

IS THE BEST GOOD ENOUGH?

BY PETER LIPTON

I

Is it ever rational to believe that a scientific theory is even approximately true? The evidence, however extensive, will not entail the theory it supports: the grounds for belief always remain inductive. Consequently, the realist who holds that there can be rational grounds for belief remains hostage to wholesale Humean scepticism about induction. The Humean argument has yet to be conclusively turned, but that project is not my present concern. Instead, I propose to consider intermediate forms of scepticism which attempt to show that, even if we grant scientists considerable inductive powers, rational belief in theory remains impossible. I will argue that some of these intermediate forms of scepticism are unstable, leading either back to radical Humean doubt or towards a moderate realism.

I will focus especially on the argument from 'underconsideration'. This argument has two premises. The *ranking* premise states that the testing of theories yields only a comparative warrant. Scientists can rank the competing theories they have generated with respect to likelihood of truth. The premise grants that this process is known to be highly reliable, so that the more probable theory is always ranked ahead of a less probable competitor and the truth, if it is among the theories generated, is likely to be ranked first, but the warrant remains comparative. In short, testing enables scientists to say which of the competing theories they have generated is likeliest to be correct, but does not itself reveal how likely the likeliest theory is. The second premise of the argument, the *no-privilege* premise, states that scientists have no reason to suppose that the process by which they generate theories for testing makes it likely that a true theory will be among those generated. It always remains possible that the truth lies rather among those theories nobody has considered, and there is no way of judging how likely this is. The conclusion of the argument is that, while the best of the

generated theories may be true, scientists can never have good reason to believe this. They know which of the competing theories they have tested is likeliest to be true, but they have no way of judging the likelihood that any of those theories is true. On this view, to believe that the best available theory is true would be rather like believing that Jones will win the Olympics when all one knows is that he is the fastest miler in Britain.

The argument from underconsideration is clearly different from the radical Humean problem. The upshot of Hume's argument is that all non-deductive evaluation is unjustifiable. By contrast, the argument from underconsideration concedes very substantial inductive powers, by granting scientists the ability to rank reliably whichever competing theories they generate. Indeed these powers are almost certainly stronger than any sensible scientific realist would wish to claim. This only seems to strengthen the underconsideration argument, however, since it appears to show that even these generous powers can not warrant belief in any scientific theory.

The argument from underconsideration is much more similar to an argument from underdetermination. According to one version of this argument, scientists are never entitled to believe a theory true because, however much supporting evidence that theory enjoys, there must exist competing theories, generated or not, that would be as well supported by the same evidence. This is an argument from inductive ties. Like the argument from underconsideration, it is an intermediate form of scepticism, since it grants scientists considerable inductive powers, but the two arguments are not the same. The argument from underconsideration does not exploit the existence of inductive ties, though it may allow them. On the other side, the argument from underdetermination does not assume any limitations on the scientists' powers of theory generation. Roughly speaking, whereas the underdetermination argument depends on the claim that scientists' inductive powers are excessively coarse-grained, the underconsideration argument focusses instead on the claim that they are only comparative. Moreover, the argument from underdetermination is in one sense more extreme than the argument from underconsideration. The underdetermination problem would remain even if scientists knew all the possible competing hypotheses and all possible data, whereas the underconsideration problem would disappear if they only knew all the competitors. Nevertheless, the similarities between the two arguments are substantial. Towards the end of this essay I will suggest that some of the objections to

the argument from underconsideration also threaten the argument from underdetermination.

II

Bas van Fraassen has recently deployed the argument from underconsideration as part of his attack on Inference to the Best Explanation (van Fraassen 1989, pp.142-50). It may therefore be useful to clarify the connections between Inference to the Best Explanation, van Fraassen's constructive empiricism, and the argument from underconsideration, before turning to a critical assessment of the argument itself. (In the same work, van Fraassen also mounts a quite different dutch-book argument against Inference to the Best Explanation: although this argument seems to me flawed, I will not criticise it here.)

Constructive empiricism is the view that an aim of science is not truth across the board, but only empirical adequacy, the truth about all observable entities and processes. Inference to the Best Explanation is an account of inductive inference. Its governing idea is that explanatory considerations are a guide to inference. In its simplest form, the account claims that scientists judge that the theory which would, if correct, provide the best explanation of the available evidence is also the theory that is likeliest to be correct. What then is the relationship between constructive empiricism and Inference to the Best Explanation? They are widely supposed to be incompatible. Certainly champions of Inference to the Best Explanation tend to be realists and van Fraassen develops his case against Inference to the Best Explanation as part of his argument for constructive empiricism. But the two views are in fact compatible, since one may have a constructive empiricist version of Inference to the Best Explanation. To do this requires only that we construe 'correct' as empirically adequate rather than as true and that we allow that false theories may explain. I see no special barrier to the former, and van Fraassen's own account of explanation allows the latter.

