The Philosophical Quarterly, Vol. 51, No. 204 July 2001
ISSN 0031–8094

WHAT DID HUME REALLY SHOW ABOUT INDUCTION?

By Samir Okasha

I. INTRODUCTION

My aim in this paper is to challenge some orthodox views concerning the problem of induction. In particular I hope to cast doubt on the opinion, widespread among contemporary philosophers, that Hume's argument against induction is irrefutable on its own terms. (A handful of Hume commentators deny that Hume *did* intend to propound a sceptical thesis concerning induction; they should read me as dealing with the argument in *Enquiry* IV which appears to be, and has been widely interpreted as, an argument against induction.) Before I embark on so immodest-sounding a project, a few words of explanation are in order.

The philosophical importance of Hume's argument against induction cannot be doubted. The impact of the argument on twentieth-century philosophy of science in particular has been profound. Popper's entire theory of knowledge was based on the supposition that Hume had proved the irrationality of using induction: hence his attempt to show that science does not actually proceed inductively at all. Reichenbach also accorded great significance to Hume's argument, which he considered the 'heaviest blow' against empiricism, though he thought induction could be given a sort of justification.[1] Carnap regarded the traditional Humean formulation of the problem of induction as 'not the right way' of asking the question, but still attributed great importance to the issue of induction's rationality.[2] The approaches to induction favoured by Popper, Reichenbach and Carnap find few adherents today, but contemporary philosophers of science continue to invest Hume's arguments with a high degree of philosophical significance.

[1] H. Reichenbach, *Experience and Prediction* (Chicago UP, 1938), p. 347.
[2] R. Carnap, 'Inductive Logic and Inductive Intuition', repr. in I. Lakatos (ed.), *The Problem of Inductive Logic* (Amsterdam: North-Holland, 1968), pp. 258–67, at p. 258.

But what is the purpose of re-examining Hume's sceptical argument, which is, after all, one of the best known arguments in the history of philosophy? My reasons are two. First, despite the attention which Hume's argument has received, there appears to be little consensus on what exactly the argument shows, nor on how the threat of inductive scepticism should be met. Secondly, many philosophers have reacted to Hume's argument in a way which in my view could not possibly be correct.

The structure of this paper is as follows. In §II, I examine three widely held opinions about the problem of induction which seem to stand in serious tension with one another. In the light of this tension, I look at Hume's argument for inductive scepticism in §III. I show that the argument as it stands is valid but unsound, for it rests on Hume's flawed description of our inductive behaviour. In §IV, I argue that Hume's argument can only be converted into a sound one if our inductive behaviour can be characterized as a process of rule-governed ampliation. Drawing on some Bayesian ideas, I argue that our inductive behaviour probably cannot be characterized in this way, so our immunity from Hume is secured. In §V, I compare my response to Hume with some other well known approaches to the problem of induction.

II. POPULAR OPINIONS ABOUT THE PROBLEM OF INDUCTION

Here are three claims which represent very widely held opinions about the problem of induction:

(a) Hume was not simply objecting to inductive inferences on the score of their logical invalidity: he was doing something much deeper than that. In other words, Hume's sceptical argument does not just amount to the boring argument which says 'Past data never *entail* anything about the future; therefore it is *possible* that the future will turn out differently from how we believe; therefore knowledge of the future is impossible'.

(b) The *description* Hume gave of our inductive practices was extremely inadequate. He claimed that in forming expectations of the future we assume that nature is uniform, i.e., that the future will resemble the past: 'All our experimental conclusions proceed upon the supposition that the future will be conformable to the past'.[3] This description of our inductive habits is a serious over-simplification at best.

(c) Hume's sceptical argument is irrefutable.

[3] Hume, *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, ed. P.H. Nidditch (Oxford: Clarendon Press, 1975), p. 36.

I shall argue that these three opinions stand in serious tension with one another. First it is worth commenting briefly on each.

Claim (a) is clearly true. The boring argument which objects to inductive inferences because they are logically invalid shows only that our beliefs about the future cannot be certain, which is obviously true but totally unthreatening. But Hume claimed to have shown that our beliefs about the future are *unreasonable*, not just that they cannot be certain: 'even after the observation of the frequent or constant conjunction of objects, we have *no reason* to draw any inference concerning any object beyond those of which we have had experience'.[4] In other words, Hume was arguing for inductive *scepticism*, not just inductive fallibilism, which is all that the boring argument establishes.

Claim (b) is clearly true too, as Goodman pointed out so forcefully. To say that we assume that nature is uniform, as Hume does, is to say nothing. Nature cannot be uniform in every respect, and cannot fail to be uniform in some respects if it is to be describable at all. We extrapolate some regularities into the future, but not others. All of the emeralds we have examined in the past have been grue, but that does not create the expectation that the next emerald we observe will be grue. Of course artificial predicates like 'grue' are not required to reveal the inadequacies of Hume's description: they simply make the point particularly dramatically. All of the Costa Ricans I have ever met are philosophers, but that does not lead me to believe that all Costa Ricans are philosophers, nor that the next Costa Rican I meet will be a philosopher. Expectation-formation takes place in the light of a vast store of background information.

An example of Putnam's neatly illustrates how the presence of background theory undermines Hume's simple model of inductive inference. Usually when you slam two small objects together, nothing in particular happens. However, when you slam two high-speed neutrons together, you get an atomic explosion. This is not a chance event: scientists predicted that it would happen. Their prediction was entirely 'theory-driven': it clearly did not involve the extrapolation of a past regularity into the future, as per Hume's account. There can be no serious doubt that Hume's description of our inductive practices was badly deficient.

Claim (c) is *not* clearly true, though it is very widely believed. With a handful of exceptions, most philosophers have despaired of finding a satisfactory answer to Hume's arguments.

The tension between (a), (b) and (c) is not difficult to see. Anyone who is committed both to (b), the claim that Hume's description of our inductive

[4] Hume, *A Treatise of Human Nature*, ed. P.H. Nidditch (Oxford: Clarendon Press, 1978), p. 139; my italics.

practices was inadequate, and to (c), the claim that Hume's argument is irrefutable, must presumably think that the following is true:

(d) The fact that Hume's description of our inductive practices was so inadequate in no way affects the correctness of his sceptical conclusion. Sceptical considerations of the sort Hume gives will apply whatever the correct description of our inductive practices happens to be.

