

X*—DOES HUME'S ARGUMENT AGAINST INDUCTION REST ON A QUANTIFIER-SHIFT FALLACY?

by Samir Okasha

ABSTRACT It is widely agreed that Hume's *description* of human inductive reasoning is inadequate. But many philosophers think that this inadequacy in no way affects the force of Hume's argument for the *unjustifiability* of inductive reasoning. I argue that this constellation of opinions contains a serious tension, given that Hume was not merely pointing out that induction is fallible. I then explore a recent diagnosis of where Hume's sceptical argument goes wrong, due to Elliott Sober. Sober argues that Hume committed a quantifier-shift fallacy, i.e. inferred a statement of $\exists\forall$ form from one of $\forall\exists$ form. The implications of this diagnosis for the traditional problem of induction are briefly examined.

The author suggests that, were his philosophy received, we must alter from the foundation the greatest part of the sciences.

Hume 1740

I

Introduction. In *Experience and Prediction*, Hans Reichenbach described Hume's sceptical reflections about induction as 'the heaviest blow against empiricism' (1938 p. 347). Despite the advances in our understanding of inductive inference since Reichenbach wrote, the notion that Hume showed something deeply troubling about induction remains widespread. Thus for example, Wesley Salmon (1978) describes Hume's problem as 'unfinished business' for philosophy, while Peter Lipton (1991) says that the Humean argument 'has yet to be overturned', adding that the prospects for overturning it are 'bleak' (p. 12). I will not try to decide whether this pessimistic viewpoint is correct. Rather, I want to explore a recent diagnosis of Hume's case for inductive scepticism due to Elliott Sober (1988). Sober suggests that Hume's argument rests on a quantifier-shift fallacy, i.e. a

*Meeting of the Aristotelian Society, held in Senate House, University of London, on Monday, 7 March, 2005 at 4.15 p.m.

move from a statement of $\forall\exists$ form to one of $\exists\forall$ form. John Norton (2003) defends a related idea.

Before turning to the details of Sober's and Norton's views, I want to draw attention to, and discuss at some length, what strikes me as a puzzling feature of the way many twentieth-century philosophers have reacted to Hume. Reflecting on this feature will help motivate the arguments that follow. First, some exegetical preliminaries.

II

Two Rival Interpretations of Hume. Hume formulated his scepticism about induction extremely succinctly. All of our arguments from experience, he says, are 'founded on the supposition' that the future will resemble the past, i.e. that nature is uniform (1777, p. 32). But if we ask what justifies this supposition, there is no satisfactory answer. There can be no demonstrative argument for nature's uniformity, for 'It implies no contradiction that the course of nature may change' (p. 35). There can be no probable argument for nature's uniformity, for probable arguments are *themselves* arguments from experience and hence presuppose that very uniformity. ('To endeavour, therefore, the proof of this supposition by probable arguments... must evidently be going in a circle, and taking that for granted, which is the very point in question,' p. 36.) Since demonstrative and probable arguments exhaust the options, Hume concludes that the uniformity of nature assumption admits of no rational justification at all. Since our inductive inferences are 'founded' on that assumption, they themselves also lack rational justification.

Despite the apparent simplicity of Hume's argument, debate has raged over how exactly it should be understood. Two rival interpretative camps have emerged. The first camp includes David Stove, author of the only book-length analysis of Hume's argument, J.L. Mackie, and others. The second camp includes Barry Stroud, Wesley Salmon, Elliott Sober, and others. The main point of dispute is whether Hume's argument contains an implicit commitment to 'deductivism'. 'Deductivism' is the thesis that the only rationally compelling arguments are deductively valid ones. Members of the first camp believe that Hume's

argument does rely on 'deductivism', while members of the second camp deny it. 'Deductivism' is not very plausible, and is not a thesis for which Hume offered any explicit defence, so members of the first camp usually take a low view of Hume's sceptical argument; indeed, Stove takes himself to have refuted Hume. By contrast, members of the second camp generally view the Humean argument more favourably.