Is Inference to the Best Explanation especially vulnerable to the argument from underconsideration, more vulnerable than other accounts of inference? Van Fraassen's discussion gives this impression, since he deploys the argument specifically against this account. Moreover, Inference to the Best Explanation does seem particularly vulnerable, since it seems that 'best theory' can only mean 'best of those theories that have been generated'. Here too, however, the appearances may be deceptive. The

governing idea of Inference to the Best Explanation, as I have said, is simply that explanatory considerations are a guide to inference, and this need not be articulated in a way that makes the evaluation comparative. That is, Inference to the Best Explanation might be more accurately if less memorably called 'Inference to the Best Explanation if the Best is Sufficiently Good'. The story one tells about the explanatory virtues may make them either comparative or absolute. In spite of my best efforts, Inference to the Best Explanation still remains at such an early stage of articulation that we cannot yet say with any confidence which version is the more promising (Lipton 1991).

Finally, what is the relationship between the argument from underconsideration and constructive empiricism? Once again, van Fraassen's discussion may give a false impression, since one might suppose that the argument forms part of his general case for favouring constructive empiricism over realism. Yet the argument seems clearly to work against the constructive empiricist too, if it works at all. The ranking premise is no less plausible for evaluation with respect to empirical adequacy than it is with respect to truth, and so far as I can tell, van Fraassen himself accepts it. Similarly, we have a constructive empiricist version of the no-privilege premise, to the effect that scientists have no reason to suppose that the means by which they generate theories for testing in itself makes it likely that an empirically adequate theory will be among those generated. Recalling that empirically adequate means adequate to everything observable and not just everything observed, this too will seem plausible to someone who endorses the realist version of the premise and, once again, van Fraassen appears to accept it. Constructive empiricism can itself be seen as being based in part on a kind of intermediate scepticism, to the effect that our inductive powers extend only to the limits of the observable, but this form of scepticism is orthogonal to the one articulated by the argument from underconsideration. The argument from underconsideration is thus especially salient neither as part of an argument for constructive empiricism nor as an argument against Inference to the Best Explanation. Its interest is more general, since it applies to many models of theory evaluation and views of the proper aims of science.

III

Let us now consider the argument from underconsideration in its own right. There are several quick replies that immediately suggest themselves. We may simply

deny either or both of the premises. That is, we may insist either that scientists are capable of absolute and not only comparative evaluation or that their methods of theory generation do sometimes provide them with good reason to believe that the truth lies somewhere among the theories they have generated. These responses may well be correct but, baldly asserted, they lead to an unsatisfying standoff between those who believe in absolute evaluation or privilege and those who do not. Moreover, it seems undeniable that scientists' actual evaluative practises do include a strong comparative element, and one that is reflected in the most popular accounts of confirmation. Examples of this include the use of 'crucial' experiments and the distribution of prior probabilities between the available hypotheses (cf. Sklar 1985, pp.151-53).

Another obvious reply would be to concede some force to the sceptical argument but to deny that it undermines the rationality of science. As we have seen, the ranking assumption grants to the scientist considerable inductive powers. In particular, it allows that theory change is a truth-tropic process, so that later theories are always likelier to be correct than those they replace. Thus we might maintain that science is a progressive activity with respect to the aim of truth, even if scientists are never in a position rationally to assert that the best theory of the moment is actually true. (This view would be a kind of inductively-boosted Popperianism.) More ambitiously, it might be argued that this truth-tropism even justifies scientific belief, by appealing to the scientist's desire to avoid ignorance as well as error. But the cost of these truth-tropic approaches is high, since there are various aspects of scientific activity that appear to require absolute evaluations. The most obvious of these is the practical application of science. In order to decide whether to administer a drug with known and serious side-effects, one needs to know how likely it is that the drug will effect a cure, not merely that it is likelier to do so than any other drug. Absolute evaluations also seem indispensable to 'pure' research, for example to the decision over whether it is better to develop the best available theory or to search for a better alternative.

IV

The quick replies I have mentioned are not to be disdained, but they concede too much to the argument from underconsideration. The nub of the argument is the claim that there is an unbridgeable gap between comparative and absolute evaluation. This gap is, however, only a plausible illusion.