But (d) appears to conflict with (a). Certainly, if Hume had simply been objecting to inductive inferences on the ground of their logical invalidity, the fact that his description of our inductive practices was so inadequate would in no way affect the cogency of his sceptical argument. But Hume was not doing that, as (a) correctly asserts. It is hard to see how there could be a sceptical argument which in no way depends on what our inductive practices are actually like without simply collapsing into the boring argument which objects to inductive inferences because they are invalid. And yet many philosophers seem to think that Hume provided just such an argument – that is what the conjunction of (a) and (d) amounts to.

One example of a philosopher who endorses both (a) and (d) is Wesley Salmon. According to him, 'Hume's arguments are not peculiar to induction by enumeration or any other special kind of inductive inference; they apply with equal force to any inference whose conclusion might be false when it has true premises'. And in the same article Salmon says 'the suggestion that Hume merely showed the fallibility of induction is a mistake'.[5]

Salmon thinks that Hume's argument creates a strong *prima facie* case for inductive scepticism, and applies to any logically invalid inference; hence it does not depend on the specific description of our inductive practices with which Hume operated. Peter Lipton is another example. He says that Hume's sceptical argument 'seems to work, whatever the details of our [inductive] inferences'; he also says that the prospects for finding a satisfying response to Hume are 'bleak'.[6] Now if Hume's argument simply amounted to the boring one that objects to inductive inferences because they are invalid, the prospects for finding a satisfying solution could not be called 'bleak', so Lipton clearly thinks that Hume's argument amounts to something more.

In my view it is far from clear how any argument could succeed in doing what Salmon and Lipton think Hume's argument does. For clearly the only property that all logically invalid inferences have in common is the property

[5] W. Salmon, 'The Justification of Inductive Rules of Inference', repr. in Lakatos (ed.), *The Problem of Inductive Logic*, pp. 24–66, at pp. 29 and 30.

[6] P. Lipton, *Inference to the Best Explanation* (New York: Routledge, 1991), pp. 14 and 12.

of logical invalidity itself. So a sceptical argument that applies to all logically invalid inferences must presumably object to these inferences *because* they are logically invalid. But objecting to an inference because it is logically invalid does not show that the inference is unreasonable; it only shows that it is not infallible, which everybody concedes. In other words, if it makes no difference to the sceptic what our inductive practices are like, beyond the fact that they are inductive, his argument can only amount to the boring one which establishes merely inductive fallibilism. No *successful* argument for inductive scepticism can be invariant with respect to the correct description of our inductive practices.

Supposing this is correct, what follows? It does not follow that our actual inductive practices are immune from a Hume-type sceptical argument simply because Hume misdescribed them. It does follow, however, that we would be wrong to assume *automatically*, as Salmon and Lipton do, that our actual inductive practices are susceptible to Humean sceptical attack simply because they are inductive (i.e., non-deductive). Clearly the crucial questions are (i) what would our inductive practices have to be like for us to be susceptible to a Hume-type sceptical argument? And (ii) are our inductive practices actually like that or not? In other words, if Hume had described our inductive behaviour accurately, would he still have been able to run his sceptical argument? Only when this question has been addressed shall we be able to decide whether Hume's inductive scepticism is really irrefutable, as popular opinion holds.

Before examining this question, it is worth asking why the Salmon–Lipton view is so widely adopted. So far as I can tell, there are two reasons. First, many philosophers seem to think that the following is a legitimate generalization of Hume's point: 'Our inductive methods cannot be justified deductively; and they cannot be justified inductively without circularity' (formulations of this sort abound in the literature). And it certainly *looks* as if this little argument works, whatever the details of our 'inductive methods'. But I shall argue that here appearances deceive. Secondly, many philosophers believe that something *like* Hume's description of our inductive behaviour must in the end be right. But I shall argue that in this they are wrong.

III. HUME'S ARGUMENT FOR INDUCTIVE SCEPTICISM

The general outline of Hume's argument is well known, so my exposition will be relatively brief. I shall focus on the version of the argument from *Enquiry* IV.

Hume starts by giving a description of how we actually arrive at our opinions about the unobserved. The description is simple: we observe past regularities, and project them into the future. He expresses this by saying that all our experimental reasoning 'proceeds upon the supposition' that nature is uniform, i.e., that the future will resemble the past (p. 17: I use the expressions 'Nature is uniform' and 'The future will resemble the past' interchangeably throughout, as does Hume). He then asks whether it is *rational* to draw conclusions from past experience in the way he has described. Is the supposition of nature's uniformity, which underlies our reasoning from experience, one that we have good reason to make? According to Hume (p. 35), reasoning can be divided into two kinds: demonstrative reasoning, which concerns 'relations of ideas', and probable reasoning, which concerns 'matters of fact and existence'. That nature is uniform cannot be demonstrated, he continues, for the only propositions susceptible of demonstration are those whose denials are self-contradictory, and it 'implies no contradiction that the course of nature may change'. That leaves us with probable reasoning, or reasoning from experience. But this Hume dismisses swiftly, as quite inadequate to the task. Probable reasoning *presumes* that nature is uniform, and cannot therefore help to establish it. Any attempt to argue for the uniformity of nature on the basis of past experience must be 'evidently going in a circle, and taking that for granted, which is the very point in question' (p. 36). Hume is thus led to the sceptical conclusion that our propensity to project past regularities into the future has no rational foundation at all, but is simply a brute fact about human nature.