In effect, the Stove and Stroud camps differ over whether Hume's sceptical worries should trouble an epistemological fallibilist or not. By a fallibilist I mean someone who accepts that our inductive methods are not *guaranteed* to yield the truth, and thus that empirical science cannot attain the type of certainty achievable in mathematics, but regards that as quite compatible with our inductively-based beliefs being justified, or reasonable, or well-confirmed. If Stove *et al.* are right that Hume's argument rests on a deductivist premise, then the most the argument can show is that our inductively-based beliefs cannot be certain. And for a fallibilist, who regards certainty about matters empirical as an impossible ideal, that is an unthreatening conclusion. But if the Stroud interpretation is right, then Hume's argument cannot be seen off so simply. For according to Stroud, Hume was not simply objecting to inductive inferences on the score of their logical invalidity. His point was not just that inductive inferences cannot be *conclusively* justified, but rather that they do not admit of any non-circular justification at all. This point, if correct, is troubling for fallibilist and infallibilist alike.

At the exegetical level, the dispute hinges on what Hume *meant* when he said that arguments from experience, i.e. inductive inferences, are 'founded on the supposition' that nature is uniform. According to Stove, Hume meant that the statement 'Nature is uniform', or 'The future will resemble the past' must be added to the premises of an argument from experience, in order to convert it into a deductively valid one. Stroud's camp disputes this. The phrase 'is founded on the supposition of' needs to be understood epistemically, not logically, Stroud holds. Hume is saying that the premises of an argument from experience only provide reason to believe the conclusion if there is reason to think that nature is uniform, but this is *not* because he thinks all reasons must be deductive. The role of the uniformity principle is not that of additional premise needed to convert

inductive inferences into deductively valid ones; rather, it is that of assumption without which the premises of an inductive inference fail to probabilify the conclusion at all. Hume had no particular interest in deductive validity, Stroud maintains.

I think there is textual evidence in favour of both Stove and Stroud. But on balance, Stroud's view seems preferable. For Stove's interpretation of Hume's claim that inductive inferences are 'founded on the supposition' of nature's uniformity requires imputing to Hume a simple logical error. Consider a typical inductive inference of the sort Hume was interested in, e.g. from 'The sun has risen every day in the past' to 'The sun will rise tomorrow.' If we add in the premise 'Nature is uniform' or 'The future will resemble the past', we may perhaps be able to make the inference deductively valid.¹ But it is not true that this premise is *needed* to convert the argument into a valid one: any number of additional premises could be invented that would do the trick. *No* logically invalid inference has the property that there is only one additional premise which can be added to make it valid; there is always an innumerable number of such premises. If Hume meant that 'Nature is uniform' must be added to the premises of an argument from experience in order to make it deductively valid, as Stove and Mackie allege, it is hard to avoid the conclusion that he was guilty of an elementary logical error.²

I do not want to delve any further into these exegetical matters here. Rather, I will assume that something like the Stroud interpretation is correct, and thus that Hume's inductive scepticism cannot be avoided simply by adopting a fallibilist epistemology. Supporters of Stove should still read on, for not everything I say below depends on assuming the correctness of Stroud's view.

1. Presuming the meaning of the statement 'nature is uniform' can be adequately spelled out, that is.

2. Bizarrely, Stove and Mackie both miss this point entirely: they appear to think it is *true* that the uniformity assumption is needed to convert inferences from observed to unobserved into deductively valid ones. Stove says that when Hume's claim about our inferences 'presupposing' nature's uniformity is interpreted as he recommends, it reads 'Inductive arguments are all invalid as they stand, and it would be necessary, in order to turn them into valid arguments, to add to their premises a proposition which asserts that unobserved instances resemble observed ones' (1973 p. 43). He continues 'And so interpreted, [the claim] is true' (p. 43). Ironically, Stove regards his interpretation of Hume as an exercise in interpretative charity when in fact it makes Hume guilty of a simple logical mistake.

III

A Puzzling Feature. I mentioned above a puzzling feature of the way many twentieth-century philosophers have reacted to Hume. To bring out this puzzling feature, a brief digression is necessary.

Philosophical enquiries into the nature of inductive inference are of two main types (Lipton 1991). First, there are attempts to specify the principles of inductive inference actually used in empirical enquiry (everyday and scientific). This is essentially a descriptive project: it aims to codify and describe the inductive practices that scientists and laymen actually engage in. Secondly, there are attempts to justify the principles of inductive inference used in everyday life and science. This project is normative or justificatory: it tries to show why it is *rational* to use the inductive principles that we do in fact use. Clearly, this is a different and much more ambitious project than the descriptive one.