The most straightforward way to eliminate a gap between comparative and absolute evaluation would be by exhaustion. If the scientist could generate all possible competitors in the relevant domain, and he knew this, then he would know that the truth is among them. Given the reliability that the ranking premise grants, he would also know that the best of them is probably true. This brute-force solution, however, seems hopeless, since it takes a wildly exaggerated view of the scientist's abilities. Even granting that we can make sense of the notion of all possible competitors, how could the scientists possibly generate them all?

But collapsing the distinction between relative and absolute evaluation does not require exhaustion. The scientist does not have to know that he has considered all the competitors, only that one of the those he has considered must be true, and for this he needs only a pair of contradictories, not the full set of contraries. It is enough that the scientist consider a theory and its negation, or the claim that a theory has a probability greater than one-half and the claim that it does not, or the claim that X is a cause of some phenomenon and the claim that it isn't, or the claim that an entity or process with specified properties exists or it doesn't. Since scientists are plainly capable of considering contradictories and the ranking premise entails that, when they do, they will be able to determine which is true, the argument from underconsideration fails.

The sceptic has two natural replies to this objection from contradictories. The first is to modify and restrict the ranking premise, so it concedes only the ability to rank contraries, not contradictories. But while the original ranking premise is epistemically over-generous, it is not clearly over-generous in this way. Scientists do, for example, compare the likelihood of the existence and non-existence of entities, causes and processes. So the sceptic would owe us some argument for denying that these comparisons yield reliable rankings while accepting the reliability of the comparisons of contraries. Moreover, it is not clear that the sceptic can even produce a coherent version of this restricted doctrine. The problem is that a pair of contraries entails a pair of contradictories. To give a trivial example, (P&Q) and -P are contraries, but the first entails P, which is the contradictory of -P. Indeed, all pairs of contraries entail a pair of contradictories, since one member of such a pair always entails the negation of the other. Suppose then that we wish to rank the contradictories T1 and -T1. If we find a contrary to T1 (say T2) that is ranked ahead of T1, then -T1 is ranked ahead of T1, since T2 entails -T1. Alternatively, if we find a contrary to -T1 (say T3) that is ranked

ahead of $-T1$, then $T1$ is ranked ahead of $-T1$, since $T3$ entails $T1$. So it is not clear how to ban the ranking of contradictories while allowing the ranking of contraries.

The second natural reply the sceptic might make to the objection from contradictories would concede contradictory ranking. For in most cases, only one of a pair of contradictories would mark a significant scientific discovery. Not to put too fine a point on it, usually one member of the pair will be interesting, the other boring. Thus if the pair consists of the claim that all planets move in ellipses and the claim that some don't, only the former claim is interesting. Consequently, the sceptic may concede contradictory ranking but maintain that the result will almost always be that the boring hypothesis is ranked above the interesting one. In short, he will claim that the best theory is almost always boring, so the scientist will almost never be in a position rationally to believe an interesting theory.

This concession substantially changes the character of the argument from underconsideration, however, and it is a change for the worse. Like most important sceptical arguments, what made the original argument from underconsideration interesting was the idea that it might rule out reasons for belief, even in cases where the belief is in fact true. (Compare Hume's general argument against induction: he does not argue that the future will not resemble the past, but that, even if it will, this is unknowable.) With the concession, however, the argument from underconsideration reduces to the claim that scientists are unlikely to think of the truth. The idea that scientists are only capable of relative evaluation no longer plays any role in the argument, since ranking of contradictory theories has collapsed the distinction between relative and absolute evaluation, and the argument reduces to the observation that scientists are unlikely to think of interesting truths, since they are hidden behind so many interesting falsehoods.

So the revised argument is substantially less interesting than the original. But the situation is worse than this. For scientists do in fact often rank interesting claims ahead of their boring contradictories. The revised argument thus faces a dilemma. If it continues to grant that scientists are reliable rankers, then the fact that interesting claims often come out ahead refutes the claim that scientists do not generate interesting truths. If, on the other hand, reliable ranking is now denied, we have lost all sense of the original strategy of showing how even granting scientists substantial inductive powers would be insufficient for rational belief.

V

The argument from underconsideration depends on a gap between relative and absolute evaluation. I have suggested that contradictory ranking closes that gap and that the argument cannot be modified to reopen it without substantial loss of interest or force. What I will argue now is that the original argument is fundamentally flawed even if we restrict our attention to the ranking of contraries. Given an uncontroversial feature of the way scientists rank theories, the two premises of the argument from underconsideration are incompatible.

