HUME AND THE PROBLEM OF INDUCTION

Marc Lange

1 INTRODUCTION

David Hume first posed what is now commonly called "the problem of induction" (or simply "Hume's problem") in 1739 — in Book 1, Part iii, section 6 ("Of the inference from the impression to the idea") of *A Treatise of Human Nature* (hereafter T). In 1748, he gave a pithier formulation of the argument in Section iv ("Skeptical doubts concerning the operations of the understanding") of *An Enquiry Concerning Human Understanding* (E).[1] Today Hume's simple but powerful argument has attained the status of a philosophical classic. It is a staple of introductory philosophy courses, annually persuading scores of students of either the enlightening or the corrosive effect of philosophical inquiry – since the argument appears to undermine the credentials of virtually everything that passes for knowledge in their other classes (mathematics notably excepted[2]).

According to the standard interpretation, Hume's argument purports to show that our opinions regarding what we have not observed have no justification. The obstacle is irremediable; no matter how many further observations we might make, we would still not be entitled to any opinions regarding what we have not observed. Hume's point is not the relatively tame conclusion that we are not warranted in making any predictions *with total certainty.* Hume's conclusion is more radical: that we are not entitled to *any degree of confidence whatever*, no matter how slight, in any predictions regarding what we have not observed. We are not justified in having 90% confidence that the sun will rise tomorrow, or in having 70% confidence, or even in being more confident that it will rise than that it will not. There is no opinion (i.e., no degree of confidence) that we are entitled to have regarding a claim concerning what we have not observed. This conclusion "leaves not the lowest degree of evidence in any proposition" that goes beyond our present observations and memory (T, p. 267). Our justified opinions must be "limited to the narrow sphere of our memory and senses" (E, p. 36).

[1]All page references to the *Treatise* are to [Hume, 1978]. All page references to the *Enquiry* are to [Hume, 1977].

[2]However, even in mathematics, inductive logic is used, as when we take the fact that a computer search program has found no violation of Goldbach's conjecture up to some enormously high number as evidence that Goldbach's conjecture is true even for higher numbers. For more examples, see [Franklin, 1987]. Of course, such examples of inductive logic in mathematics must be sharply distinguished from "mathematical induction", which is a form of deductive reasoning.

Handbook of the History of Logic. Volume 10: Inductive Logic. Volume editors: Dov Gabbay, Stephan Hartmann and John Woods.
General Editors: Dov M. Gabbay, Paul Thagard and John Woods.

Hume's problem has not gained its notoriety merely from Hume's boldness in denying the epistemic credentials of all of the proudest products of science (and many of the humblest products of common-sense). It takes nothing for someone simply to declare himself unpersuaded by the evidence offered for some prediction. Hume's problem derives its power from the strength of Hume's *argument* that it is impossible to justify reposing even a modest degree of confidence in any of our predictions. Again, it would be relatively unimpressive to argue that since a variety of past attempts to justify inductive reasoning have failed, there is presumably no way to justify induction and hence, it seems, no warrant for the conclusions that we have called upon induction to support. But Hume's argument is much more ambitious. Hume purports not merely to show that various, apparently promising routes to justifying induction all turn out to fail, but also to exclude every *possible* route to justifying induction.

Naturally, many philosophers have tried to find a way around Hume's argument — to show that science and common-sense are justified in making predictions inductively. Despite these massive efforts, no response to date has received widespread acceptance. Inductive reasoning remains (in C.D. Broad's famous apothegm) "the glory of Science" and "the scandal of Philosophy" [Broad, 1952, p. 143].

Some philosophers have instead embraced Hume's conclusion but tried to characterize science so that it does not involve our placing various degrees of confidence in various predictions. For example, Karl Popper has suggested that although science refutes general hypotheses by finding them to be logically inconsistent with our observations, science never confirms (even to the smallest degree) the predictive accuracy of a general hypothesis. Science has us make guesses regarding what we have not observed by using those general hypotheses that have survived the most potential refutations despite sticking their necks out furthest, and we make these guesses even though we have no good reason to repose any confidence in their truth:

> I think that we shall have to get accustomed to the idea that we must not look upon science as a 'body of knowledge,' but rather as a system of hypotheses; that is to say, a system of guesses or anticipations which in principle cannot be justified, but with which we work as long as they stand up to tests, and of which we are never justified in saying that we know that they are 'true' or 'more or less certain' or even 'probable'. [Popper, 1959, p. 317; cf. Popper, 1972]

However, if we are not justified in having any confidence in a prediction's truth, then it is difficult to see how it could be rational for us to rely upon that prediction [Salmon, 1981]. Admittedly, "that we cannot give a justification ... for our guesses does not mean that we may not have guessed the truth." [Popper, 1972, p. 30] But if we have no good reason to be confident that we have guessed the truth, then we would seem no better justified in being guided by the predictions of theories that have passed their tests than in the predictions of theories that have failed their tests. There would seem to be no grounds for calling our guesswork "rational", as Popper does.

Furthermore, Popper's interpretation of science seems inadequate. Some philosophers, such as van Fraassen [1981; 1989], have denied that science confirms the truth of theories about unobservable entities (such as electrons and electric fields), the truth of hypotheses about the laws of nature, or the truth of counterfactual conditionals (which concern what would have happened under circumstances that actually never came to pass — for example, "Had I struck the match, it would have lit"). But these philosophers have argued that these pursuits fall outside of science because we need none of them in order to confirm the empirical adequacy of various theories, a pursuit that is essential to science. So even these interpretations of science are not nearly as austere as Popper's, according to which science fails to accumulate evidence for empirical predictions.

In this essay, I will devote sections 2, 3, and 4 to explaining Hume's argument and offering some criticism of it. In section 6, I will look at the conclusion that Hume himself draws from it. In sections 5 and 7-11, I will review critically a few of the philosophical responses to Hume that are most lively today.[3]

2 TWO PROBLEMS OF INDUCTION

Although Hume never uses the term "induction" to characterize his topic, today Hume's argument is generally presented as targeting inductive reasoning: any of the kinds of reasoning that we ordinarily take as justifying our opinions regarding what we have not observed. Since Hume's argument exploits the differences between induction and deduction, let's review them.

For the premises of a good deductive argument to be true, but its conclusion to be false, would involve a contradiction. (In philosophical jargon, a good deductive argument is "valid".) For example, a geometric proof is deductive since the truth of its premises ensures the truth of its conclusion by a maximally strong (i.e., "logical") guarantee: *on pain of contradiction*! That deduction reflects the demands of non-contradiction (a semantic point) has a metaphysical consequence — in particular, a consequence having to do with necessity and possibility. A contradiction could not come to pass; it is impossible. So it is impossible for the premises of a good deductive argument to be true but its conclusion to be false. (That is why deduction's "guarantee" is maximally strong.) It is impossible for a good deductive argument to take us from a truth to a falsehood (i.e., to fail to be "truth-preserving") because such failure would involve a contradiction and contradictions are impossible. A good deductive argument is *necessarily* truth-preserving.