Some theories of inductive inference incorporate both descriptive and normative elements. Bayesianism is an example. Bayesians have tried to show that actual scientific reasoning does in fact conform to Bayesian principles, and they generally regard this as supplying a vindication of the reasoning in question. The idea is that *by* describing actual inductive practices in a certain way, the rationality of those practices can be exhibited. But even here, the descriptive and normative issues can in principle be disentangled—asking whether scientists *do* update their beliefs by Bayesian conditionalization is distinct from asking whether they *should*. Other theories of induction explicitly restrict themselves to the descriptive project. An example is Clark Glymour's (1980) 'bootstrapping' account of theory confirmation. Glymour attempts to codify a commonly used pattern of non-deductive inference found in the physical sciences. But as regards the justification for employing this pattern of inference, he admits he has nothing to say. Glymour writes: 'There is nothing in this book that corresponds to an attempt to show that the methods I have described are justified or uniquely rational' (p. 377).

As Glymour's work illustrates, it is possible to engage in the descriptive project while leaving the normative project untouched. The converse possibility, by contrast, seems clearly impossible. If we want to try to justify the principles of inductive inference, however exactly that might be achieved, we obviously

need to know what those principles *are*, at least in rough outline. It would be odd indeed to attempt a justification of our 'inductive methods' unless one had a reasonable idea of what those methods were.³ Justificatory projects in philosophy are hard enough as it is; their chances of success must be especially slim if we do not even know what it is we want to justify.

This last point applies not just to positive justificatory projects, which try to show that our inductive methods *are* justified, but also to justificatory projects that have a sceptical intent, i.e. which try to show that our inductive methods are *not* justified. Just as the would-be defender of induction needs to know what it is he wishes to defend, so the would-be inductive sceptic needs to know what it is he wishes to be sceptical about. This may all sound too obvious to be worth stating, but I will argue that a certain, quite standard way of thinking about Hume's problem falls foul of it.

Where does Hume's own discussion of induction stand with respect to the distinction between description and justification? I think Hume was concerned with both. He was certainly interested in asking whether the methods by which we infer from observed to unobserved could be rationally justified; he concluded that they could not. But he was also interested in trying to say what those methods are, i.e. trying to describe how our inductive practices do in fact work. The description he gave was a simple one: we notice regularities in our past experience and then project them into the future. In Hume's words: 'We always presume, when we see like sensible qualities . . . that effects, similar to those which we have experienced, will follow from them, and that the course of nature will continue uniformly the same' (1777 p. 33). So Hume's project contains both descriptive and normative elements. He offers both a description of the principles which do in fact guide our inductive inferences, and a sceptical argument to the effect that those principles cannot be rationally justified.

Leaving aside the question of justification, it is clear that Hume's description is not adequate. There are two basic sources of inadequacy. First, Hume focused exclusively on one very

3. Here and throughout, I use the phrase 'inductive methods' in a deliberately vague way, to mean the methods we use to infer from observed to unobserved, whatever they are.

simple type of inductive inference—extrapolating past regularities into the future. Such inferences are arguably important, in both science and everyday life, but they do not exhaust our repertoire of non-deductive methods. Many philosophers think a complete account of inductive reasoning must also include some form of hypothetico-deduction, or 'inference to the best explanation', and Hume said nothing about this. Secondly and more seriously, Hume's description is inadequate even on its own terms. It simply is not true that we extrapolate *every* regularity in our past experience into the future. As Goodman (1954) famously observed, you can find regularities wherever you please, for any set of objects will resemble one another in innumerable respects. So Hume's description of how our inductive practices work is clearly unsatisfactory as it stands.

This is a very familiar point. But what interests me is the following. Although it is obvious that Hume's description of how we induce was flawed, many philosophers apparently think this has *no bearing at all* on his sceptical argument. Wesley Salmon is an example. He says that 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' (1968 p. 29). Another example is Peter Lipton. He says that Hume's sceptical argument 'seems to work, whatever the details of our [inductive] inferences' (1991 p. 14). So even though Salmon and Lipton agree that Hume got the 'details' wrong, they think this doesn't affect the force of his sceptical argument at all. A Humean sceptical argument will apply, they believe, whatever the correct description of our inductive practices turns out to be.