Faced with the problem of justifying scientists' methods of evaluation, one may forget how difficult it is even to describe them. This is exacerbated by the general tendency of epistemologists to focus on normative issues at the expense of descriptive ones. In any event, the descriptive project has turned out to be enormously challenging. As the paradox of the ravens and the new riddle of induction illustrate, most standard accounts are remarkably crude, leading to the absurd consequence that almost anything is evidence for anything else. Moreover, as one might expect for any inquiry at such a primitive stage, there is little consensus about even the most basic features that a correct account should include. At least one feature of theory evaluation, however, is almost universally acknowledged, not least among those eager to cast doubt on the possibility of rational belief in science. This is the essential role played by background theories: theories already accepted, if only tentatively, at the time when a new theory is tested. These theories influence the scientists' understanding of the instruments they use in their tests, the way the data themselves are to be characterised, the prior plausibility of the theory under test, and bearing of the data on the theory. (The importance of background theories and their bearing on realism have been emphasised by Richard Boyd in many articles (e.g. Boyd (1985).)

Scientists rank new theories with the help of background theories. According to the ranking premise of the argument from underconsideration, this ranking is highly reliable. For this to be the case, however, it is not enough that the scientists have any old background theories on the books with which to make the evaluation: these theories must be *probably true*, or at least probably approximately true. If most of the background theories were not even approximately true, they would skew the ranking, leading in some cases to placing an improbable theory ahead of a probable competitor, and perhaps leading generally to true theories, when generated, being ranked below

falsehoods. The ranking premise would be violated. So the ranking premise entails that the background is probably (approximately) true. The problem for the argument from underconsideration then appears on iteration. These background theories are themselves the result of prior generation and ranking, and the best of the theories now being ranked will form part of tomorrow's background. Hence, if scientists are highly reliable rankers, as the ranking premise asserts, the highest ranked theories have to be absolutely probable, not just more probable than the competition. This is only possible if the truth tends to lie among the candidate theories the scientists generate, which contradicts the no-privilege premise. Hence, if the ranking premise is true, the no-privilege assumption must be false, and the argument from underconsideration self-destructs.

Given the role of background in theory evaluation, the truth of the ranking premise entails the falsity of the no-privilege premise. Moreover, since the ranking premise allows not only that scientists are reliable rankers, but also that they know this, the situation is even worse. If a scientist knows that her method of ranking is reliable, then she is also in a position to know that her background is probably true, which entails that she is capable of absolute evaluation. Thus knowing that she is capable of comparative evaluation (and, perhaps, reading this essay) enables the scientist to know that she is capable of absolute evaluation, and the claim of the ranking premise that the scientist knows that she is only capable of reliable comparative evaluation must be false.

So the initially plausible idea that scientists might be completely reliable rankers yet arbitrarily far from the truth is an illusion. Might the sceptic salvage his case by weakening the ranking premise, as he was tempted to do in response to the objection from contradictories? I do not think this will help. Of course, if ranking were completely unreliable, the sceptic would have his conclusion, but this just takes us back to Hume. The point of the argument from underconsideration was rather to show that the sceptical conclusion follows even if we grant scientists considerable inductive powers. So the sceptic needs to argue that, if scientists were moderately but not completely reliable rankers, the connection between the best theory and the truth would be severed. Our sceptic has not, however, provided us with such an argument, and there is good reason to believe that no sound argument of this sort exists. For the level of reliability seems to depend, not on just on the degree of reliability of the prior ranking of background theories, but on their verisimilitude.

To see this, suppose that reliability did depend only on the reliability of the prior ranking process by which the background theories were selected. Consider now two isolated scientific communities that are equally reliable rankers, but who in the past generated quite different ranges of candidate theories and so come to have quite different backgrounds. One community was lucky enough to generate true theories, while the other was uninspired enough to generate only wildly erroneous ones. If present reliability depended only prior ranking, we would have to suppose that these two communities are now equally reliable rankers of new theories, which is clearly incorrect. The general point is that the level of reliability a background confers depends on its *content*, not just on the method by which it was generated, and that what matters about the content is, among other things, how close it is to the truth. Consequently, even though scientists are in fact only moderately reliable rankers, this does not sever the connection between relative and absolute evaluation. Even moderately reliable ranking is not compatible with the claim that scientists' methods may leave them with theories that are arbitrarily far from the truth. In other words, even moderately reliable ranking requires moderate privilege.

