general facts can be known, and being knowable all these kinds of proposition are potentially evidence in further inferences. As suggested earlier, scientific progress involves the accumulation of knowledge, and all this knowledge may be drawn upon in generating new knowledge. This is not to deny that what we call observations are, as a matter of contingent fact, very important in science. Often the facts we need to test a theory cannot be inferred from what we already know, and so we must make observations. But it should equally be noted that such observations may have precious little to do with perception. The observational evidence used in constructing weather models is now frequently collected and processed automatically. Modern radio telescopy delivers nothing of significance that is perceptible, but rather the observations in the form of data collected in a computer, which, furthermore, has been processed by statistical software. According to this view observation has nothing to do with perception. Rather, what makes something an item of observational knowledge is the reliability and uncontentious nature of the mechanism which produces it. As such, what counts as an observation depends on the current state of science.

I do not regard the arguments sketched above as knock-down refutations of empiricism; but I do think that they provide a significant line of attack for the realist philosopher of science. While there may be an epistemologically significant observable–unobservable distinction, it does not align with the perceptible–imperceptible distinction. Nor does it support any kind of scepticism with respect to the unobservable. Our evidence is not limited to the observable, even when understood in this way, for both theoretical knowledge and observational knowledge are used together as the premises in knowledge-extending arguments. Thinking more broadly about realism, the claim that progress in science is the accumulation of scientific knowledge can be used as the basis of an account of what scientific realism is: the scientific realist claims that science really does progress—it does add to our stock of scientific knowledge. More specific realist versus anti-realist debates can focus on the kinds of knowledge that is added. Thus a constructive empiricist could might be a partial realist—science adds knowledge of the empirical adequacy of theories while being a partial anti-realist—science does not add knowledge of 'unobservable' (*viz.* imperceptible) entities.

[6] See Bird 2004 for further details.

4 Social knowing

Evidence that is collected, processed, and stored automatically might seem to be a counterexample to the claim that all evidence is knowledge, since this evidence does not seem to be known by anyone. To respond to this, we need to distinguish between, on the one hand, what is known by and is evidence for some individual, and, on the other hand, what is known by and is evidence for scientists in a field, regarded as a collective entity. I think that there is a social sense of the term 'knowledge' as used, for example, in the phrase 'the growth of knowledge'. This sense is not reducible to what is known by any individual person. Indeed I argue that what is socially known is not even supervenient on the mental states of individuals. Social knowledge can be stored, for example, not only in peoples' mind, but also in libraries and electronic databases; it can be produced not only by scientists but by automated systems, as just mentioned.

My preferred conception of social knowing is a functional analogue to individual knowing (Bird 2009b). As I mentioned above, the function of individual capacities for knowledge is to generate or store inputs to action or processes of reasoning. Likewise for social knowing, where the actions are socially coordinated actions (such as building the LHC or, one might hope, a plan for combating climate change) and the processes of reasoning are social ones (e.g. the social process whereby new hypotheses are tested, published, debated, and accepted). What goes to make the social does not necessarily have to be human: what is social is not just a collection of individuals, nor is it a collection of individuals who happen to be coordinated in a certain way. The material means of coordination are also part of the social. And these can play a role in the generation and storage of social knowledge.

5 Inference to the Only Explanation

Above I factored, crudely perhaps, the process of reaching a conclusion by inference into an evidential component and an inferential component. We have discussed the former. How about the latter?

There are many ways of making inferences. New theoretical knowledge permits new knowledge-generating inferences. For example the Born rule and Hund's rules in physics should be regarded as subject-specific inference rules. Certain inferences may be drawn using various techniques and apparatus which may be designed using theoretical knowledge; these might vary from the simple, such as inferring that a liquid is acid from the fact that it turns litmus paper red, to the more sophisticated, such as inferring the presence of tuberculosis from a chest X-ray. There is no single scientific method, but a range of scientific methods, which grows as our knowledge grows (Bird 1998). New knowledge begets new techniques for generating knowledge and so on. This is why (and the sense in which) epistemology, the epistemology of science in particular, much be naturalized (see Kitcher 1993 for a naturalized epistemology of science that eschews the single Scientific Method of 'Legend').