In contrast, no contradiction is involved in the premises of a good inductive argument being true and its conclusion being false. (Indeed, as we all know, this

[3]My critical review is hardly exhaustive. For an admirable discussion of some responses to Hume in older literature that I neglect, see [Salmon, 1967].

sort of thing is a familiar fact of life; our expectations, though justly arrived at by reasoning inductively from our observations, sometimes fail to be met.) For example, no matter how many human cells we have examined and found to contain proteins, there would be no contradiction between our evidence and a given as yet unobserved human cell containing no proteins. No contradiction is involved in a good inductive argument's failure to be truth-preserving. Once again, this semantic point has a metaphysical consequence if every necessary truth is such that its falsehood involves a contradiction (at least implicitly): even if a given inductive argument is in fact truth-preserving, it could have failed to be. It is not *necessarily* truth-preserving.[4]

These differences between deduction and induction lead to many other differences. For example, the goodness of a deductive argument does not come in degrees; all deductive arguments are equally (and maximally) strong. In contrast, some inductive arguments are more powerful than others. Our evidence regarding the presence of oxygen in a room that we are about to enter is much stronger than our evidence regarding the presence of oxygen in the atmosphere of a distant planet, though the latter evidence may still be weighty. As we examine more (and more diverse) human cells and find proteins in each, we are entitled to greater confidence that a given unobserved human cell also contains proteins; the inductive argument grows stronger. Furthermore, since the premises of a good deductive argument suffice to ensure its conclusion on pain of contradiction, any addition to those premises is still enough to ensure the conclusion on pain of contradiction. In contrast, by adding to the premises of a good inductive argument, its strength may be augmented or diminished. By adding to our stock of evidence the discovery of one human cell that lacks proteins, for example, we may reduce the strength of our inductive argument for the prediction of proteins in a given unobserved human cell.

That inductive arguments are not deductive — that they are not logically guaranteed to be truth-preserving — plays an important part in Hume's argument (as we shall see in a moment). But the fact that the premises of a good inductive argument cannot give the same maximal guarantee to its conclusion as the premises of a deductive argument give to its conclusion should not by itself be enough to cast doubt on the cogency of inductive reasoning. That the premises of an inductive argument fail to "demonstrate" the truth of its conclusion (i.e., to show that the conclusion could not possibly be false, given the premises) does not show that its premises fail to *confirm* the truth of its conclusion — to warrant us (if

[4]Sometimes it is said that since the conclusion of a good deductive argument is true given the premises on pain of contradiction, the conclusion is implicitly contained in the premises. A good deductive argument is not "ampliative". It may make explicit something that was already implicit in the premises, and so we may learn things through deduction, but a deductive argument does not yield conclusions that "go beyond" its premises. In contrast, a good inductive argument is ampliative; it allows us to "go beyond" the evidence in its premises. This "going beyond" is a metaphor that can be cashed out either semantically (the contrary of an inductive argument's conclusion does not contradict its premises) or metaphysically (it is possible for the conclusion to be false and the premises true).

our belief in the premises is justified) in placing greater confidence (perhaps even great confidence) in the conclusion. (Recall that Hume purports to show that even modest confidence in the conclusions reached by induction is not justified.) That the conclusion of a given inductive argument can be false, though its premises are true, does not show that its premises fail to make its conclusion highly plausible.

In short, inductive arguments take risks in going beyond our observations. Of course, we all know that some risks are justified, whereas others are unwarranted. The mere fact that inductive arguments take risks does not automatically show that the risks they take are unreasonable. But Hume (as it is standard to interpret him) purports to show that we are not justified in taking the risks that inductive inferences demand. That induction is fallible does not show that inductive risks cannot be justified. To show that, we need Hume's *argument*. It aims to show that any scheme purporting to justify taking those risks must fail.

It is important to distinguish two questions that could be asked about the risks we take in having opinions that go beyond the relatively secure ground of what we observe:

1. Why are we justified in going beyond our observations at all?
2. Why are we justified in going beyond our observations in a certain specific way: by having the opinions endorsed by inductive reasoning?

To justify our opinions regarding what we have not observed, it would not suffice merely to justify having some opinions about the unobserved rather than none at all. There are many ways in which we could go beyond our observations. But we believe that only certain ways of doing so are warranted. Our rationale for taking risks must be *selective*: it must reveal that certain risks are worthy of being taken whereas others are unjustified [Salmon, 1967, p. 47]. In other words, an adequate justification of induction must justify *induction* specifically; it must not apply equally well to all schemes, however arbitrary or cockeyed, for going beyond our observations. For example, an adequate justification of induction should tell us that as we examine more (and more diverse) human cells and find proteins in each, we are entitled (typically) to greater confidence that a given unobserved human cell also contains proteins, but not instead to lesser confidence in this prediction — and also not to greater confidence that a given unobserved banana is ripe.

In short, an answer to the first of the two questions above that does not privilege induction, but merely supports our taking some risks rather than none at all, fails to answer the second question. An adequate justification of induction must favor science over guesswork, wishful thinking, necromancy, or superstition; it cannot place them on a par.

3 HUME'S FORK: THE FIRST OPTION

Consider any inductive argument. Its premises contain the reports of our observations. Its conclusion concerns something unobserved. It may be a prediction regarding a particular unobserved case (e.g., that a given human cell contains proteins), a generalization concerning all unobserved cases of a certain kind (that all unobserved human cells contain proteins), or a generalization spanning all cases observed and unobserved (that all human cells contain proteins) — or even something stronger (that it is a law of nature that all human cells contain proteins). Although Hume's argument is not limited to relations of cause and effect, Hume typically gives examples in which we observe a cause (such as my eating bread) and draw upon our past experiences of events that have accompanied similar causes (our having always derived nourishment after eating bread in the past) to confirm that a similar event (my deriving nourishment) will occur in this case. Another Hume favorite involves examples in which a body's presenting a certain sensory appearance (such as the appearance of bread) and our past experiences confirm that the body possesses a certain disposition (a "secret power," such as to nourish when eaten). How can the premises of this inductive argument justify its conclusion?

Hume says that in order for the premises to justify the conclusion, we must be able to reason from the premises to the conclusion in one of two ways:

> All reasonings may be divided into two kinds, namely demonstrative reasoning, or that concerning relations of ideas, and moral reasoning, or that concerning matter of fact and existence. (*E*, p. 22)

By "demonstrative reasoning", Hume seems to mean deduction. As we have seen, deduction concerns "relations of ideas" in that a deductive argument turns on semantic relations: certain ideas contradicting others. Then "moral reasoning, or that concerning matter of fact and existence" would apparently have to be induction. ("Moral reasoning", in the archaic sense that Hume uses here, does not refer specifically to reasoning about right and wrong; "moral reasoning" could, in the strongest cases, supply "moral certainty", a degree of confidence beyond any reasonable doubt but short of the "metaphysical certainty" that a proof supplies.[5]) I will defer the half of Hume's argument concerned with non-demonstrative reasoning until the next section. indexmoral reasoning

Is there a deductive argument taking us from the premises of our inductive argument about bread (and only those premises) to the argument's conclusion? We cannot think of one. But this does not show conclusively that there isn't one. As we all know from laboring over proofs to complete our homework assignments for high-school geometry classes, we sometimes fail to see how a given conclusion can be deduced from certain premises even when there is actually a way to do it. But Hume argues that even if we used greater ingenuity, we could not find a way to reason deductively from an inductive argument's premises to its conclusion. No way exists. Here is a reconstruction of Hume's argument.

> If the conclusion of an inductive argument could be deduced from its premises, then the falsehood of the conclusion would contradict the

[5]For other examples of this usage, see the seventh definition of the adjective "moral" in *The Oxford English Dictionary*.

> truth of the premises.
>
> But the falsehood of its conclusion does not contradict the truth of its premises.
>
> So the conclusion of an inductive argument cannot be deduced from its premises.