Hume's sceptical argument can therefore be broken down as follows:

1. Arguments from past experience, or 'probable arguments', 'proceed upon the supposition' that nature is uniform
2. Unless we have reason to believe that nature really is uniform, we have no reason to believe the conclusions of arguments from past experience
3. The proposition 'Nature is uniform' cannot be established by demonstrative argument, as its negation is not self-contradictory
4. The proposition 'Nature is uniform' cannot be established by probable argument, as we only have reason to believe the conclusions of probable arguments if we already have reason to believe that nature is uniform, as (2) says
5. If a proposition cannot be established by either demonstrative or probable argument, then we have no reason to believe it
6. So we have no reason to believe that nature is uniform [from (3)–(5)]
7. So we have no reason to believe the conclusions of probable arguments [from (6) and (2)], i.e., past experience gives us no reason to believe anything about the future.

This argument is valid; the issue is whether the premises are true. I shall assume that (5) is true – probable and demonstrative arguments exhaust the means by which we can come to know the truth of propositions. (A demonstrative argument, for Hume, has *a priori* premises and is deductively valid, while a probable argument has one or more empirical premises. Premise (5) says, in effect, that the class of synthetic *a priori* truths is empty.) It is clear that (3) is true: 'Nature is uniform' cannot be demonstrated. Furthermore it is clear that if (2) is true, then (4) is true too: if the reasonableness of probable arguments is conditional on our having reason to believe that nature is uniform, then no probable argument can give us reason to believe that nature is uniform. So everything hangs on premises (1) and (2).

Premises (1) and (2) have been the focus of much exegetical controversy, so it is worth commenting on them briefly.[7] Clearly the premises are not independent: the first is supposed to support the second. It is because Hume thinks (1) that arguments from past experience 'proceed upon the supposition' of nature's uniformity that he thinks (2) we need reason to believe in nature's uniformity, if arguments from past experience are to be reasonable. Appreciating this helps us to see what Hume meant by (1), and why he asserted (2).

In my view, (1) is basically just Hume's description of how we arrive at our expectations of the future. Obviously, past data can be extrapolated into the future in many different ways. As Hume says (p. 29), 'when I see ... a billiard-ball moving in a straight line towards another ... may I not conceive, that a hundred different events might as well follow from that cause?'. However, Hume was impressed by the fact that typically we select one of the possible extrapolations as the one most likely to happen – we believe the first billiard-ball will impart motion to the second. He offers an account of why this is, namely, because we follow the rule that past regularities will continue to hold in the future, i.e., that nature is uniform. In saying that our arguments from past experience 'proceed upon the supposition' that nature is uniform, Hume is thus saying that the uniformity principle is the rule that guides us in our formation of future expectations.

Hume's reason for asserting (2) is obvious. We need a reason to believe in nature's uniformity, if our beliefs about the future are to be reasonable, because those beliefs have been formed according to the principle that nature *is* uniform. Using this principle is not guaranteed to lead to successful predictions – there is 'no contradiction' in thinking that the course of nature

[7] What Hume meant by premise (1), and why he asserted premise (2), are the main points of dispute between, on the one hand, D. Stove, *Probability and Hume's Inductive Scepticism* (Oxford: Clarendon Press, 1973), and J.L. Mackie, *The Cement of the Universe* (Oxford: Clarendon Press, 1974), and on the other, B. Stroud, *Hume* (London: Routledge & Kegan Paul, 1977). My interpretation of Hume's argument is similar to Stroud's.

may change. Thus Hume thinks that some reason must be given for employing this principle and not another one. Corresponding to the inductive principle which Hume says we use – 'Assume that nature is uniform' – is an empirical assumption about the world, that nature *is* uniform. If we are to be justified in using this inductive principle, rather than any other, we must have some reason for thinking the empirical assumption true. Hence (2).

The problem with Hume's argument as it stands is obvious. The argument rests on his flawed description of our inductive behaviour. The inductive principle which he says we employ is one which we most certainly do not employ, as Goodman showed. We believe that some observed regularities will continue into the future, but not others. The first two premises of Hume's argument are therefore false. Hume has not found any inductive rule which we use to extrapolate past data into the future, and so he has not found any empirical assumption about the world on which the reasonableness of our arguments from experience is conditional. The sceptical argument Hume actually gave is valid but unsound.

The key question is then this: if Hume's deficient description of our inductive behaviour were replaced with the 'correct' description, could his sceptical argument be converted from a valid one into a sound one? Presumably those philosophers who endorse the Salmon–Lipton position criticized earlier would answer 'Yes'.

But I think the answer is actually 'Not necessarily'. Only if our inductive practices can be characterized in a certain fairly specific way will it be possible to construct a Hume-type sceptical argument to undermine them. The abstract form of Hume's argument is this: he isolates an empirical assumption about the world and claims that arguments from past experience are only reasonable if there is reason to make this assumption; but only an argument from past experience could give us reason to do that. The empirical assumption corresponds to the rule of inductive inference which Hume says we use, a rule which in point of fact we do not use. (In effect, the empirical assumption says 'The world is such that the rule of inference is truth-conducive'.) Now it is tempting to assume that when someone uncovers the 'real' inductive rule (or rules) which we use, an empirical assumption will correspond to *it* (or them), and an argument perfectly analogous to Hume's will immediately fall into place. But this line of reasoning rests on a crucial assumption, that we do actually use rules to form new beliefs from data which do not entail them. If this assumption fails, then a Humean sceptical argument cannot be run.

The point stressed earlier is crucial here: the inductive sceptic's argument cannot just turn on the fact that we do have opinions that go beyond what our data entail. That point establishes only inductive fallibilism. According

to my analysis, what is required to bridge the gap between inductive fallibilism and inductive scepticism is the assumption that our opinions have been arrived at by application of one or more rules of inductive inference. If that assumption were true, then Hume's sceptical argument would go through. For using an inductive rule requires having reason to believe that the world is arranged in one way rather than in another – for in most possible worlds the rule will lead you astray. And what licenses you, Hume will ask, in thinking that the world *is* like that? But if we do *not* in fact arrive at our opinions by using inductive rules, then Hume's sceptical argument cannot be converted into a sound one.

IV. NO RULES OF INDUCTION, NO HUMEAN ARGUMENT

In this section I argue that our actual inductive practices are not susceptible to Humean attack, because they do not involve the use of inductive or ampliative rules. First, I sketch a picture of what our inductive practices *could* be like, if we did not use inductive rules. Next, I offer some reasons for thinking that this is what our inductive practices *are* actually like. Finally, I explain in more detail why this secures our immunity from Humean inductive scepticism.