VI

The moral of the story is that certain kinds of intermediate scepticism, of which the argument from underconsideration is one example, are incoherent. Because of the role of background beliefs in theory evaluation, what we cannot have are inductive powers without inductive achievements. At the beginning of this essay, I distinguished the argument from underconsideration from the better known argument from underdetermination. Having seen what is wrong with the former, however, it appears that a similar objection applies to the latter, and I want now briefly to suggest why this may be so.

The central claim of the argument from underdetermination is sometimes expressed by saying that, however much evidence is available, there will always be many theories that are incompatible with each other but compatible with the evidence. This version of underdetermination, however, ought not to bother the realist, since it amounts only to the truism that the connection between data and theory is and always will be inductive. Like the argument from underconsideration, an interesting version of the underdetermination argument is an intermediate scepticism which attempts to show

that rational belief is impossible even granting the scientist considerable inductive powers. Such a version of the underdetermination argument is an argument from inductive ties. The central claim is that, although some theories are better supported by the evidence than others, for any theory there must exist a competitor (which scientists may not have generated) that is equally well-supported, and this situation remains however much evidence the scientist has. The argument thus allows that scientists are reliable rankers, but insists that the ranking will not discriminate between every pair of competing theories. In particular, it is claimed that this 'coarse' ranking is such that, however much evidence a scientist has, there exist competitors to the highest ranked theory which, if they were considered, would do just as well. Consequently, even if one of the theories the scientist has actually generated is ranked ahead of all the others, he has no reason to believe that this one is true, since he has only avoided a tie through lack of imagination.

Coarse ranking is not quite the same as moderately reliable ranking; the difference is roughly that between a degree of ignorance and a degree of error. Nevertheless, the objection from the background seems to apply here too. Even coarse ranking requires that most of the background theories be close to the truth. If they were not, we would have more than a failure of discrimination; we would have misranking. In other words, even if the underdeterminationist is correct in claiming that there will always in principle be ties for the best theory, this does not support the conclusion that the theories we accept may nonetheless be arbitrarily far from the truth. To get that conclusion would require abandoning the concession that coarse ranking is reliable, as far as it goes, and we are back to an indiscriminating Humean scepticism about non-demonstrative inference.

The underdeterminationist might respond to the objection from the background by 'going global'. He could take the unit of evaluation to be the full set of candidate beliefs a scientist might endorse at one time, rather than a particular theory. The point then would be that there are always ties for the best total set of beliefs. By in effect moving the background into the foreground, the objection from the background appears to be blocked, since what is evaluated now always includes the background and so cannot be relative to it. At the same time, the argument appears able to grant the scientist considerable inductive powers, since it can allow that not all consistent sets are equally likely or equally ranked, and that the higher ranked sets are more likely to be correct than those ranked below them.

I do not think this response is successful. One difficulty is that the global version of the underdetermination argument does not respect the fact that scientists' actual methods of evaluation are local and relative to a (revisable) background. Consequently, although the argument makes a show of granting scientists some sort of inductive powers, it does not grant reliability to the methods scientists actually employ. The reliability of the actual practice of local ranking relative to background cannot be accommodated within this global version without ruining the argument since, as we have seen, local reliability requires that the background be approximately true, the consequence the underdeterminationist is trying to avoid.

A further and related difficulty with the global argument is that it appears tacitly to rely on an untenable distinction between methodological principles and substantive belief. The argument suggests a picture in which the principles of evaluation float somehow above the varying global sets of candidate beliefs, permitting a common scheme of ranking to be applied to them all. Since beliefs about inductive support (such as what is evidence for what) are themselves part of the scientists' total set of beliefs, however, this picture is untenable. What are we to put in its place? If we could say that all the sets shared the same principles, this would perhaps suffice for the argument, but we cannot say this. The problem is not simply that these principles will in fact vary, but that the very notion of a division of the elements of a global set into those that are methodological principles and those that are substantive beliefs is suspect.

There are two reasons for this suspicion. Notice first that, unlike the principles of deductive inference, reliable principles of induction are *contingent*. (This is the source of the Humean problem.) A pattern of non-demonstrative inference that generally takes us from truth to truth in this world would not do so in some other possible worlds. Moreover, although this is perhaps somewhat more controversial, the principles also appear to be *a posteriori*. Given all this, it is difficult to see why they are not tantamount of substantive claims about our world. A second reason for treating the distinction between principle and belief with suspicion pushes from the other side and appeals to the main theme of this essay: the role of the background in evaluation. Given this role, it is unclear on what basis one is to deny that the substantive theories in a global set are themselves also principles of evaluation.