How does Hume know that the conclusion's falsehood does not contradict the truth of the argument's premises? Hume says that we can form a clear idea of the conclusion's being false along with the premises being true, and so this state of affairs must involve no contradiction. Here is the argument in some of Hume's words:

> The bread, which I formerly eat, nourished me; that is, a body of such sensible qualities, was, at that time, endued with such secret powers: But does it follow, that other bread must also nourish me at another time, and that like sensible qualities must always be attended with like secret powers? The consequence seems nowise necessary. ... That there are no demonstrative arguments in the case, seems evident; since it implies no contradiction, that the course of nature may change, and that an object, seemingly like those which we have experienced, may be attended with different or contrary effects. May I not clearly and distinctly conceive, that a body, falling from the clouds, and which, in all other respects, resembles snow, has yet the taste of salt or the feeling of fire? ... Now whatever is intelligible, and can be distinctly conceived, implies no contradiction, and can never be proved false by any demonstrative argument or abstract reasoning a priori. (*E*, pp. 21-2)

This passage takes us to another way to understand "Hume's fork": the choice he offers us between two different kinds of reasoning for taking us from an inductive argument's premises to its conclusion. We may interpret "demonstrative reasoning" as reasoning *a priori* (that is, reasoning where the step from premises to conclusion makes no appeal to what we learn from observation) and reasoning "concerning matter of fact and existence" as reasoning empirically (that is, where the step from premises to conclusion depends on observation). According to Hume, we can reason *a priori* from p to q only if "If p, then q" is necessary — i.e., only if it could not have been false. That is because if it could have been false, then in order to know that it is true, we must check the actual world — that is, make some observations. If "If p, then q" is not a necessity but merely happens to hold (i.e., holds as a "matter of fact"), then we must consult observations in order to know that it is the case. So reasoning "concerning matter of fact and existence" must be empirical.

Now Hume can argue once again that by reasoning *a priori*, we cannot infer from an inductive argument's premises to its conclusion. Here is a reconstruction:

> If we could know *a priori* that the conclusion of an inductive argument is true if its premises are true, then it would have to be necessary for the conclusion to be true if the premises are true.
>
> But it is not necessary for the conclusion to be true if the premises are true.
>
> So we cannot know *a priori* that the conclusion of an inductive argument is true if its premises are true.

Once again, Hume defends the middle step on the grounds that we can clearly conceive of the conclusion being false while the premises are true. Hence, there is no contradiction in the conclusion's being false while the premises are true, and so it is not necessary for the conclusion to be true if the premises are true. Hume says:

> To form a clear idea of any thing, is an undeniable argument for its possibility, and is alone a refutation of any pretended demonstration against it. (*T*, p. 89; cf. *T*, pp. 233, 250)

Of course, if the premises of our inductive argument included not just that bread nourished us on every past occasion when we ate some, but also that all bread is alike in nutritional value, then there would be an *a priori* argument from the premises to the conclusion. It would be a contradiction for all bread to be nutritionally alike, certain slices of bread to be nutritious, but other slices not to be nutritious. However, Hume would ask, how could we know that all bread is alike in nutritional value? That premise (unlike the others) has not been observed to be true. It cannot be inferred *a priori* to be true, given our observations, since its negation involves no contradiction with our observations.

Earlier I said that Hume aims to show that we are not entitled even to the smallest particle of confidence in our predictions about (contingent features of) what we have not observed. But the arguments I have just attributed to Hume are directed against conclusions of the form "p obtains", not "p is likely to obtain" or "p is more likely to obtain than not to obtain". How does Hume's argument generalize to cover these conclusions? How does it generalize to cover opinions short of full belief — opinions involving a degree of confidence less than certainty?

The appropriate way to extend Hume's reasoning depends on what it is to have a degree of belief that falls short of certainty. Having such a degree of belief in p might be interpreted as equivalent to (or at least as associated with) having a full belief that p has a given objective chance of turning out to be true, as when our being twice as confident that a die will land on 1, 2, 3, or 4 than that it will land on 5 or 6 is associated with our believing that the die has twice the objective chance of landing on 1, 2, 3, or 4 than on 5 or 6. As Hume says,

> There is certainly a probability, which arises from a superiority of chances on any side; and accordingly as this superiority increases, and

> surpasses the opposite chances, the probability receives a proportionable increase, and begets still a higher degree of belief or assent to that side, in which we discover the superiority. If a die were marked with one figure or number of spots on four sides, and with another figure or number of spots on the two remaining sides, it would be more probable, that the former would turn up than the latter; though, if it had a thousand sides marked in the same manner, and only one side different, the probability would be much higher, and our belief or expectation of the event more steady and secure. (*E*, p. 37; cf. *T*, p. 127)

Suppose, then, that our having $n\%$ confidence in p must be accompanied by our believing that p has $n\%$ chance of obtaining. Then Hume could argue that since there is no contradiction in the premises of an inductive argument being true even while its conclusion lacks $n\%$ chance of obtaining — for any n, no matter how low — we cannot proceed *a priori* from an inductive argument's premises to even a modest degree of confidence in its conclusion. For example, there is no contradiction in a die's having landed on 1, 2, 3, or 4 twice as often as it has landed on 5 or 6 in the many tosses that we have already observed, but its not having twice the chance of landing on 1, 2, 3, or 4 as on 5 or 6; the die could even be strongly biased (by virtue of its mass distribution) toward landing on 5 or 6, but nevertheless have happened "by chance" to land twice as often on 1, 2, 3, or 4 as on 5 or 6 in the tosses that we have observed.

Hume sometimes seems simply to identify our having n% confidence in p with our believing that p has $n\%$ chance of obtaining. For example, he considers this plausible idea:

> Shou'd it be said, that tho' in an opposition of chances 'tis impossible to determine *with certainty*, on which side the event will fall, yet we can pronounce with certainty, that 'tis more likely and probable, 'twill be on that side where there is a superior number of chances, than where there is an inferior: ... (*T* , p. 127)

Though we are not justified in having complete confidence in a prediction (e.g., that the die's next toss will land on 1, 2, 3, or 4), we are entitled to a more modest degree of belief in it. (One paragraph later, he characterizes confidence in terms of "degrees of stability and assurance".) He continues:

> Shou'd this be said, I wou'd ask, what is here meant by *likelihood and probability?* The likelihood and probability of chances is a superior number of equal chances; and consequently when we say 'tis likely the event will fall on the side, which is superior, rather than on the inferior, we do no more than affirm, that where there is a superior number of chances there is actually a superior, and where there is an inferior there is an inferior; which are identical propositions, and of no consequence. (*T*, p. 127)

In other words, to have a greater degree of confidence that the die's next toss will land on 1, 2, 3, or 4 than that it will land on 5 or 6 is nothing more than to believe that the former has a greater chance than the latter.

However, it is not the case that our having $n\%$ confidence that p must be accompanied by our believing that p has $n\%$ chance of obtaining. Though the outcomes of die tosses may be governed by objective chances (dice, after all, is a "game of chance"), some of our predictions concern facts that we believe involve no objective chances at all, and we often have non-extremal degrees of confidence in those predictions. For instance, I may have 99% confidence that the next slice of bread I eat will be nutritious, but I do not believe that there is some hidden die-toss, radioactive atomic decay, or other objectively chancy process responsible for its nutritional value. For that matter, I may have 90% confidence that the dinosaurs's extinction was preceded by an asteroid collision with the earth (or that the field equations of general relativity are laws of nature), but the objective chance right now that an asteroid collided with the earth before the dinosaurs' extinction is 1 if it did or 0 if it did not (and likewise for the general-relativistic field equations).

Suppose that our having $n\%$ confidence that the next slice of bread I eat will be nutritious need not be accompanied by any prediction about which we must have full belief — such as that the next slice of bread has $n\%$ objective chance of being nutritious, or that $n\%$ of all unobserved slices of bread are nutritious, or that n equals the limiting relative frequency of nutritious slices among all bread slices. Since there is no prediction q about which we must have full belief, Hume cannot show that there is no *a priori* argument from our inductive argument's premises to $n\%$ confidence that the next slice of bread I eat will be nutritious by showing that there is no contradiction in those premises being true while q is false.