An inductive or ampliative rule, if such a thing existed, would be a rule for forming new beliefs on the basis of evidence, where the evidence does not entail the belief. One type of inductive rule would be a rule for forming future predictions on the basis of past data, of the sort that Hume envisaged. Another type would be a rule for forming empirical generalizations on the basis of positive instances, or for estimating the frequency of a given trait in a population on the basis of an observed sample. Yet another type would be a rule for arriving at theoretical hypotheses on the basis of observational data. In response to the suggestion that we do not actually follow any such rules, it may seem tempting to reply as follows: 'Clearly we have opinions of various sorts about the unobserved. Clearly these are based somehow on past experience – they are not random. Therefore there *must* be some rules which we use to infer these opinions from past experience.' I suspect that this line of reasoning underpins the widespread view that it makes sense to talk of the 'inductive method', even if no one knows quite what it is. But the reasoning is flawed, as we shall see.

The idea that there are no inductive rules which we follow is tantamount to the idea that we do not strictly engage in inductive *inference* at all – for where there is inference, there is a rule of inference. This idea has been defended before, by Bayesians and by Popper and his followers. Recently

van Fraassen has defended the idea at length from a Bayesian viewpoint,[8] arguing that there exists no 'viable inductive method, inductive logic, or set of canons of induction' (*Laws and Symmetry*, p. 279). Van Fraassen is persuaded in part by theoretical arguments, designed to cast doubt on the idea of 'inductive logic'; but he is also impressed by the repeated failure of philosophers to come up with the inductive rules which we allegedly follow. He stresses (p. 296), rightly in my view, that it would be a mistake to say there *must* exist inductive rules even if no one has discovered them yet, just because we do adopt beliefs that go beyond our data and alter those beliefs in the light of experience in ways that transcend deductive logic. For, he maintains, this ignores the Bayesian model of belief change.

On the Bayesian model, each agent begins with an allocation of subjective probabilities to a set of propositions, an allocation that conforms, if the agent is rational, to the axioms of the probability calculus. The basic way of learning from experience, according to Bayesians, is to conditionalize one's prior probability function on the deliverances of experience. Updating by Bayesian conditionalization does not count as inductive inference in the strict sense, of course. For a rule of inductive inference is supposed to tell you what beliefs you should have, given your data, and the rule of conditionalization does not do that. If you apply the rule of conditionalization to your data, the state of opinion you end up in depends on the state you were in previously; whereas if you apply an inductive rule to your data, the state of opinion you end up in depends on the instructions contained in the rule. So there is the following crucial difference between the rule of conditionalization and the inductive rules whose existence I am disputing: an inductive rule is only going to be truth-conducive if the world is arranged in a particular way, but the same is not true of the rule of conditionalization. The propriety of employing an inductive or ampliative rule, therefore, depends on having reason to believe that the world *is* arranged in that way; not so with the rule of conditionalization. Of course this is not to say that updating by Bayesian conditionalization is guaranteed to take one closer to the truth. It is not: one's prior opinions could have been entirely wrong. But the mere possibility of error has no sceptical implications in itself – it only highlights the obvious fact that empirical enquiry is a fallible business.

According to the strictest Bayesians, conditionalization is the *only* rational way of responding to new evidence. However, the Bayesians' own attempts to justify the rule of conditionalization do not warrant a claim this strong.

[8] B. van Fraassen, 'Empiricism in the Philosophy of Science', in P. Churchland and C. Hooker (eds), *Images of Science* (Univ. of Chicago Press), pp. 245–308; *Laws and Symmetry* (Oxford: Clarendon Press, 1989); and 'The False Hopes of Traditional Epistemology', *Philosophy and Phenomenological Research*, 60 (2000), pp. 253–79.

The most that is shown by the standard diachronic Dutch-book argument, due to Lewis and Teller, is that *if* you adopt a preset plan for changing your opinions, that plan had better be conditionalization.[9] This point is the basis of the liberal version of Bayesianism defended by van Fraassen. Van Fraassen adopts what he calls a 'permissive' conception of rationality: rationality concerns what you may believe, not what you must. Given this conception, van Fraassen is able to countenance all sorts of opinion-change. Conditionalization is the usual way of changing opinion, he thinks, but we are never rationally *compelled* to conditionalize: the scientific imagination can be granted untrammelled freedom in deciding how to modify the overall world-picture in the light of experience, without transgressing the limits of rationality. This concession is vital to the plausibility of Bayesianism, in my view. For sometimes our response to new evidence is to invent a hypothesis or prediction which had not occurred to us before; such responses obviously cannot be modelled as cases of conditionalization. Van Fraassen does think, however, that *rule-governed* violations of conditionalization are irrational – for he thinks the Lewis–Teller theorem shows that. So he argues (*Laws and Symmetry*, ch. 5) that inductive or ampliative rules of the sort criticized above actually violate the demands of Bayesian rationality. Whether this is a legitimate reading of the Lewis–Teller theorem is not a question I wish to broach here. (The answer depends on whether inductive rules can be equated with the 'preset plans' for opinion-change to which the Lewis–Teller theorem applies.) The point I wish to stress is just this: given a liberalized version of Bayesianism of the sort van Fraassen offers, there is no *need* for rules of inference in order to get from data to opinions not entailed by the data, whether or not it is true that there is no *room* for rules. Bayesian conditionalization reveals the dichotomy 'Either there are inductive rules or else our responses to experience are random' to be a false one.

The foregoing provides only the briefest of sketches of what our inductive practices might be like, if we used no inductive rules. Whether the Bayesian model is adequate in all respects is a complex issue, which I do not wish to tackle here. For even if it is granted that Bayesianism shows how we *might* do without inductive rules, it does not follow that we *do* actually do without them; this requires independent argument.

The basic case against the idea that we use inductive rules is straightforward: no one has yet come close to saying what these rules actually are. But there are other reasons for doubt too.