The intermixture of methodological principle and substantive belief, in part a consequence of the essential role of background belief in theory evaluation, makes it

unclear how even to formulate the global argument, and in what sense the argument grants the scientist reliable inductive powers. The intermixture of principle and belief is also perhaps the root cause of the failure of the two forms of intermediate scepticism I have considered in this essay: it explains why it proves so difficult to grant the reliability of evaluation without also admitting the correctness of theory.

VII

'Of course! Why didn't I think of that!' The distinction between being able to generate the right answer and seeing that an answer is correct once someone else has proposed it is depressingly familiar. Meno's slave-boy (or the reader of the dialogue) might never have thought of doubling the square by building on its diagonal, but he has no trouble seeing that the answer must be correct, once Socrates suggests it. And it is apparently no great leap from this truism that there is a distinction between generation and evaluation, between the context of discovery and the context of justification, to the thought that powers of evaluation are quite distinct from powers of generation, that we might be good at evaluating the answers we generate, yet bad at generating correct answers. Hence the thought that scientists might be reliable rankers of the conjectures they generate, yet hopeless at generating conjectures that are true, or close to the truth. Yet this thought turns out to be mistaken, falling to the elementary observation that the scientists' methods of evaluation work relative to a set of background beliefs and that these methods can not be even moderately reliable unless that background is close to the truth. Hence the failure of the argument from underconsideration and of at least some versions of the argument from underdetermination. Of course, in particular cases scientists fail to generate answers that are even approximately correct, but the idea that they might always so fail even though their methods of evaluation are reliable is incoherent. Scientists who did not regularly generate approximately true theories could not be reliable rankers.

What is the bearing of these considerations on scientific realism? Both the arguments from underconsideration and from underdetermination threaten the view that scientists may have rational grounds for believing that a theory is at least approximately true; insofar as these arguments have been turned, the realist who believes in the existence of such grounds will be comforted. It is important, however, to emphasise what has not been shown. I have argued against certain intermediate scepticisms, but

have suggested no answer here to wholesale inductive scepticism. Moreover, I have not tried to show that all intermediate arguments are untenable. In particular, I have not argued against van Fraassen's own intermediate position, which depends in part on the claim that scientists' inductive powers extend only to statements about the observable. On this view, what the scientist is entitled to believe is not that theories are true, but only that they are empirically adequate, that their observable consequences are true. The objection from the background would gain purchase here if it could be shown that, in order for scientists reliably to judge the empirical adequacy of their theories, their background theories must themselves be true, not just empirically adequate. I suspect that this is the case, but I have not attempted to argue it here.

The role of the background in theory evaluation is something of a two-edged sword. It defeats some sceptical arguments, but it also shows both that the realist must take care not to exaggerate scientists' inductive powers and how much even a modest realism entails. Even the most fervent realist cannot afford to claim that scientists are completely reliable thinkers since this would require that all their background beliefs be true, a hopelessly optimistic view and one that is incompatible with the way the scientific background changes over time. The objection from the background drives home the point that realists must also be thoroughgoing fallibilists, allowing the possibility of error not just about theory and the data which support it, but also about the assessment of support itself. The argument of this paper also shows that the realist cannot maintain that scientists are good at evaluation while remaining agnostic about their ability to generate true theories. Reliable evaluation entails privilege, so the realist must say that scientists do have the knack of thinking of the truth. This ability is, from a certain point of view, somewhat surprising, but it remains in my view far more plausible than the extreme ignorance, substantive and methodological, that a coherent critic must embrace.

Department of History and Philosophy of Science
Cambridge University
Free School Lane
Cambridge CB2 3RH

FOOTNOTE

1 I am grateful for the very helpful comments of Jeremy Butterfield, Gavin Ferris, Chris Daly, Michael Gaylard, Mary Hesse, and Tim Williamson.

REFERENCES

Richard Boyd (1985) '*Lex Orandi est Lex Credendi*', in Paul Churchland & Clifford Hooker (eds.) *Images of Science*, (Chicago: University of Chicago Press).

Peter Lipton (1991) *Inference to the Best Explanation*, (London: Routledge).

Lawrence Sklar (1985) *Philosophy and Spacetime Physics*, (Berkeley: University of California Press).

Bas van Fraassen (1989) *Laws and Symmetry*, (Oxford: Oxford University Press).