We have here a significant gap in Hume's argument [Mackie, 1980, pp. 15-16; Stove, 1965]. If our degrees of belief are personal (i.e., subjective) probabilities rather than claims about the world, then there is no sense in which the truth of an inductive argument's premises fail to contradict the falsehood of its conclusion — since there is no sense in which its conclusion can be false (or true), since its conclusion is a degree of belief, not a claim about the world. (Of course, the conclusion involves a degree of belief in the truth of some claim about the world. But the degree of belief itself is neither true nor false.) Hence, Hume cannot conclude from such non-contradiction that there is no *a priori* argument from the inductive argument's premises to its conclusion. Of course, no *a priori* argument could demonstrate that the premises' truth logically guarantees the conclusion's truth — since, once again, the conclusion is not the kind of thing that could be true (or false). But there could still be an *a priori* argument from the opinions that constitute the inductive argument's premises to the degrees of belief that constitute its conclusion — an argument showing that holding the former opinions requires holding the latter, on pain of irrationality.

This *a priori* argument could not turn entirely on semantic relations because a degree of belief is not the sort of thing that can be true or false, so it cannot be that one believes a contradiction in having the degrees of belief in an inductive argument's premises without the degree of belief forming its conclusion. Thus, the *a priori* argument would not be deductive, as I characterized deduction in section 2. Here we see one reason why it is important to distinguish the two ways that Hume's fork may be understood: (i) as deduction versus induction, or (ii) as *a priori* reasoning versus empirical reasoning. Hume apparently regards all *a priori* arguments as deductive arguments, and hence as arguments that do not yield mere degrees of belief, since degrees of belief do not stand in relations of contradiction and non-contradiction. (At E, p. 22, Hume explicitly identifies arguments that are "probable only" with those "such as regard matter of fact and real existence, according to the division above mentioned" — his fork. See likewise T, p. 651.)

If degrees of belief can be interpreted as personal probabilities, then there are *a priori* arguments purporting to show that certain degrees of belief cannot rationally be accompanied by others: for example, that 60% confidence that p is true cannot be accompanied by 60% confidence that p is false — on pain not of contradiction, but of irrationality ("incoherence"). Whether such *a priori* arguments can resolve Hume's problem is a question that I will take up in section 9.

On the other hand, even if our degrees of belief are personal probabilities rather than claims about the world, perhaps our use of induction to generate our degrees of belief must (on pain of irrationality) be accompanied by certain full beliefs about the world. Suppose we regard Jones as an expert in some arcane subject — so much so that we take Jones' opinions on that subject as our own. Surely, we would be irrational to regard Jones as an expert and yet not believe that there is a higher fraction of truths among the claims in Jones' area of expertise about which Jones is highly confident than among the claims in Jones' area of expertise about which Jones is highly doubtful (presuming that there are many of both). If we did not have this belief, then how could we consider Jones to be an expert? (A caveat: Perhaps our belief that Jones is an expert leaves room for the possibility that Jones has a run of bad luck so that by chance, there is a higher fraction of truths among the claims about which Jones is doubtful than among the claims about which Jones is highly confident. However, perhaps in taking Jones to be an expert, we must at least believe there to be a high objective chance that there is a higher fraction of truths among the claims about which Jones is highly confident than among the claims about which Jones harbors grave doubts.) We use induction to guide our predictions. In effect, then, we take induction as an expert; we take the opinions that induction yields from our observations and make them our own. Accordingly, we must believe that there is (a high chance that there is) a higher fraction of truths among the claims to which induction from our observations assigns a high degree of confidence than among the claims to which induction from our observations assigns a low degree of confidence (presuming that there are many of both). (Perhaps we must even believe that there is (a high chance that there is) a high fraction of truths among the claims to which induction from our observations assigns a high degree of confidence. Otherwise, why would we have such a high degree of confidence in their truth?)

We may now formulate an argument [Skyrms, 1986, pp. 25—27] in the spirit of Hume's. To be justified in using induction to generate our degrees of belief, we must be justified in believing that there is (a high chance that there is) a higher fraction of truths among the claims to which induction from our observations assigns a high degree of confidence than among the claims to which induction from our observations assigns a low degree of confidence. But the falsehood of this claim does not contradict our observations. So we cannot know *a priori* (or deductively) that this claim is true given our observations.

For us to be justified in using induction, would it suffice that we justly possess *a high degree of confidence* that there is (a high chance that there is) a higher fraction of truths among the claims to which induction from our observations assigns a high degree of confidence than among the claims to which induction from our observations assigns a low degree of confidence? Perhaps.[6] If so, then once again, our Humean argument is vulnerable to the reply that there may be an *a priori* argument for our having this high degree of confidence, given our observations, even if there is no contradiction between our observations and the negation of the claim in which we are placing great confidence.

On the other hand, consider our expert Jones. Suppose we merely possess a high degree of confidence that there is a higher fraction of truths among the claims in Jones' area of expertise about which Jones is highly confident than among the claims in Jones' area of expertise about which Jones is highly doubtful. Then although we might give great weight to Jones' opinions, we might well not take Jones' opinions as our own. We should, if possible, consult many other experts along with Jones and weigh each one's opinion regarding p by our confidence in the expert who holds it in order to derive our own opinion regarding p. We should take into account whether we believe that a given expert is more likely to err by placing great confidence in claims about which he should be more cautious or by having grave doubts regarding claims in which he should place greater confidence. But our relation to the expert Jones would then be very different from our relation to our "in-house" expert Induction. In contrast to Jones' opinions, the opinions that induction generates from our observations we take unmodified as our own. If we possessed merely a high degree of confidence that there is a higher fraction of truths among the claims to which induction from our observations assigns a high degree of confidence than among the claims to which induction from our observations assigns a low degree of confidence, then we would have to take the degrees of belief recommended by induction and amend them in light of our estimates of induction's tendency to excessive confidence and tendency to excessive caution. We do not seem to rest our reliance on induction upon any balancing (or even contemplation) of these correction factors.

---

[6]Though Hume doesn't seem to think so: "If there be any suspicion, that the course of nature may change, and that the past may be no rule for the future, all experience becomes useless, and can give rise to no inference or conclusion." (*E*, p. 24)

4 HUME'S FORK: THE SECOND OPTION

Let's now turn to the second option in Hume's fork: Is there an inductive (rather than deductive) — or empirical (rather than *a priori*) — argument taking us from the premises of a given inductive argument to its conclusion? Of course there is: the given inductive argument itself! But since that is the very argument that we are worrying about, we cannot appeal to it to show that we are justified in proceeding from its premises to its conclusion. Is there any *independent* way to argue inductively (or empirically) that this argument's conclusion is true if its premises are true?

Hume argues that there is not. He believes that any inductive (or empirical) argument that we would ordinarily take to be good is of the same kind as the argument that we are worrying about, and so cannot be used to justify that argument on pain of circularity:

> [A]ll experimental conclusions [what Hume on the following page calls "inferences from experience"] proceed upon the supposition that the future will be conformable to the past. To endeavour, therefore, the proof of this last supposition by probable arguments, or arguments regarding existence, must be evidently going in a circle, and taking that for granted, which is the very point in question. (E, p. 23)

> [P]robability is founded on the presumption of a resemblance betwixt those objects, of which we have had experience, and those, of which we have had none; and therefore 'tis impossible this presumption can arise from probability. (T, p. 90)[7]

Since all non-deductive arguments that we consider good are based on the "principle of the uniformity of nature" (that unexamined cases are like the cases that we have already observed), it would be begging the question to use some such argument to take us from the premises to the conclusion of an inductive argument.