The Salmon/Lipton view seems to me very widespread in philosophy.⁴ I do not claim that the view is definitely mistaken. But it does sit badly with the Stroud interpretation of Hume's argument. Certainly, if Hume had simply been objecting to inductive inferences on the score of their logical invalidity, as proponents of the *Stove* interpretation hold, then the inadequacy of his

4. Some such view is implicitly held by those who think that the essence of Hume's sceptical argument can be characterised abstractly as follows: 'Our inductive methods cannot be justified deductively, and they cannot be justified inductively without circularity.' Characterisations of this sort abound in the literature, but they should be regarded with suspicion by anyone favouring the Stroud interpretation of Hume.

description would not matter to his sceptical argument. But proponents of the Stroud interpretation agree that Hume was not doing that. Indeed Salmon himself is one such proponent. He writes ‘The suggestion that Hume merely showed the fallibility of induction is a mistake’—which is precisely to say that the Stove interpretation is wrong (1968 p. 30). So Salmon thinks that Hume provided a sceptical argument which (i) applies to *every* deductively invalid inference, hence does not depend on the specific description of our inductive practices with which Hume operated; and (ii) does *more* than merely show the fallibility of induction.

This is a strange combination of views. For clearly, the only property that all deductively invalid inferences have in common is the property of deductive invalidity itself. So a sceptical argument that applies to *all* deductively invalid inferences must presumably be objecting to these inferences *because* they are deductively invalid. But the Stroud interpretation holds that Hume was *not* doing this—he was not a deductivist. So if the Stroud interpretation is correct, and Hume’s argument does more than merely show the fallibility of our inductive methods, then the argument had better depend in some way on what our ‘inductive methods’ actually *are*. Put differently, if it makes no difference to the Humean sceptic what our inductive practices are actually like, beyond the fact that they are inductive, then it is hard to see how the Stroud rather than Stove interpretation can be correct. And if the Stove interpretation is correct, then Hume’s problem disappears as soon as we adopt a fallibilist epistemology, of course.

The puzzling feature is therefore as follows. Many philosophers think that although Hume’s description of our inductive practices was inadequate, this doesn’t affect the force of his sceptical argument. Many philosophers also think, following Stroud, that Hume’s sceptical argument doesn’t just turn on the point that our inductive inferences are deductively invalid. But these two opinions can’t both be right. If Hume’s argument is really invariant with respect to the correct description of our inductive practices, then it must ultimately turn on the fact that those practices are non-deductive—there is nothing else it could turn on. Conversely, if Hume’s argument *doesn’t* turn on that fact, as Stroud insists, then it can’t be invariant with respect to

what our inductive practices are actually like. And in that case, the 'details' of how our inductive inferences work may indeed be relevant to assessing the case for inductive scepticism.⁵

If this is correct, it suggests there is something seriously wrong with the way many philosophers have reacted to Hume. I believe this to be the case. But rather than attempt to substantiate this opinion further, I want now to examine a recent diagnosis of Hume's sceptical argument which does *not* fall foul of the dilemma noted above.

IV

Elliott Sober on Humean Scepticism. Sober (1988) emphasises the abstract form of Hume's sceptical argument. Hume isolates a *single* empirical assumption—that nature is uniform—and claims that all inductive inferences rely on this assumption, in the sense that without it, their premises will not support their conclusions. But the assumption cannot be proved true, since it is contingent, and cannot be inductively justified itself, since all inductive inferences rely on it. Unlike the authors criticised above, Sober realises that this argument depends on Hume's particular conception of how inductive reasoning proceeds. It is *because* Hume thought of induction as a matter of extrapolating observations according to the principle that the future will resemble the past, that he thought we need reason to believe in nature's uniformity, if our inductive inferences are to be rationally justified. The sceptical argument Hume gave is directly dependent on the simplistic model of inductive inference with which he operated.

As I have stressed, real-life inductive inference is much more complicated, and harder to describe, than Hume realised. Sober emphasises a feature of real-life inductive inference that Hume ignored: our inferences take place in the light of our background beliefs, and the evidential significance of our observations depends on the content of those beliefs. Relative to one background theory, a given observation might lead us to make one prediction, but relative to a different background theory it might lead to a different prediction, or to no prediction

5. See Okasha (2001) for a more detailed elaboration of this paragraph's argument.

at all.⁶ Sober stresses that it is these background theories, not the blanket assumption that nature is uniform, that mediate our inferences from observed to unobserved. He writes: ‘When asked to say why we take past observations to support the belief that the sun will rise tomorrow, we answer by citing our theory of planetary motion, not Hume’s principle of the uniformity of nature’ (1988, p. 66). Inferences from observed to unobserved are always dependent on our background theory, Sober claims; observations on their own have no evidential significance.