It is important to be clear about the concept of using an inductive or ampliative rule. That we accept some substitution-instances of a particular

[9] P. Teller, 'Conditionalization and Observation', *Synthese*, 26 (1973), pp. 218–58.

rule does not mean that we are following this rule. For example, if we believe that some observed correlations between properties will continue to hold in the future, that does not mean we are using the inductive rule 'Infer that properties correlated in the past will be correlated in the future', for there are plenty of instances of that rule we do not accept. With rules of inference we really do use, e.g., *modus ponens*, we accept *all* substitution-instances that we can recognize as such. Though elementary, this point appears to be overlooked by many who try to specify the inductive rules and principles we allegedly use. For example, Frank Jackson, in a discussion of 'straight-rule' induction, describes as 'undeniable' the claim that 'we all use [the straight-rule] on occasion and take it as rational to do so', where using the straight-rule involves arguing from 'certain Fs being G to certain other Fs being G'.[10] But this talk of 'using a rule on occasion' stretches the concept of a rule beyond breaking point. One might just as well say that a heavy smoker obeys the rule 'Don't smoke' on occasion, simply because he takes breaks between cigarettes. It *is* true that we sometimes increase our confidence in statements of the form 'The next F will be G' when we learn the truth of statements of the form 'Some Fs are G'. Bayesian analysis yields a clear account of the conditions under which this is a rational thing to do.[11] Correctly described, Jackson's 'undeniable fact' is easily accounted for by the supposition that we are ordinary Bayesian conditionalizers. It lends no credence to the notion that we use rules of inductive inference, either the 'straight-rule' or any other.

One powerful consideration against 'inductive rules' concerns the role that background beliefs play in mediating our responses to experience. The same experience affects our expectations of the future in different ways, depending on which background beliefs we hold. Bayesian models accommodate this point easily, of course, in the form of prior probability assignments. But the point is highly problematic for the 'inductive rules' idea. For any inductive rule that we allegedly follow, it is easy to imagine background beliefs which would lead us to disregard it. Advocates of inductive rules have typically responded by saying that their rules only apply from a position of informational neutrality. For example, Salmon claimed that the straight-rule was the 'basic inductive rule', but argued that we only use this rule where there are 'meagre data' and no other relevant information to

[10] F. Jackson, 'Grue', *Journal of Philosophy*, 72 (1975), pp. 113–31, at p. 113. Jackson is using the expression 'straight-rule' in a loose and generic sense; what he calls the 'straight-rule' is not identical with the straight-rule of Reichenbach, who introduced the expression.

[11] Suppose that h is 'All Fs are G', e_n is 'The nth F is G', and k is background information. Since h entails e_n for all n, then so long as $P(h/k) > 0$, it follows that $P(e_n/e_1 \,\&\, e_2 \,\&\, \ldots \,\&\, e_{n-1} \,\&\, k)$ must tend to 1 as n grows without limit, as J. Earman shows: see his *Bayes or Bust: a Critical Examination of Bayesian Confirmation Theory* (MIT Press, 1992), p. 78.

hand.[12] Hempel made a similar move in defence of his claim that we always take observational generalizations to be confirmed by their instances. Faced with apparent counter-examples devised by I.J. Good, Hempel argued that these examples depended on an illicit appeal to background information: Good was failing to adhere to the 'methodological fiction' that we know nothing at all about the world beyond the information provided by the observed instance.[13] Elsewhere, Hempel wrote that his principles of confirmation were only supposed to apply in 'theoretically barren contexts'. In my view we should be extremely wary of this Salmon–Hempel move. For one thing, it is far from clear that the idea of informational neutrality, or of a 'theoretically barren context', is actually coherent. In any case, Salmon's and Hempel's admissions that their principles of induction and confirmation only apply from a hypothetical epistemic starting-point, very different from the one in which we find ourselves in real life, is an enormous concession. For this is little different from saying that we do not actually follow the principles in question at all. Perhaps Salmon and Hempel will reply that the background information whose presence renders their principles inapplicable to real-life contexts was itself arrived at by application of those very principles, all the way back to infancy. But one wonders how they could possibly know that this is true. Rather than falling into speculation of this sort, it seems preferable simply to admit that we do *not* use any of the inductive principles that have so far been proposed. That is quite compatible with allowing that certain aspects of our inductive practices appear to be quite well captured by straight-rule induction, 'enumerative' induction, 'hypothetico-deduction', etc., for elegant Bayesian analyses have been given of the conditions under which, when they are stated sufficiently precisely, such principles are valid. For example, Bayesian theory explains why an agent will often increase his confidence in 'All ravens are black' upon observing a black raven. Presuming the agent's initial degree of belief in 'All ravens are black' is neither 0 nor 1, then if the logically strongest proposition he learns is that a given object, already known to be a raven, is black, updating by Bayesian conditionalization will lead him to become more confident of the generalization in question. No inductive rule needs to be posited to explain the change of opinion.

Above I suggested that one reason why Hume's inadequate description of our inductive behaviour is not usually taken to undermine his sceptical

[12] Salmon, 'Replies', in Lakatos (ed.), *The Problem of Inductive Logic*, pp. 74–97, at pp. 77 and 83.

[13] I.J. Good, 'The White Shoe is a Red Herring', *British Journal for the Philosophy of Science*, 17 (1967), p. 322; C. Hempel, 'The White Shoe: No Red Herring', *British Journal for the Philosophy of Science*, 18 (1967), pp. 239–40.

argument is that many philosophers think that something *like* Hume's description must in the end be right. Ironically, Goodman's own work may have contributed to this erroneous view. For Goodman thinks that the basic idea of induction by enumeration is actually correct, so long as it is restricted to the so-called 'projectible' predicates, and he thinks that the extension of the set of projectible predicates is easily specified. But as numerous Bayesian authors have argued, this is just wrong – the distinction between projectible and non-projectible predicates is mistaken in principle: Goodman radically misdiagnosed his own problem.[14] Whether we take 'positive instances' to support universal generalizations or predictions depends on (a) the precise informational content provided by the observation of the 'positive instance', and (b) our background beliefs; it has nothing to do with the type of vocabulary used to describe the data. Goodman's critique of Hume's description of our inductive habits was well taken, but the alternative he proposed was not. If anything like the Bayesian conception of learning from experience is right, no minor modifications of Hume's description will prove correct. I concur fully with Roger Rosenkrantz (p. 84): 'there is no sound principle of induction as traditionally understood. Induction, or learning from experience, is just the process of revising probability assignments in the light of additional information.'