For example, suppose we argued as follows for a high degree of confidence that the next slice of bread to be sampled will be nutritious:

1. We have examined many slices of bread for their nutritional value and found all of them to be nutritious.
2. (from 1) If unobserved slices of bread are like the slices of bread that we have already examined, then the next slice of bread we observe will be nutritious.
3. When in the past we examined things that had not yet been observed, we usually found them to be like the things that we had already observed.

[7]Although Hume's is the canonical formulation of the argument, the ideas behind it seem to have been in the air. In 1736, Joseph Butler [1813, p. 17] identified the probability "that all things will continue as we experience they are" as "our only natural reason for believing the course of the world will continue to-morrow, as it has done as far as our experience or knowledge of history can carry us back."

4. So (from 3) unobserved slices of bread are probably like examined slices of bread.

5. Therefore (from 2 and 4) it is likely that the next slice of bread we observe will be nutritious.

But the step from (3) to (4) is based on our confidence that unobserved things are like observed things, which — had we been entitled to it — could have gotten us directly from (1) to (5) without any detour through (2), (3), and (4). As Hume wrote,

> Shou'd it be said, that we have experience, that the same power continues united with the same object, and that like objects are endow'd with like powers, I wou'd renew my question, *why from this experience we form any conclusion beyond those past instances, of which we have had experience*. If you answer this question in the same manner as the preceding, your answer gives still occasion to a new question of the same kind, even *in infinitum*; which clearly proves, that the foregoing reasoning had no just foundation. (*T*, p. 91)

To justify induction by arguing that induction is likely to work well in the future, since it has worked well in the past, is circular.[8]

It might be suggested that although a circular argument is ordinarily unable to justify its conclusion, a circular argument is acceptable in the case of justifying a fundamental form of reasoning. After all, there is nowhere more basic to turn, so all that we can reasonably demand of a fundamental form of reasoning is that it endorse itself. However, certain ludicrous alternatives to induction are also self-supporting. For instance, if induction is based on the presupposition that unexamined cases are like the cases that we have already observed, then take "counterinduction" to be based on the opposite presupposition: that unexamined cases are *unlike* the cases that we have already observed. For example, induction urges us to expect unexamined human cells to contain proteins, considering that

[8]Moreover, surely we did not have to wait to accumulate evidence of induction's track record in order to be justified in reasoning inductively.

It has sometimes been suggested (for instance, by [Black, 1954]) that an inductive justification of induction is not viciously circular. Roughly speaking, the suggestion is that the argument from past observations of bread to bread predictions goes by a form of reasoning involving only claims about bread and other concrete particulars, whereas the argument justifying that form of reasoning ("It has worked well in past cases, so it will probably work well in future cases") goes by a form of reasoning involving only claims about forms of reasoning involving only claims about bread and the like. In short, the second form of reasoning is at a higher level than and so distinct from the first. Therefore, to use an argument of the second form to justify an argument of the first form is not circular.

This response to the problem of induction has been widely rejected on two grounds [BonJour, 1986, pp. 105—6]: (i) Even if we concede that these two forms of argument are distinct, the justification of the first form remains conditional on the justification of the second form, and so on, starting an infinite regress. No form ever manages to acquire unconditional justification. (ii) The two forms of argument do not seem sufficiently different for the use of one in justifying the other to avoid begging the question.

every human cell that has been tested for proteins has been found to contain some. Accordingly, given that same evidence, counterinduction urges us to expect unexamined human cells *not* to contain proteins.[9] Counterinduction is plainly bad reasoning. However, just as induction supports itself (in that induction has worked well in the past, so by induction, it is likely to work well in the future), counterinduction supports itself (in that counterinduction has not worked well in the past, so by counterinduction, it is likely to work well in the future). If we allow induction to justify itself circularly, then we shall have to extend the same privilege to counterinduction (unless we just beg the question by presupposing that induction is justified whereas counterinduction is not). But as I pointed out at the close of section 2, an adequate justification of induction must justify *induction* specifically; it must not apply equally well to all schemes, however arbitrary or cockeyed, for going beyond our observations. Even counterinduction is self-supporting, so being self-supporting cannot suffice for being justified. [Salmon, 1967, pp. 12—17]

It might be objected that there are many kinds of inductive arguments — not just the "induction by enumeration" (taking regularities in our observations and extrapolating them to unexamined cases) that figures in Hume's principal examples, but also (for example) the hypothetico-deductive method, common-cause inference [Salmon, 1984], and inference to the best explanation [Harman, 1965; Thagard, 1978; Lipton, 1991]. Does this diversity undermine Hume's circularity argument?

One might think not: even if an inference to the best explanation could somehow be used to support the "uniformity assumption" grounding one of Hume's inductions by enumeration, we would still need a justification of inference to the best explanation in order to justify the conclusion of the inductive argument.

There is some justice in this reply. However, this reply also misunderstands the goal of Hume's argument. Hume is not merely demanding that we justify induction, pointing out that we have not yet done so, and suggesting that until we do so, we are not entitled to induction's fruits. Hume is purporting to show that it is *impossible* to justify induction. To do that, Hume must show that any possible means of justifying induction either cannot reach its target or begs the question in reaching it. The only way that inference to the best explanation (or some other

[9] Of course, expressed this crudely, "counterinduction" would apparently lead to logically inconsistent beliefs — for instance, that that the next emerald we observe will be yellow (since every emerald we have checked so far has been found not to be yellow) and that the next emerald we observe will be orange (since every emerald we have checked so far has been found not to be orange). One way to reply is to say: so much the worse, then, for any argument that purports to justify counterinduction! Another reply is to say that like induction, counterinduction requires that we form our expectations on the basis of all of our evidence to date, so we must consider that every emerald we have checked so far has been found not merely to be non-yellow and non-orange, but to be green, so by counterinduction, we should expect only that the next emerald to be observed will not be green. Finally, we might point out that induction must apply the principle of the uniformity of nature selectively, on pain of leading to logically inconsistent beliefs, as Goodman's argument will show (in a moment). "Counterinduction" must likewise be selective in applying the principle of the non-uniformity of nature. But no matter: let's suppose that counterinduction allows that principle to be applied in the argument that I am about to give by which counterinduction supports itself.

non-deductive kind of inference) can beg the question is if it, too, is based on some principle of the uniformity of nature. That it has not yet itself been justified fails to show that induction cannot be justified.

In other words, Hume's argument is not that if one non-deductive argument is supported by another, then we have not yet justified the first argument because the second remains ungrounded. Rather, Hume's argument is that every non-deductive argument that we regard as good is of the same kind, so it would be circular to use any of them to support any other. In other words, Hume is arguing that there is a single kind of non-deductive argument (which we now call "induction") that we consider acceptable.

Consequently, it is misleading to characterize Hume's fork as offering us two options: deduction and induction. To put the fork that way gives the impression that Hume is entitled from the outset of his argument to presume that induction is a single kind of reasoning. But that is part of what Hume needs to and does argue for:

> [A]ll arguments from experience are founded on the similarity, which we discover among natural objects, and by which we are induced to expect effects similar to those, which we have found to follow from such objects. (*E*, p. 23)

If some good non-deductive argument that does not turn on a uniformity-of-nature presumption could be marshaled to take us from the premises to the conclusion of an inductive argument, then we could invoke that argument to justify induction. As long as the argument does not rely on a uniformity-of-nature presumption, we beg no question in using it to justify induction; it is far enough away from induction to avoid circularity. Hume's point is not that any non-deductive scheme for justifying an induction leaves us with another question: how is that scheme to be justified? Hume's point is that any non-deductive scheme for justifying an induction leaves us with a question *of the same kind* as we started with, because every non-deductive scheme is fundamentally the same kind of argument as we were initially trying to justify.