On the basis of this account of how inductive inference proceeds, Sober offers an attractive diagnosis of where Hume’s sceptical argument goes wrong. Sober *agrees* with Hume that observations alone cannot justify any prediction about the future—something else is needed too. However, where Hume thought that what was needed was the assumption that nature is uniform, Sober holds that what is needed is some background empirical theory or other, and not necessarily the same one for every inductive inference. According to Sober, Hume thus committed a simple quantifier-shift fallacy: ‘Hume’s mistake was to think that since every inductive inference requires assumptions additional to observations, there must be an additional assumption that every inductive inference requires’ (ibid. p. 68). In short, Hume asserted a thesis of the form $\exists\forall$ where all he was entitled to was the corresponding $\forall\exists$ thesis.

Two points about Sober’s argument bear emphasis. First, his point is *not* that additional assumptions are needed to make inferences from observed to unobserved deductively valid; like Hume (interpreted à la Stroud), Sober is not a deductivist. Rather, additional assumptions are needed if observations are to confer any degree of support at all on conclusions about the unobserved. Secondly, these ‘additional assumptions’ are theoretical in character—they are not simply more observations. To take an example, suppose DNA evidence indicates extensive genetic similarities between two species, and we infer from this that the species are closely related. The relevant additional

6. This point is fully consonant with Bayesian confirmation theory. Letting ‘o’ denote our past observations, ‘u’ our prediction about an unobserved case, and K1 and K2 two alternative background theories, it is quite possible that $P(u/o \ \& \ K1) > P(u/K1)$ but $P(u/o \ \& \ K2) < P(u/K2)$, i.e. o supports u conditional on K1 but not conditional on K2.

assumptions might include a theory about how the process of speciation generates genetic differentiation. This theory is obviously not reducible to a mere list of observations, hence is potentially in need of inductive justification itself.

Why does Sober think that inferences from observed to unobserved always require background theory? He is not wholly explicit on this point, but his thought is essentially as follows. Suppose we have a set of observations and we wish to make a prediction about an unobserved case. There is any number of different predictions we could make. (As Hume said, '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?' 1777, p. 29.) So to be justified in making the prediction we do, we need some reason to prefer it over the other candidates. Arguably, such a reason can only stem from an empirical theory about what the world is like. Consider again the DNA example. If someone held very different biological theories, they might take the genetic similarities to show that the species in question are *unrelated*, not closely related. So if our inference from 'genetic similarity' to 'closely related' is to be justified, at the very minimum we need reason to believe that this person's background theory is wrong; and the negation of a theoretical belief is itself theoretical.

Sober's quantifier-shift diagnosis of Hume's argument has much to recommend it, in my view. For one thing, it avoids the difficulty that Salmon, Lipton and others ran into, discussed above. That difficulty, to recall, stemmed from combining the Stroud interpretation with the claim that Hume's inadequate description of our inductive practices has no bearing at all on his sceptical argument. Sober's view avoids the difficulty by rejecting the latter claim (in effect). As Sober notes, his own account of how our inductive inferences work precludes a sceptical argument of Hume's *form* being run. Since different inductive inferences rely on different empirical assumptions, there is no one assumption that cannot be justified by any inductive inference because it is required by every one. Reversing the order of the quantifiers means that Hume's charge of circularity will no longer go through. So the 'details' of our inductive inferences are indeed relevant to assessing the Humean argument, contra Salmon and Lipton.

However, Sober does not think that he has refuted inductive scepticism—Hume was ‘half-right’, he holds. Hume’s fundamental point—that observations *alone* cannot justify conclusions about the unobserved—was correct, according to Sober. Additional assumptions are always required, and these assumptions themselves need inductive justification. So a sceptical problem does arise. Whenever someone makes an inference from observed to unobserved, we can challenge them to justify the additional assumption which underpins their inference. In the biological example above, what justifies the theory about how speciation generates genetic differentiation? Presumably, the answer will be that the theory is itself well-supported by other observations. But then an infinite regress of justification looms. For these further observations will only support the theory in question *modulo* other empirical assumptions. And what justifies *those* assumptions? And so on. Sober writes: ‘To hold true to the principle of empiricism—that beliefs about the future must be justified in terms of present observation *alone*—is, as Hume rightly saw, to be led to straight to scepticism’ (ibid., p. 68).

In effect, Sober is suggesting that although Hume’s actual sceptical argument will not go through, since it is premised on a flawed account of how our inductive inferences work, his sceptical conclusion is correct nonetheless. Hume wrongly thought that all inductive inferences rely on a single empirical assumption, ‘Nature is uniform’, which therefore cannot be justified without circularity. Since, according to Sober, different inductive inferences rely on different empirical assumptions, no such circularity arises, but a regress of justification arises instead. So the opponent of scepticism can take no solace from the failure of Hume’s actual argument: we have simply traded a vicious circularity for an infinite regress.