That said, I think that there is a form of inference that is particularly important in the production of new theoretical knowledge. This is *inference to the best explanation* (IBE). The basic idea behind IBE is that if a putative hypothesis would explain some evidence, then that evidence provides some degree of confirmation to that hypothesis. Thus the fact that Einstein's general theory of relativity could explain the anomalous precession of the perihelion of Mercury was a reason to think that

Einstein's theory is correct. In some cases, several competing hypotheses each provide possible explanations of the evidence. IBE tells us that, subject to various constraints, the evidence most strongly confirms that hypothesis which *best* explains the evidence.[7] Lipton (2004) makes a useful distinction between a potential explanation and the actual explanation. A potential explanation is a hypothesis about what could possibly explain certain pertinent evidence (the explanandum); the actual explanation is the hypothesis that *correctly* explains the explanandum. Lipton takes IBE to be a two stage process. In the first stage scientists develop various potential explanations of the explanandum. In the second stage the hypotheses considered in stage one are compared with respect to how well they explain this evidence. We can understand IBE as saying that the best potential explanation of the evidence has the high likeliness of being the actual explanation, and hence of being true or close to the truth. Semmelweis was able to show that infectious material carried on the hands of medical students from dissections was the cause of puerperal fever in post-partum women whom they later examined. Lipton argues that Semmelweis's inference was that this hypothesis was a better explanation of the puerperal fever than its competitor hypotheses (birth position, presence of a priest, overcrowding, rough examinations).

Understood thus, IBE is subject to three objections. The first points to stage one. At this stage the scientists think up explanations for consideration at stage two. In so doing there will be a multitude of potential explanations that exist in the abstract but which the scientists never even think of and so never evaluate for explanatory goodness at stage two. If the actual explanation is among the multitude that is not considered, then the ranking at stage two will be to no avail. This is the problem of underconsideration. At stage two, a pair of problems arise, which Lipton calls 'Hungerford's objection' and 'Voltaire's objection'. According to the former, 'goodness' is too subjective a notion to have any correlation with the truth. According to the latter, even if there is an objective ranking according to explanatory goodness, there is no reason to suppose that the best explanation is the actual explanation: to assume so is to assume that ours is the best of all possible worlds, by explanatory standards, but that is not a safe assumption.

In my view, there is a version of IBE that avoids the two problems at stage two: Inference to the *only* explanation (Bird 2005a, 2007a). Note that in a realistic scientific case, evidence will be sought that will differentiate between the competing hypotheses. That is compatible with Lipton's IBE—such evidence may indicate that one explanation is better than a competitor. However, such evidence is often evidence that will *refute* one of the competing explanations; I argue that Semmelweis was able to show that the competing hypothesis were not simply poor explanations of the evidence—rather they *could not* be explanations of it (Bird 2009a). If such evidence were to refute all but one of the competing potential explanations, the remaining one can be inferred to be the actual explanation; such is the import of the dictum put into the mouth of Sherlock Holmes by Sir Arthur Conan Doyle (1953: 94,118): "Eliminate the impossible, and whatever remains, however improbable, must be the truth."[8]

The format of a Holmesian train of reasoning can be captured thus:

[7] Something like this converse relationship between confirmation and explanation is also a consequence of combining the deductive-nomological model of explanation with the hypothetico-deductive model of confirmation.

[8] Holmes's eliminativism is referred to also by Kitcher (1993: 239) and by Earman (1992).

(i) S knows that one of the hypotheses $h_1, h_2, \ldots, h_n$ is true;
(ii) S gathers evidence inconsistent with hypotheses $h_1, h_2, \ldots, h_{n-1}$;
(iii) S thereby comes to know that hypotheses $h_1, h_2, \ldots, h_{n-1}$ are false and that hypothesis h_n is true.

Note that in order for this to hold, evidence must meet the high Williamsonian standard, that of constituting knowledge, as I discussed earlier. For if evidence is less than knowledge, then the fact that h_1 is inconsistent with the evidence cannot suffice for knowing that h_1 is false, and so on. And a fortiori one cannot know that the remaining hypothesis, h_n is true.

While inference to the only explanation constitutes the best case of inference to the best explanation, I accept also that we cannot always achieve this optimal inferential position, in which case we must fall back on comparative IBE of the kind Lipton describes. While inference to the only explanation can, it should be clear, give us knowledge, it is an open question whether comparative IBE can do so also? I am sympathetic to the mildly sceptical concern that if some plausible live explanatory option has not been refuted, how can we know that it is false? Nonetheless, IBE might give us justification for preferring one potential explanation over another, if not permitting us to believe the preferred option outright.

6 Kuhnian psychology of discovery

Inference to the only explanation still leaves the problem of underconsideration unanswered. And if comparative IBE is to guide rational preference, then there ought to be some degree of correlation between explanatory goodness and truth, in which case Hungerford's objection and Voltaire's objection are still in play.