Let me put my point in one final way. Hume has sometimes been accused of setting an unreasonably high standard for inductive arguments to qualify as justified: that they be capable of being turned into deductive arguments [Stove, 1973; Mackie, 1974]. In other words, Hume has been accused of "deductive chauvinism": as presupposing that only deductive arguments can justify. But Hume does not *begin* by insisting that deduction is the only non-circular way to justify induction. Hume *argues* for this by arguing that every non-deductive argument is of the same kind. If there were many distinct kinds of non-deductive arguments, Hume would not be able to suggest that any non-deductive defense of induction is circular.

Hume's argument, then, turns on the thought that every inductive argument is based on the same presupposition: that unobserved cases are similar to the cases that we have already observed. However, Nelson Goodman [1954] famously showed that such a "principle of the uniformity of nature" is empty. No matter what the unobserved cases turn out to be like, there is a respect in which they are similar to the cases that we have already observed. Therefore, the "principle of the uniformity of nature" (even if we are entitled to it) is not sufficient to justify making one prediction rather than another on the basis of our observations. Different possible futures would continue different past regularities, but any possible future would continue some past regularity. [Sober, 1988, pp. 63—69]

For example, Goodman says, suppose we have examined many emeralds and found each of them at the time of examination to be green. Then each of them was also "grue" at that time, where

> Object x is grue at time t iff x is green at t where t is earlier than the year 3000 or x is blue at t where t is during or after the year 3000.[10]

Every emerald that we have found to be green at a certain moment we have also found to be grue at that moment. So if emeralds after 3000 are similar to examined emeralds in their grueness, then they will be blue, whereas if emeralds after 3000 are similar to examined emeralds in their greenness, then they will be green. Obviously, the point generalizes: no matter what the color(s) of emeralds after 3000, there will be a respect in which they are like the emeralds that we have already examined. The principle of the uniformity of nature is satisfied no matter how "disorderly" the world turns out to be, since there is inevitably some respect in which it is uniform. So the principle of the uniformity of nature is necessarily true; it is knowable *a priori*. The trouble is that it purchases its necessity by being empty.

Thus, we *can* justify believing in the principle of the uniformity of nature. But this is not enough to justify induction. Indeed, by applying the "principle of the uniformity of nature" indiscriminately (both to the green hypothesis and to the grue hypothesis), we make inconsistent predictions regarding emeralds after 3000. So to justify induction, we must justify expecting certain sorts of past uniformities rather than others to continue.

The same argument has often been made in terms of our fitting a curve through the data points that we have already accumulated and plotted on a graph. Through any finite number of points, infinitely many curves can be drawn. These curves disagree in their predictions regarding the data points that we will gather later. But no matter where those points turn out to lie, there will be a curve running through them together with our current data. Of course, we regard some of the curves passing through our current data as making arbitrary bends later (at the year 3000, for instance); we would not regard extrapolating those curves as justified. To justify induction requires justifying those extrapolations we consider

[10]Here I have updated and simplified Goodman's definition of "grue." He defines an object as "grue" if it is green and *examined* before a given date in the distant future, or is blue otherwise. My definition, which is more typical of the way that Goodman's argument is presented, defines what it takes for an object to be grue *at a certain moment* and does without the reference to the time at which the object was examined. Notice that whether an object is grue at a given moment before 3000 does not depend on whether the object is blue after 3000, just as to qualify as green now, an object does not need to be green later.

"straight" over those that make unmotivated bends.

It might be alleged that of course, at any moment at which something is green, there is a respect in which it is like any other thing at any moment when it is green, whereas no property is automatically shared by any two objects while they are both grue; they must also both lie on the same side of the year 3000, so that they are the same color. Thus, "All emeralds are grue" is just a linguistic trick for papering over a non-uniformity and disguising it as a uniformity. But this move begs the question: why do green, blue, and other colors constitute respects in which things can be alike whereas grue, bleen, and other such "schmolors" do not? Even if there is some metaphysical basis for privileging green over grue, our expectation that unexamined emeralds are green, given that examined ones are, can be based on the principle of the uniformity of nature only if we already know that all green objects are genuinely alike. How could we justify that without begging the question?

The principle of the uniformity of nature does much less to ground some particular inductive inference than we might have thought. At best, each inductive argument is based on some narrower, more specific presupposition about the respect in which unexamined cases are likely to be similar to examined cases. Therefore, Hume is mistaken in thinking that all inductive arguments are of the same kind in virtue of their all turning on the principle of the uniformity of nature. Hence, Hume has failed to show that it is circular to use one inductive argument to support another.

Of course, even if this gap in Hume's argument spoils Hume's own demonstration that there is no possible way to justify induction, it still leaves us with another, albeit less decisive argument (to which I alluded a moment ago) against the possibility of justifying any particular inductive argument. Rather than argue that any inductive justification of induction is *circular*, we can offer a *regress* argument. If observations of past slices of bread justify our high confidence that the next slice of bread we eat will be nutritious, then what justifies our regarding those past observations of bread as confirming that the next slice of bread we eat will be nutritious? Apparently, other observations justify our believing in this link between our bread observations and our bread predictions — by justifying our high confidence that unexamined slices of bread are similar in nutritional value to already observed slices of bread. But whatever those other observations are, this bread uniformity does not follow deductively from them. They confirm it only by virtue of still other observations, which justify our believing in this link between certain observations and the bread uniformity. But what justifies our regarding those observations, in turn, as confirming this link, i.e., as confirming that the bread uniformity is likely if the first set of observations holds? We are off on a regress. A given inductive link is justified (if at all) only by observations, but those observations justify that link (if at all) only through another inductive link, which is justified (if at all) only by observations, which justify that link (if at all) only through another inductive link... . How can it end? If all non-deductive arguments can be justified only by way of observations, then any argument that we might use to justify a given non-deductive argument can itself be justified only by way of observations, and those observations could justify that argument only by an argument that can be justified only by way of other observations, and those observations could justify that argument only by an argument that can be justified only by way of still other observations, and so on infinitely. No bottom ever seems to be reached, so none of these arguments is actually able to be justified.

In other words, any non-deductive argument rests upon high confidence in some contingent (i.e., non-necessary) truth (such as that slices of bread are generally alike nutritionally) linking its conclusion to its premises. We cannot use a non-deductive argument to justify this confidence without presupposing high confidence in some other contingent truth.

We have here a close cousin of Hume's argument — one that leads to the same conclusion, but through a regress rather than a circle.[11] Even if there is no single "principle of the uniformity of nature" on which every inductive argument rests, there remains a formidable argument that no inductive argument can be justified inductively.[12]

5 THREE WAYS OF REJECTING HUME'S PROBLEM

Let's sum up Hume's argument. We cannot use an inductive argument to justify inferring an inductive argument's conclusion from its premises, on pain of circularity. We cannot use a deductive argument to justify do so, because there is no deductive argument from the premises of an inductive argument to its conclusion. So there is no way to justify the step from an inductive argument's premises to its conclusion.

Hume, according to the standard interpretation, holds that we are not entitled to our opinions regarding what we have not observed; those opinions are unjustified. This is a bit different from the conclusion that there is no way to justify that step — that no successful (e.g., non-question-begging) argument can be given for it. Perhaps there is no argument by which induction can be justified, but we are nevertheless justified in using induction, and so we are entitled to the opinions that we arrive at inductively. In this section, I shall look briefly at some forms that this view has taken.

It has often been thought that certain beliefs are justified even though there is no argument by which they acquire their justification; they are "foundational". Many epistemologists have been foundationalists, arguing that unless certain beliefs are

[11] Perhaps it is even Hume's argument. See the passage I quoted earlier from *T*, p. 91, where Hume's argument takes the form of a regress.