I think there may actually be more anti-sceptical mileage in Sober’s position than he himself recognises. Before explaining why, I want to look briefly at some of John Norton’s ideas about induction.

V

Norton on Formal versus Material Theories of Induction. In an interesting recent paper, John Norton (2003) outlines what

he calls a 'material' theory of induction, contrasting it with 'formal' theories of induction. Formal theories try to reduce all non-deductive inferences to one or more universal schema, applicable to every subject matter. 'Straight-rule induction', inference to the best explanation, and Glymour's 'bootstrapping' confirmation are all examples of formal theories, in this sense: they all identify 'good' inductive inferences with those exemplifying a certain abstract pattern.⁷ Norton argues persuasively that formal theories are inherently flawed. Inductive inferences, where rational, do not derive their licence from universal schemas, but rather from facts about the world. If an inference from 'Some Xs are Y' to 'All Xs are Y' is warranted, that is *not* because of the form of the inference—many other inferences of the same form are not warranted. Rather, it is empirical facts about the world, concerning the actual Xs and Ys in question, which license the inference. Someone not apprised of those facts would not perform the inference, and would be right not to. Norton writes: 'All inductions ultimately derive their licenses from facts pertinent to the matter of the induction' (p. 650). He calls these facts 'material postulates' of induction.

Though independently arrived at, there is an obvious parallel between Norton's ideas and Sober's.⁸ Interestingly, Norton argues that his material theory of induction has the resources to defeat, or at least 'elude', the Humean sceptic. A version of Hume's problem *does* arise for the material theory, he claims, but it is far less problematic than the traditional problem. Since Norton, like Sober, holds that different inductive inferences derive their licences from different facts about the world, not all from the same fact, the sceptical threat is of a regress, not a vicious circularity. The threat arises in exactly the same way as it did for Sober: inferences from observed to unobserved rely on

7. Norton also identifies Bayesianism as a 'formal' theory in his sense, but this seems to me questionable.

8. One possible difference is that Sober tends towards epistemological internalism, Norton towards externalism. Norton says that empirical facts 'license' inductive inferences whether the inducing scientist knows it or not, implying that the justification for the inference depends only on these facts obtaining; Sober talks about background empirical *theories*, rather than facts, providing the ground for our inductive inferences, implying that the justification depends on what is believed to obtain. Though important, this difference shouldn't obscure the fundamental congruence between Sober's and Norton's positions.

material postulates; these postulates need inductive justification themselves, hence the need for further material postulates, and so-on. But unlike Sober, Norton regards this regress as relatively benign—it is ‘neither infinite nor demonstrably harmful’, he claims (p. 666).

Norton’s attempt to show that the regress is not infinite is not especially convincing. He says that some chains of justification ‘may just terminate in brute facts of experience that do not need further justification’ (p. 668). But if this is so, it is hard to see what motivates the thesis that *every* inductive inference needs to be licensed by some material postulate or other; indeed, the claim that ‘brute facts of experience’ may terminate the regress of justification is tantamount to a denial of that thesis. It cannot be true that inferences from observed to unobserved always require additional empirical assumptions *and* that ‘brute facts of experience’ can justify conclusions about the unobserved on their own.

However, Norton does raise one interesting point about the analogue of the Humean argument that arises for his material theory. He stresses that the demands for justification that arise, when inductive inference is thought of materially, are of a quite ordinary character; they are demands that arise in empirical enquiry all the time. ‘As we trace back the justifications of justifications of inductions, we are simply engaged in the repeated exercise of displaying the reasons for why we believe this or that fact within our sciences’ (p. 668). This contrasts sharply with the justificatory demand that arises on a formal picture of inductive inference. Here, the demand is for the justification of some grand overarching fact about the world, e.g. that nature is uniform, which would license all inferences satisfying a certain schema. Whether or not such justification could be given, it is certainly true that scientists do not spend their time seeking, or demanding, such justifications. I return to this point below.

VI

Consequences. Let us assume that something like the Sober/Norton $\forall\exists$ picture of inductive inference is correct, and thus that the quantifier-shift diagnosis of Hume’s sceptical argument is correct too. What philosophical consequences follow?

We have touched on one consequence already: the Salmon/Lipton claim that Hume's argument doesn't depend on the 'details' of how our inductive inferences work is shown to be mistaken. The details *do* matter. Even if Sober is right that Hume's sceptical conclusion still stands, the actual argument Hume gave for that conclusion does not go though.