As it happens, I believe that we can sometimes be in a position where the problem of underconsideration does not arise. In some cases of inference to the best explanation we can know both that we refuted all considered hypotheses bar one and also that we have refuted all unconsidered hypotheses. Controlled experiments are designed precisely so that the hypotheses in play are reduced to two, one of which is the null hypothesis.

Not all experimental or observational circumstances can be understood in such a way that in effect every possible hypothesis is considered directly or indirectly. But that does not matter if the possibilities that remain unconsidered are remote. We need not require that our evidence refutes all false hypotheses, but only that it refutes the false hypotheses that hold in nearby possible worlds.[9] So what we need is some mechanism for ensuring that we do at least consider hypotheses that do hold in all the nearby worlds, and in particular in the actual world.

I suggest that Kuhn's psychology of scientific research, and its core notion of an exemplar may provide part of the answer here (Kuhn 1970: 175–6, Bird 2005b). The general idea, which is backed by a considerable quantity of more recent psychological research (e.g. Dunbar 1996), is that scientists look for answers to scientific problems by considering potential solutions that are similar to existing, successful, exemplary problem solutions (which Kuhn initially called 'paradigms'). Now let us add somewhat to the Kuhnian framework by assuming that a certain set of exemplars are all true or approximately true. When we consider potential solutions to a new, but similar problem, and these potential solutions are correspondingly sim-

[9] This is the principle of safety; cf. Williamson (2000: 123–30), Pritchard (2003).

ilar to the problems solved by those exemplars, we are then considering potential solutions that hold in nearby possible worlds.

The fact that a true exemplar anchors our hypothesizing to the actual world, and so typically encourages consideration of nearby possible worlds, does not guarantee that the hypothesis that is true of the actual world is always considered, nor that every nearby possible world is considered. I conjecture that it is in fact frequently the case that we consider the relevant hypotheses that hold of all nearby possibly worlds; but clearly it is not *always* the case. That this is so is shown, in Kuhnian terms, by periods of crisis, when solutions to scientific problems that are based on existing exemplars fail to work. That in turn necessitates the development of new exemplars, which may require the rejection of existing ones. Such events are scientific revolutions. A revolution must cast doubt on claims to knowledge made for scientific theories and problem-solutions that are based on the now-rejected exemplars. In my view such revolutions, on the large-scale, happen relatively rarely. They include the Copernican revolution, and Lavoisier's rejection of phlogiston. They continue to happen on a smaller scale, for example the overturn of the theory that stress is a major factor in the aetiology of peptic ulcers in favour of the view that they are mostly caused by a bacterium, *Helicobater pylori*. Significant Kuhnian revolutions tend more to be the discovery of limitations to existing exemplars. This will have ramifications for existing problem solutions based on those exemplars, but the ramifications may be a matter of limiting the applications of those solutions rather than rejecting them altogether. The revolutions in modern physics in the first decades of the twentieth century are more of this kind. While Kuhn emphasized the problems of incommensurability, especially in his later writings, he also stressed (1970: 160–73) that there is continuity and progress even through scientific revolutions, which will retain the bulk of the problem-solving power of the exemplars in their replacements.

Kuhnian exemplars also help us with Hungerford's and Voltaire's objections. The idea here is that exemplars set standards of explanatory goodness, or what Lipton calls 'loveliness'. If so, both problems are answered.[10] For loveliness will not be purely subjective, but will be determined by objectively successful scientific achievements. Furthermore, on the assumption that such achievements have a large truth-content, the standards of loveliness thus set will have some degree of correlation with the truth. If the standards are not subjective, neither are they sempiternal, necessary, world-independent standards. If they were the latter, then their correlation with truth in the actual world would require the actual world to be very special. But the objective standards set by our exemplars are not like this; they are world-relative, being based on what has been shown to work in the actual world. So the actual world may be the best of all possible worlds, as regards explanatory loveliness, but only because the standards are this-world standards. The standards set by exemplars are both objective and relative. In this sense they are rather like the standards we apply to marking exams; the standards may change over time and from subject to subject, but they are not merely a matter of purely subjective opinion.