[12] Contrast John Norton [2003], who gives similar arguments that there is no general rule of inductive inference. Norton contends that different inductions we draw are grounded on different opinions we have regarding various particular contingent facts (e.g., that samples of chemical elements are usually uniform in their physical properties). He concludes that there is no special problem of *induction*. There is only the question of how there can be an end to the regress of justifications that begins with the demand that we justify those opinions regarding particular contingent facts on which one of our inductive arguments rests.

justified without having to inherit their justification by inference from other beliefs that already possess justification, none of our beliefs is justified. (After all, a regress seems to loom if every belief acquires its justification from other beliefs that acquire their justification from other beliefs... . How could this regress end except with beliefs that are justified without having to have acquired their justification from other beliefs? This regress argument poses one of the classic problems of epistemology.) The beliefs that we acquire directly from making certain observations have often been considered foundational. Another kind of belief that has often been considered foundational consists of our beliefs in certain simple propositions that we know *a priori*, from which we infer the rest of our *a priori* knowledge — for example, that a person is tall if she is tall and thin. We rest our knowledge of this fact on no argument. It has sometimes been maintained that we just "see" — by a kind of "rational insight" — that this fact obtains (indeed, that it is necessary).

Some philosophers have suggested that the proper lesson to take from Hume's argument is that induction is likewise foundational. For instance, Bertrand Russell [1959, pp. 60-69] offers an inductive principle and suggests that it does not need to rest on anything to be justified. It is an independent, fundamental rule of inference.[13] But this approach has all of the advantages of theft over honest toil.[14] It fails to explain why induction rather than some alternative is a fundamental rule of inference. It does not tell us why we should expect the products of inductive reasoning to be true. It tries to make us feel better about having no answer to Hume's problem — but fails. As Wesley Salmon writes:

> This is clearly an admission of defeat regarding Hume's problem, but it may be an interesting way to give up on the problem. The search for the weakest and most plausible assumptions sufficient to justify alternative inductive methods may cast considerable light upon the logical structure of scientific inference. But, it seems to me, admission of unjustified and unjustifiable postulates to deal with the problem is tantamount to making scientific method a matter of faith. [Salmon, 1967, pp. 47–8]

When philosophers have identified certain sorts of beliefs as foundational, they have generally offered some positive account of how those beliefs manage to be

[13]Russell offers the following as a primitive inductive principle: "(a) When a thing of a certain sort A has been found to be associated with a thing of a certain other sort B, and has never been found dissociated from a thing of the sort B, the greater the number of cases in which A and B have been associated, the greater is the probability that they will be associated in a fresh case in which one of them is known to be present; (b) Under the same circumstances, a sufficient number of cases of association will make the probability of a fresh association nearly a certainty, and will make it approach certainty without limit." [1959, p. 66; cf. Russell, 1948, pp. 490–1] Of course, this principle is vulnerable to Goodman's "grue" problem. Other approaches offering primitive inductive principles are Mill's [1872] "axiom of the uniformity of the course of nature" and Keynes' [1921] presumption of "limited independent variety".

[14]The phrase is Russell's: "The method of "postulating" what we want has many advantages; they are the same as the advantages of theft over honest toil." [1919, p. 71]

non-inferentially justified (e.g., of how certain of us qualify as able to make certain kinds of observations, or of how we know certain facts *a priori*). Simply to declare that induction counts as good reasoning seems arbitrary.

A similar problem afflicts the so-called "ordinary-language dissolution" of the problem of induction. Many philosophers have suggested that induction is a fundamental kind of reasoning and that part of what we *mean* by evidence rendering a given scientific theory "justified", "likely", "well supported", and so forth is that there is a strong inductive argument for it from the evidence. Hence, to ask "Why is inductive reasoning able to justify?" is either to ask a trivial question (because by definition, inductive reasoning counts as able to justify) or to ask a meaningless question (because, in asking this question, we are not using the word "justify" in any familiar, determinate sense). As P.F. Strawson remarks:

> It is an analytic proposition that it is reasonable to have a degree of belief in a statement which is proportional to the strength of the evidence in its favour; and it is an analytic proposition, though not a proposition of mathematics, that, other things being equal, the evidence for a generalization is strong in proportion as the number of favourable instances, and the variety of circumstances in which they have been found, is great. So to ask whether it is reasonable to place reliance on inductive procedures is like asking whether it is reasonable to proportion the degree of one's convictions to the strength of the evidence. Doing this is what 'being reasonable' *means* in such a context. ... In applying or withholding the epithets 'justified', well founded', &c., in the case of specific beliefs, we are appealing to, and applying, inductive standards. But to what standards are we appealing when we ask whether the application of inductive standards is justified or well grounded? If we cannot answer, then no sense has been given to the question. [Strawson, 1952, pp. 256-7][15]

In contending that it is either trivial or meaningless to ask for a justification of induction, the ordinary-language approach does not purport to "solve" the problem of induction, but rather to "dissolve it": to show that the demand for a justification of induction should be rejected.

One might reply that this line of thought offers us no reason to believe that the conclusions of strong inductive arguments from true premises are likely to be true. But the ordinary-language theorist disagrees: that these conclusions are the conclusions of strong inductive arguments from true premises is *itself* a good reason to believe that they are likely to be true. What else could we mean by a "good reason" than the kind of thing that we respect as a good reason, and what's more respectable than induction? In his *Philosophical Investigations*,

[15]Cf. [Horwich, 1982, pp. 97—98; Salmon, Barker, and Kyburg, 1965]. For critique of this view, I am especially indebted to [BonJour, 1998, pp. 196—199; Salmon, 1967, pp. 49—52; and Skyrms, 1986, pp. 47—54].

Ludwig Wittgenstein recognizes that we may feel the need for a standard that grounds our standards for belief, but he urges us to resist this craving:

> 480. If it is now asked: But how *can* previous experience be a ground for assuming that such-and-such will occur later on? — the answer is: What general concept have we of grounds for this kind of assumption? This sort of statement about the past is simply what we call a ground for assuming that this will happen in the future...
>
> 481. If anyone said that information about the past could not convince him that something would happen in the future, I should not understand him. One might ask him: what do you expect to be told, then? What sort of information do you call a ground for such a belief? ... If *these* are not grounds, then what are grounds? — If you say these are not grounds, then you must surely be able to state what must be the case for us to have the right to say that there are grounds for our assumption...
>
> 482. We are misled by this way of putting it: 'This is a good ground, for it makes the occurrence of the event probable.' That is as if we had asserted something further about the ground, which justified it as a ground; whereas to say that this ground makes the occurrence probable is to say nothing except that this ground comes up to a particular standard of good grounds — but the standard has no grounds!.....
>
> 484. One would like to say: 'It is a good ground only because it makes the occurrence *really* probable.'...
>
> 486. Was I justified in drawing these consequences? What is *called* a justification here? — How is the word 'justification' used?... [Wittgenstein, 1953; cf. Rhees and Phillips, 2003, pp. 73-77]

But this argument makes the fact that induction counts for us as "good reasoning" seem utterly arbitrary. We have not been told why we *should* respect induction in this way. We have simply been reminded that we *do*.