So much for the case against inductive rules. Of course these considerations do not *prove* that our beliefs are not formed from data by using inductive rules. But they do put the burden of proof on the shoulders of those who believe in such rules. I turn now to defending the claim that this has a significant bearing on Humean inductive scepticism.

The basic argument has already been given. Hume argued that the reasonableness of our opinions about the unobserved is conditional on an empirical assumption for which we have no non-circular justification. He argued in that way because he thought the opinions in question were arrived at by extrapolating from past experience using a rule of inference, a rule which would only take us towards the truth if the empirical assumption were true. Quite reasonably, Hume demanded some reason for thinking the empirical assumption *was* true before he would rationally sanction the opinions to which the rule gave rise. Hume's inductive scepticism would be quite right, if we did use rules of inductive inference. But if we do not, then a Humean sceptical argument will not go through.

[14] See in particular R. Rosenkrantz, 'Does the Philosophy of Induction Rest on a Mistake?', *Journal of Philosophy*, 79 (1982), pp. 78–97; van Fraassen, 'Theory Confirmation: Tension and Conflict', in P. Weingartner and M. Czermak (eds), *Epistemology and Philosophy of Science: Proceedings of the Seventh International Wittgenstein Symposium* (Vienna: Hölder-Pichler-Tempsky, 1983), pp. 319–29; and E. Sober, 'A Bayesian Primer on the Grue Problem', in D. Stalker (ed.), *Grue! the New Riddle of Induction* (La Salle: Open Court, 1994), pp. 225–40.

It should not be surprising that the possibility of running a Humean sceptical argument turns out to depend on the actual details of our inductive practices. On the contrary, it would be something of a coincidence if this dependency did *not* obtain, given that Hume was not just making the banal point that some of our beliefs go beyond what the data entail. Furthermore, the specific form of dependency advocated here – that a Humean argument would only work if we used inductive or ampliative rules – does not seem to me counter-intuitive. For in a way there is something inherently irrational about following an inductive rule: it deprives one of the flexibility needed for rationally adjusting one's opinions to experience. To use an inductive rule is to assume that the world is arranged in a particular way, as I have stressed. But presumably, for any assumed arrangement of the world, it should be possible to *imagine* evidence for the world's not being like this – that is the nature of empirical assumptions. If such evidence materialized, the rational thing to do would be to give up the assumption that the world is arranged in the way in question, i.e., to stop using the inductive rule. So following any particular inductive rule does seem less than fully rational. It embodies a fixed commitment to the world's being in a certain state; but *qua* empiricists we should undertake such commitments only provisionally, not hold on to them at all costs. From this point of view, it is no great surprise that the use of inductive or ampliative rules opens the door to Humean scepticism.

Bayesianism has sometimes been regarded as impotent in the face of the problem of induction, because of the ineliminable role of prior probabilities in Bayesian models. This view allows that Bayesian theory provides a good account of how we should *change* our opinions in the light of experience, but notes that the theory does not tell us what our initial opinions should be, beyond requiring that they be probabilistically coherent, and so it is powerless against the Humean sceptic. In my view this pessimistic view is flawed, for it rests on a mistaken assimilation of inductive fallibilism to inductive scepticism. This point is worth spelling out.

It is obvious that, in real life, how we respond to a new datum depends fairly heavily on our background beliefs. If observing a black raven leads me to predict that the next raven will be black, that is because my background beliefs include things like 'Conspecific birds are usually uniform in colour'. So by catering for prior opinions Bayesianism displays fidelity to how our inductive practices actually operate. But the inductive sceptic sees this reliance on background beliefs as question-begging: in tackling the problem of induction, surely we cannot help ourselves to information that is itself inductively based? Thus the sceptic demands that we suppress our background beliefs, adopt the notorious 'theoretically barren' starting-point, and try to justify our response to the data without the help of those beliefs.

But this demand depends crucially on the coherence of the notion of a 'theoretically barren' or information-free starting-point. Within a probabilistic framework, the idea of an information-free starting-point becomes the idea of a prior probability distribution which represents a state of ignorance. Now if there were such a prior probability distribution, i.e., if there were a viable principle of indifference, then the sceptic would have a point. For in that case those who started work with a prior distribution *different* from the distribution corresponding to the state of ignorance would be helping themselves to information they are simply not entitled to. But the idea of a probabilistic representation of ignorance is one of the most widely discredited ideas in probability theory, whose shortcomings do not need to be repeated here – it is very widely accepted that the principle of indifference leads directly to contradiction.[15] Abandoning the principle of indifference means accepting that all coherent prior probability distributions are equally arbitrary. This implies, obviously, that the actual prior distribution we start with cannot be shown to be 'better' than any other. But this simply reflects the fact that there is an inevitably conjectural element, and thus an element of uncertainty, in all empirical enquiry – if there were not, it would not be empirical. To demand that our actual prior probability distribution must be uniquely justified is to demand the impossible; it is the analogue within a probabilistic epistemology of the rationalistic demand that we must *prove* that our beliefs about the world are true. Obviously the demand cannot be met, but that is a point in favour of inductive fallibilism, not inductive scepticism. Once we accept that the notion of a prior distribution which reflects a state of ignorance is chimerical, then adopting any particular prior distribution does not constitute helping ourselves to empirical information which should be suppressed; it simply reflects the fact that an element of guesswork is involved in all empirical enquiry.