7 Conclusion

Science is in the business of generating knowledge, principally through inference from evidence. While science uses a variety of inference patters, some specific and

[10]See Walker (2009) for an extended discussion.

some general, one of the most general and important inferences schemas is inference to the best explanation. This produces knowledge at least when employed in the Holmesian form, inference to the only explanation. Inference to the best explanation requires scientists to be able to think up potential explanations of phenomena. For that reason, theoretical science cannot be algorithmic. And for the same reason there is a grey area in the distinction between the context of discovery and the context of justification. Given the importance of IBE, epistemology cannot ignore the processes whereby we generate hypotheses. I have suggested that Kuhn's insights concerning the function of exemplars, especially when combined with recent research about the nature of analogical thinking, can provide an understanding of these processes. For inference to the only explanation to generate knowledge, what counts as our evidence must itself be known. Only then can we know that a hypothesis inconsistent with our evidence is refuted. More generally indeed, only if the evidence itself is known, can inference lead to knowledge. If we regard that as the function of evidence—evidence is that which leads to knowledge by inference—then we have an account of the concept of evidence that also vindicates Williamson's equation of evidence and knowledge. The thought that what we want from a scientific inference is knowledge is itself supported by arguments for the view that what constitutes scientific progress is the accumulation of knowledge and more generally that the aim of belief is knowledge.[11]

References

Bird, A. 1998. *Philosophy of Science*. London: University College London Press.

Bird, A. 2004. Is evidence non-inferential? *The Philosophical Quarterly* **54**: 252–65.

Bird, A. 2005a. Abductive knowledge and Holmesian inference. In T. S. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology*, pp. 1–31. Oxford: Oxford University Press.

Bird, A. 2005b. Naturalizing Kuhn. *Proceedings of the Aristotelian Society* **105**: 109–27.

Bird, A. 2007a. Inference to the only explanation. *Philosophy and Phenomenological Research* **74**: 424–32.

Bird, A. 2007b. Justified judging. *Philosophy and Phenomenological Research* **74**: 81–110.

Bird, A. 2007c. Underdetermination and evidence. In B. Monton (Ed.), *Images of Empiricism*, pp. 62–82. Oxford: Oxford University Press.

Bird, A. 2007d. What is scientific progress? *Noûs* **41**: 64–89.

Bird, A. 2008. Scientific progress as accumulation of knowledge: a reply to Rowbottom. *Studies In History and Philosophy of Science Part A* **39**: 279–81.

Bird, A. 2009a. Eliminative abduction—examples from medicine. *Studies in History and Philosophy of Science (forthcoming)*.

[11] I am very grateful to Luciano Floridi and members of the research seminar "The Nature and Scope of Information" at the Department of Engineering Science, University of Oxford, for their very helpful comments on an early draft of this paper.

Bird, A. 2009b. Social knowing. *Unpublished article.*

Conan Doyle, A. 1953. The sign of four. In *The Complete Sherlock Holmes*, Volume I. New York: Doubleday.

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.

Earman, J. 1992. *Bayes or Bust?* Cambridge MA: Bradford.

Kitcher, P. 1993. *The Advancement of Science.* New York: Oxford University Press.

Kornblith, H. 2002. *Knowledge and its Place in Nature.* New York: Oxford University Press.

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

Laudan, L. 1977. *Progress and its Problems: Toward a theory of scientific growth.* Berkeley: University of California Press.

Lipton, P. 2004. *Inference to the Best Explanation* (2nd ed.). London: Routledge.

Maher, P. 1996. Subjective and objective confirmation. *Philosophy of Science* **65**: 149–74.

Neta, R. 2007. How to naturalize epistemology. In *New Waves in Epistemology*, pp. 324–53. Ashgate.

Owens, D. J. 2003. Does belief have an aim? *Philosophical Studies* **115**: 283–305.

Papineau, D. 1993. *Introduction to Philosophical Naturalism.* Oxford: Blackwell.

Popper, K. 1963. *Conjectures and Refutations: The Growth of Scientific Knowledge.* London: Routledge and Kegan Paul.

Pritchard, D. 2003. Virtue epistemology and epistemic luck. *Metaphilosophy* **34**: 106–30.

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

Rowbottom, D. 2008. N-rays and the semantic view of scientific progress. *Studies in History and Philosophy of Science* **39**: 277–8.

Stanley, J. and T. Williamson 2001. Knowing how. *Journal of Philosophy* **98**: 411–44.

Velleman, J. D. 2000. On the aim of belief. In *The Possibility of Practical Reason*, pp. 244–81. Oxford: Oxford University Press.

Walker, D. 2009. *A Kuhnian Defence of Inference to the Best Explanation.* Ph. D. thesis, University of Bristol.

Wedgwood, R. 2002. The aim of belief. *Philosophical Perspectives* **16**: 267–97.

Williams, B. 1973. Deciding to believe. In *Problems of the Self*, pp. 136–51. Cambridge: Cambridge University Press.

Williamson, T. 2000. *Knowledge and its Limits.* Oxford: Oxford University Press.