Another consequence, I suggest, is this. When we reverse the order of the quantifiers, inductive reasoning comes to seem less puzzling, and empirical enquiry less dogmatic, than on the Humean $\exists\forall$ view, for the following reason. On the Humean view, all empirical enquiry rests on a single overarching assumption about how the world is, namely that nature is uniform. Now as Hume is at pains to point out, there is no necessity that this assumption holds true—in many possible worlds it doesn't. So it looks as if scientists are blithely assuming that the actual world is not one of these worlds, every time they reason inductively. But this looks like a most dogmatic procedure. Empiricists are not supposed to make assumptions about the world and hold onto them at all costs, whatever the evidence—they should always be ready to revise their worldly beliefs. But on the Humean picture, the uniformity of nature is a presupposition of all empirical enquiry, hence not something that that could itself be subject to empirical confirmation or disconfirmation. Thus the whole practice of reasoning inductively seems to be premised on an enormous, untestable assumption about the way the world is. So when the sceptic asks, 'But *why* do you assume the world is that way?', and fails to receive a satisfactory answer, she seems to have highlighted a crippling difficulty at the heart of the scientific enterprise.

On the alternative $\forall\exists$ picture, however, induction doesn't seem quite so baffling. Granted, inferences from observed to unobserved always require further empirical assumptions, which themselves require inductive justification. But since different inferences require different assumptions, it is not true that the whole scientific enterprise rests on a single overarching assumption about how the world is constituted, which must be accepted as an article of faith before empirical enquiry can get started. Any given background assumption might eventually be overthrown, and the inductive inferences that it licensed would then have to be abandoned. So whether the chain of

justification comes to end, as Norton holds, or regresses to infinity, as Sober holds, *in either case* inductive reasoning is more comprehensible than on the Humean picture. It does not require a fixed commitment to the world being one way rather than another, of the sort that should be anathema to an empiricist, but rather a series of empirical assumptions that are provisional, not fixed, and can be revised in the course of enquiry.

Hume is not the only author to defend an $\exists\forall$ picture of induction. Other examples include Mill, Russell and Keynes. Each of these authors thought they had identified one or a few empirical presuppositions on which all inductive inference rests. Mill (1872), like Hume, talked of an ‘axiom of the uniformity of the course of nature’, Russell (1912) talked of four ‘postulates of human knowledge’ while Keynes (1921) talked of the ‘law of independent variety’. Ironically, these presuppositions were normally adumbrated in the course of *defending* induction’s rationality. But in fact, they permit the inductive sceptic’s case to be brought in a particularly damaging way. For as with Hume, they immediately invite the challenge ‘But how do you know these presuppositions hold true?’. Since this question cannot be answered satisfactorily, the rationality of induction is immediately called into question. On the $\forall\exists$ picture, by contrast, inductive reasoning relies no overarching presumptions of this sort, so the impossibility of justifying them does not matter.

This suggests that the $\forall\exists$ picture may offer more solace for the anti-sceptic than Sober allows. The *form* which the inductive sceptic’s argument takes on the $\forall\exists$ picture—pushing the demand for justification further and further back—seems somehow less problematic than in the $\exists\forall$ case. As Norton stresses, such demands are a real feature of empirical enquiry, and are met by explaining why it is we believe this or that empirical fact. But on the $\exists\forall$ picture, the sceptic’s argument assumes a far more pernicious form. If all our inductive inferences really did rest on a single overarching assumption about the world, the sceptic would surely be right to demand some justification for it. And the impossibility of supplying such a justification then makes science look deeply irrational, reliant on a blind faith that the world is one way rather than another. Even if the chains of justification that arise on the $\forall\exists$ picture do not terminate benignly the way Norton hopes, the resulting sceptical threat

still seems less serious. Perhaps infinite regresses are less bad than vicious circles after all.

A final consequence of reversing the order of the quantifiers concerns the project of 'justifying induction', a project with a long though dubious philosophical pedigree. Would-be justifiers of induction, including Reichenbach, Russell and Max Black, have usually proceeded by identifying the 'fundamental inductive principle' on which all inductive reasoning rests, and then trying somehow to justify it. (Thus for example Black tried to justify the rule 'Argue from "Most instances of A examined in a wide variety of conditions have been B" to "Probably the next A to be examined will be B"': Black 1974, p. 128.) The standard objection to this project is that it is impossible. But if the $\forall\exists$ conception of induction is correct, there is an even more fundamental problem. The principles in question are *not* universal presuppositions of all inductive reasoning—different inductive inferences presuppose different things about the world. *No-one* would use Black's rule for all choices of A and B, irrespective of what else they believed. So there is no point trying to justify this rule, nor worrying whether the proffered justification is any good, for we simply do not use the rule. Like Hume, these authors have premised their discussion on an excessively simplistic account of how our inductive practices actually work.