This important point can be seen from a different angle. When I rely on my background belief that conspecific birds are typically uniform in colour to help me predict the colour of the next raven, the inductive sceptic objects. Now obviously the sceptic's objection cannot be that my background belief *might* be false – that is true, but irrelevant. The argument must be that the background belief is *unjustified*. What would such an argument look like from within a Bayesian framework, where belief is not an all-or-nothing affair, but comes in degrees? Presumably the sceptic will argue that the high prior probability I allocate to the statement 'Conspecific birds are usually uniform

[15] See van Fraassen, *Laws and Symmetry*, and C. Howson and P. Urbach, *Scientific Reasoning: the Bayesian Approach* (La Salle: Open Court, 1989), for good accounts of the insuperable difficulties that the principle of indifference faces. But see Rosenkrantz, *Inference, Method and Decision* (Dordrecht: Reidel, 1977), for a dissenting view.

in colour' is unwarranted. But we need only ask the sceptic 'What prior probability do you recommend?'. If we could make good probabilistic sense of a state of ignorance, or of a 'theoretically barren context', the sceptic would have a plausible answer to this question. The sceptic could say 'Your prior probability in that statement should be x', where x is the number dictated by the prior probability function that represents the state of ignorance. But if, as I am assuming, the principle of indifference does not work, then the sceptic will be unable to provide an answer on these lines. So his argument fails. It does not beg the question to operate with some particular prior probability distribution if there is no alternative to doing so. Only if the inductive sceptic can show that there *is* an alternative, i.e., that 'information-free' priors do exist, would adopting some particular prior distribution beg the question. Thus we see that the role of prior probabilities in Bayesian models does not imply their impotence in the face of the problem of induction, once we pay close heed to the distinction between fallibilism and scepticism.

To return to the Salmon–Lipton position criticized previously, Salmon and Lipton maintained that a Hume-type sceptical argument can be applied to any inductive, i.e., non-deductive, inference. In a way I actually agree with Salmon and Lipton, but the agreement does not run deep. If the phrase 'inductive inference' is employed strictly, as implying the existence of a rule of inference, then I agree that a Hume-type argument applies to all inductive inferences; but in this sense it is far from clear that we actually engage in inductive inference at all, in my view. Obviously Salmon and Lipton use 'inductive inference' much more liberally, for they take it as an obvious fact that such inferences take place all the time. But, as I have stressed, what is really obvious is just that we do hold opinions that go beyond our data, and that we respond to experience in ways that transcend deductive logic. If that is all that is meant by 'inductive inference' then I do not doubt that such an animal exists; but in this sense of the phrase I would dispute that a Humean argument does apply to all inductive inferences, for the reasons I have given.

In §II I suggested that many philosophers endorse the Salmon–Lipton position because they regard the following as a legitimate generalization of Hume's point: 'Our inductive methods cannot be justified deductively; and they cannot be justified inductively without circularity'. We are now in a position to see where these philosophers go wrong. For this little argument presumes that there is such a thing as an 'inductive method' which we use, and which stands in need of justification. But that is a substantive philosophical assumption, which in no way follows from the mere fact that we hold opinions about the world which are not entailed by our data. It is an

assumption which those philosophers who believe in the irrefutability of Hume's argument have typically failed to call into question.

My suggestion that we use no 'inductive method' that stands in need of justification will no doubt sound strange to some ears. I should stress that I am not trying to cast doubt on the efforts of those philosophers, past and present, who attempt to describe what they call 'scientific method'. For the array of methodological canons and strategies they have unearthed in no way approximates to the ampliative rules whose existence I have disputed, despite all the talk of 'straight-rule induction', 'the hypothetico-deductive method', and the like. Indeed, it is striking how many of the canons that go by the name of 'scientific method' can be explained from a Bayesian point of view. The value of surprising predictions, the desire for evidential variety, the distrust of *ad hoc* hypotheses, are just some of the methodological preferences which can be explained by assuming that scientists are ordinary Bayesian conditionalizers, who start with appropriate prior probabilities.[16] The existence of these preferences, and others like them, lends no credence to the idea that scientists are employing a pattern of inference whose truth-conduciveness depends on which possible world we happen to inhabit.

To conclude, a Humean sceptical argument will only work if our inductive behaviour can be characterized as a process of rule-governed ampliation. There is no necessity that our inductive behaviour can be so characterized. I have offered reasons for thinking that it cannot be. If this is correct, then Hume's argument cannot be converted from a valid one into a sound one, and the threat of inductive scepticism is successfully parried.

V. SOME REFLECTIONS ON THE 'PROBLEM OF INDUCTION'

It is worth briefly comparing my approach to the problem of induction with certain other approaches. Some philosophers have tried to block Hume's argument by arguing that although it is quite true that the use of 'induction' will only be successful if the world is arranged in a certain way, we do not need to have reason to think that the world is like that before reasoning 'inductively' can be rationally sanctioned. This move has obvious affinities with externalism in the theory of epistemic justification. Strawson made a similar move, arguing that since 'induction' is one of the basic standards of rationality, it is true *by definition* that the use of 'induction' is rational, though

[16] See M. Hesse, *The Structure of Scientific Inference* (Univ. of California Press, 1974), P. Horwich, *Probability and Evidence* (Cambridge UP, 1982), Rosenkrantz, *Inference, Method and Decision*, Howson and Urbach, and Earman, for Bayesian explanations of these methodological preferences, and of others like them.

not that it will be successful.[17] But this way with Hume's argument presumes that there is indeed a method called 'induction' we use, whose success is contingent on the world's being arranged in a particular way. So it contains an implicit commitment to the 'inductive rules' picture attacked above. In my view, Hume would be quite right, if we did use inductive rules: Hume's line of thought 'But if you are going to use *that* rule, rather than any other, you need a reason to think the world is like *this*' seems to me entirely convincing. The externalist and Strawsonian approaches both concede to Hume the crucial assumption that we use inductive rules, and then attack his argument at a point where it is perfectly solid. Many philosophers feel that Strawson's response ultimately fails to come to grips with Hume's argument, and that externalism avoids the problem of induction only by shifting the goalposts. These sentiments are correct, in my view, and easily explained by the diagnosis of Humean scepticism offered above.