This point is relevant to the currently fashionable view that induction can be justified so long as we adopt an externalist theory of justification (Papineau 1993, van Cleve 1984). Advocates of this view claim that reasoning inductively is justified so long as it is in fact reliable, whether or not anyone knows this. But they invariably work with an $\exists\forall$ conception of induction. They claim reliability for some abstractly characterised inference pattern, which allegedly subsumes all the actual inductive inferences we make. But if the Sober/Norton $\forall\exists$ view is right, then our actual inductive practices cannot be characterised in this way. There simply is no abstract inference schema which all inductive inferences fit, for our inferences depend so heavily on background beliefs about the world. So the reliabilist approach to the problem of induction looks unpromising even if we accept the externalist starting point, for it is unclear how even to identify the inference rule for which reliability is to be claimed.

VII

Conclusion. I began by emphasising a peculiarity in the way many philosophers have responded to Hume: accepting the Stroud interpretation while simultaneously maintaining that Hume's sceptical argument works, whatever the details of our inductive inferences. The quantifier-shift diagnosis shows this conjunction of views to be wrong: the details do matter. Given the Sober/Norton account of inductive inference, no sceptical argument of Hume's form will work. Reversing the order of the quantifiers may have other significant implications too, as we have seen. A final point to note is this. Both the Humean inductive sceptic *and* those anti-sceptics who seek somehow to 'justify induction' need to operate with a plausible account of how our inductive inferences actually work. Without such an account, claims about induction's justification, whether sceptical or anti-sceptical, cannot sensibly be made. It is striking how many philosophical discussions have fallen foul of this simple methodological moral.

*Department of Philosophy
University of Bristol
9 Woodland Road
Bristol, BS8 1TB
samir.okasha@bristol.ac.uk*

REFERENCES

- Black, M. (1974) 'Self-Supporting Inductive Arguments', in R. Swinburne (ed.) *The Justification of Induction*, Oxford: Oxford University Press.
- Glymour, C. (1980) *Theory and Evidence*, Princeton: Princeton University Press.
- Goodman, N. (1954) *Fact, Fiction and Forecast*, Cambridge MA: Harvard University Press.
- Hume, D. (1740) *A Treatise of Human Nature*, L.A. Selby-Bigge (ed.) (1978) Oxford: Clarendon Press.
- Hume, D. (1777) *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, L.A. Selby-Bigge (ed.) (1975) Oxford: Clarendon Press.
- Keynes, J.M. (1921) *A Treatise on Probability*, London: Macmillan.
- Lipton, P. (1991) *Inference to the Best Explanation*, New York: Routledge.
- Mackie, J.L. (1974) *The Cement of the Universe*, Oxford: Clarendon Press.
- Mill, J.S. (1972) *A System of Logic*, London: Longman.
- Norton, J. (2003) 'A Material Theory of Induction', *Philosophy of Science* 70, 647-670.
- Okasha, S. (2001) 'What did Hume Really Show about Induction?', *Philosophical Quarterly*, 51, 307-328.
- Papineau, D. (1993) *Philosophical Naturalism*, Oxford: Oxford University Press.

- Reichenbach, H. (1938) *Experience and Prediction*, Chicago: University of Chicago Press.
- Russell, B. (1948) *Human Knowledge: Its Scope and Limits*, New York: Simon and Schuster.
- Salmon, W. (1968) 'The Justification of Inductive Rules of Inference', in I. Lakatos (ed.) *The Problem of Inductive Logic*, Amsterdam: North-Holland.
- Salmon, W. (1978) 'Unfinished Business: the Problem of Induction', *Philosophical Studies*, 33, 1–19.
- Sober, E. (1988) *Reconstructing the Past*, Cambridge MA: MIT Press.
- Stove, D. (1973) *Probability and Hume's Inductive Scepticism*, Oxford: Clarendon Press.
- Stroud, B. (1977) *Hume*, London: Routledge.
- van Cleve, J. (1984) 'Reliability, Justification and the Problem of Induction', *Midwest Studies in Philosophy IX*, Minneapolis: University of Minnesota Press.