Other philosophers have tackled the problem of induction not by trying to locate a flaw in Hume's reasoning, but by trying to produce a positive 'justification of induction'. The standard objection to this project is that it is impossible. Would-be justifiers of induction typically defend their project's feasibility by saying that circular justification is actually perfectly acceptable, or that a good justification of induction need not be convincing to those who do not already accept induction. These claims are so counter-intuitive, and so palpably *ad hoc*, that they merely serve to strengthen the aura of unassailability that often surrounds the inductive sceptic's citadel. In my opinion, the project of trying to 'justify induction' is misguided, simply because it assumes that there is some well defined practice called 'induction' which we engage in and which needs justification. This project too involves a commitment to the 'inductive rules' picture, which is often quite explicit. For example, Max Black, a keen advocate of 'circular justification', tried to justify the rule 'Argue from "Most instances of *A*s examined in a variety of conditions have been *B*" to (probably) "The next *A* to be encountered will be *B*"'.[18] But since no one would employ this rule for all choices of *A* and *B*, irrespective of what else they believed, it is hard to avoid the conclusion that Black was wasting his efforts. We simply do not use the inductive rule which he sought to justify. By focusing on the rules of inductive inference which we allegedly employ, the project of 'justifying induction' concedes the crucial assumption which paves the way for Humean scepticism. So it is no surprise that the project's supporters have had to resort to such counter-intuitive claims about epistemic justification to keep their hopes alive.

[17] P.F. Strawson, *Introduction to Logical Theory* (London: Methuen, 1952).

[18] M. Black, 'Self-Supporting Inductive Arguments', in R. Swinburne (ed.), *The Justification of Induction* (Oxford UP, 1974), pp. 127–44, at p. 128.

The approach I have developed bears a certain affinity to Popper's response to the problem of induction.[19] Popper maintained that science does not proceed inductively at all, so Hume's critique, though correct, was simply inapplicable to science as we know it. To many this claim seemed extraordinary – contemporary philosophers are virtually unanimous in their view that Popper was just evading the problem of induction. Of course many of Popper's contemporaries shared this view too. Reviewing the first edition of *The Logic of Scientific Discovery* in *Erkenntnis* for 1935, Reichenbach wrote 'The results of this book seem to me completely untenable; I cannot understand how Popper could possibly believe that with respect to the problem of induction, his investigations mean even the slightest advance'. I think history has judged Popper *slightly* unfairly. For his thesis that 'Science does not use induction' seems to me to contain two elements, one plausible, the other wholly implausible, which most authors, including Popper himself, have failed to keep distinct. On the one hand, Popper castigated those who held that scientists use an 'inductive principle' to arrive at beliefs about the world, pointing out, quite fairly, that believers in this 'inductive principle' have habitually failed to make its content explicit. Popper insisted that scientific theories are the product of *conjecture*, not inference. This aspect of Popper's thought seems to me correct – it is entirely consonant with the view defended above, that we use no rules of inductive inference. On the other hand, Popper argued that scientists spend their time trying to falsify rather than confirm their theories; this was just as well, he thought, since confirmation was impossible, as Hume had allegedly shown. But this part of the idea that 'Science does not use induction' is obviously untenable. When a scientist puts a theory to the test, he is almost always trying to find evidence in favour of his theory, not against it. From a Bayesian point of view, the fact that scientists habitually do what Popper said was impossible is easily accounted for. When a scientist learns that an implication of his theory is true, and learns nothing logically stronger, obeying Bayesian principles will almost always lead him to raise his degree of confidence in the theory. (The point is that if h entails e, then $P(h/e) > P(h)$, so long as neither $P(h)$ nor $P(e)$ equals 0 or 1.) Popper ran together these two quite different aspects of the idea that 'Science does not use induction', because of his refusal to see a distinction between Bayesianism and the 'inductivism' which he abhorred. This was unfortunate, as the implausibility of his thesis that confirmation is impossible has obscured the fact that his attack on the idea that scientists use an 'inductive principle' was basically well taken. Popper's underlying error, though, was simply to set too much store by Hume's sceptical argument, like so many others. Had he appreciated that Hume's argument itself rests on

[19] K.R. Popper, *The Logic of Scientific Discovery* (London: Hutchinson, 1959).

the assumption that empirical enquiry involves 'inductive inference' and 'inductive principles', notions he rightly attacked, twentieth-century philosophy of science might have evolved very differently.

VI. CONCLUSION

I began by noting a tension in the way many philosophers regard Hume's problem. Many philosophers accept that his description of our inductive practices was inadequate, but see the force of his sceptical argument as entirely unaffected, despite the agreed fact that Hume was not simply objecting to our beliefs because they go beyond what the data entail. Exploring this tension led to the question 'If Hume had described our inductive practices "correctly", would his sceptical argument have worked?'. The answer turned out to be 'Not necessarily': Hume's argument rests essentially on the idea that we use inductive rules. If that were true, he would be right. For an inductive rule will only be truth-conducive if the world is arranged in a certain way; to use the rule thus requires having reason to think the world *is* arranged in that way. But that is an empirical assumption, and thus in need of inductive support itself. However, I have argued that in point of fact we do not use any inductive rules, so a Humean argument cannot apply. The answer to the question 'What did Hume really show?' is therefore 'Inductive fallibilism, but no more'. Hume showed that our empirical beliefs cannot be certain; he did not show that they are unreasonable.

It would be foolish to claim the last word on a problem as old and multifaceted as the problem of induction, and I do not. But I hope to have shown that the prospects of subjecting Hume's sceptical argument to a head-on rational response are very much better than has often been thought.[20]

University of York

[20] For comments on and discussion of this material I would like to thank Bill Newton-Smith, Dorothy Edgington, Peter Lipton, Steve Barker, John Worrall, Colin Howson, Alex Rosenberg, Elliott Sober, Carl Hoefer, Scott Campbell, Nancy Cartwright, Atocha Aliseda, Sergio Martinez, Miguel Angel Fernandez, Claudia Lorena Garcia, David Papineau, Darragh Byrne, Pablo Mitnik, James Ladyman and an anonymous referee for this journal. Thanks also to audiences in London, Cambridge, Mexico, Nottingham and Birmingham, where versions of this paper were presented. Research for this paper was carried out while I was the holder of a University of London Jacobsen Fellowship, which I gratefully acknowledge.