If part of what "good reason" *means* is that inductive reasons qualify as good, then so be it. We can still ask why we ought to have a term that applies to inductive arguments (and not to bogus arguments instead or in addition) and where a consequence of its applying to some argument is that we ought to endorse that argument. The mere fact that these circumstances of application and consequences of application are coupled in the meaning of "good reason" cannot prevent us from asking why they ought to be coupled — just as (to use Michael Dummett's example) the term "Boche" has "German" as its circumstance of application and "barbaric" as its consequence of application, so that it is contradictory to say "The Boche are not really barbaric" or "The Germans are not really the Boche", but we can still ask why we ought (not) to have "Boche" in our language. [Dummett, 1981, p. 454; Brandom, 1994, pp. 126–127]. It is not contradictory to conclude that we should not use the term "Boche" because it is not true that the Germans are barbaric. Analogously, we can ask why we ought to have a term like "good argument" if an inductive argument automatically qualifies as a "good argument" and any "good argument" is automatically one that conveys justification from its premises to its conclusion. Without some account of why those circumstances of application *deserve to* go with those consequences of application, we have no reason to put them together — and so no reason to think better of arguments that qualify by definition as "good." As Salmon says,

> It sounds very much as if the whole [ordinary-language] argument has the function of transferring to the word 'inductive' all of the honorific connotations of the word 'reasonable', quite apart from whether induction is good for anything. The resulting justification of induction amounts to this: If you use inductive procedures you can call yourself 'reasonable' — *and isn't that nice!* [Salmon, 1957, p. 42; cf. Strawson, 1958]

It does not show us why we ought to be "reasonable".

However, the ordinary-language dissolutionist persists, to ask why we *ought* to use the term "good reason" — why we ought to couple its circumstances and consequences of application — is just to ask for *good reasons* for us to use it. We cannot find some point outside of all of our justificatory standards from which to justify our standards of justification. What standards of justification do we mean when we demand a justification of induction? If an argument meets some "standards of justification" that are not our usual ones, then it does not qualify as a justification. On the other hand, if it meets our usual standards of justification, then (since, Hume showed, no deductive argument can succeed in justifying induction) the argument will inevitably be inductive and so beg the question, as Hume showed.

But we do *not* need to specify our standards of justification in advance in order for our demand for a justification of induction to make sense. We know roughly what a justification is, just as we know roughly what it is for an argument to beg the question. A justification of induction would consist of an argument that we believe is properly characterized as justificatory — an argument that, we can show, meets the same standards as the familiar arguments that we pretheoretically recognize as justificatory. In showing this, we may be led to new formulations of those standards — formulations that reveal features that had been implicit in our prior use of "justification". When we ask for a justification of induction, we are not trying to step entirely outside of our prior standards of justification, but at the same time, we are not asking merely to be reminded that induction is one of the kinds of reasoning that we customarily recognize as good. Rather, we are asking for some independent grounds for recognizing induction as good — a motivation for doing so that is different enough to avoid begging the question, but not so different that it is unrecognizable as a justification. Of course, it may be unclear what sort of reasoning could manage to walk this fine line until we have found it. But that does not show that our demand for a justification of induction is trivial or meaningless.

Here is an analogy. Suppose we want to know whether capital punishment counts as "cruel" in the sense in which the United States Constitution, the English Bill of Rights, and the Universal Declaration of Human Rights outlaw cruel punishment. One might argue that since capital punishment was practiced when these documents were framed (and is practiced today in the United States), "cruel" as they (and we) mean it must not apply to capital punishment. But the question of whether capital punishment is cruel cannot be so glibly dismissed as trivial (if we mean "cruel" in our sense) or meaningless (if we mean "cruel" in some other, unspecified sense) — and could not be so dismissed even if no one had ever thought that capital punishment is cruel. What we need, in order to answer the question properly, is an independent standard of what it takes for some punishment to qualify as "cruel". The standard must fit enough of our pretheoretic intuitions about cruel punishment (as manifested, for example, in legal precedents) that we are justified in thinking that it has managed to make this notion more explicit, rather than to misunderstand it. Furthermore, the standard must derive its credentials independently from whatever it says about capital punishment, so that we avoid begging the question in using this standard to judge whether capital punishment qualifies as cruel. Of course, prior to being sufficiently creative and insightful to formulate such a standard, we may well be unable to see how it could be done. But that is one reason why it takes great skill to craft good arguments for legal interpretations — and, analogously, why it is difficult to address Hume's problem.

Another approach that deems induction to be good reasoning, even though no non-question-begging argument can be given to justify it, appeals to epistemological naturalism and externalism. On this view, if inductive reasoning from true premises does, in fact, tend to lead to the truth, then an inductive argument has the power to justify its conclusion even though the reasoner has no non-circular basis for believing that the conclusions of inductive arguments from true premises are usually true [Brueckner, 2001; Kornblith, 1993; Papineau, 1993, pp. 153–160; Sankey, 1997; van Cleve, 1984].

To my mind, this approach simply fails to engage with Hume's problem of induction. The externalist believes that we qualify as having good reasons for our opinions regarding the future as long as inductive arguments from true premises do in fact usually yield the truth regarding unexamined cases. But the problem of induction was to offer a good reason to believe that a given inductive argument from true premises will likely yield the truth regarding unexamined cases. Suppose the externalist can persuade us that to be justified in some belief is to arrive at it by reliable means. Then we are persuaded that *if* induction is actually reliable, then the conclusion of an inductive argument (from justified premises) is justified. We are also persuaded that *if* induction actually is reliable, then an inductive reasoner is justified in her belief (arrived at inductively, from the frequent success of past inductive inferences) that induction will continue to be reliable. Nevertheless, the externalist has *not* persuaded us that induction *is* reliable.

6 HUME'S CONCLUSION

Hume, according to the standard interpretation of his view, is an "inductive skeptic": he holds that we are not entitled to our opinions regarding what we have not observed. There are plenty of textual grounds for this interpretation. For example, in a passage that we have already quoted (*T*, p. 91), he says that an inductive argument for induction has "no just foundation", suggesting that his main concern is whether induction has a just foundation. Sometimes Hume appears to concede induction's justification:

> I shall allow, if you please, that the one proposition [about unexamined cases] may justly be inferred from the other [about examined cases]: I know in fact, that it always is inferred. (*E*, p. 22).

But his "if you please" suggests that this concession is merely rhetorical – for the sake of argument. His point is that someone who believes that we are justified in having expectations regarding unexamined cases should be concerned with uncovering their justification. When Hume finds no justification, he concludes that these expectations are unjustified.

Accordingly, Hume says, our expectations are not the product of an "inference" (*E*, p. 24) or some "logic" (*E*, p. 24) or "a process of argument or ratiocination" (*E*, p. 25):

> it is not reasoning which engages us to suppose the past resembling the future, and to expect similar effects from causes, which are, to appearance, similar. (*E*, p. 25)

Rather, Hume says, our expectations regarding what we have not observed are the result of the operation of certain innate "instincts" (*E*, pp. 30, 37, 110). When (for example) the sight of fire has generally been accompanied by the feeling of heat in our past experience, these instinctual mental mechanisms lead us, when we again observe fire, to form a forceful, vivid idea of heat — that is, to expect heat:

> What then is the conclusion of the whole matter? A simple one; though, it must be confessed, pretty remote from the common theories of philosophy. All belief of matter of fact or real existence is derived merely from some object, present to the memory or senses, and a customary conjunction between that and some other object. Or in other words; having found, in many instances, that any two kinds of objects, flame and heat, snow and cold, have always been conjoined together; if flame or snow be presented anew to the senses, the mind is carried by custom to expect heat or cold, and to *believe*, that such

> a quality does exist, and will discover itself upon a nearer approach. This belief is the necessary result of placing the mind in such circumstances. It is an operation of the soul, when we are so situated, as unavoidable as to feel the passion of love, when we receive benefits; or hatred, when we meet with injuries. All these operations are a species of natural instincts, which no reasoning or process of the thought and understanding is able, either to produce, or to prevent. (*E*, p. 30)

With prose like that, is it any wonder that Hume's argument has become such a classic?

Hume believes that we cannot help but form these expectations, in view of the way our minds work. So Hume does not recommend that we try to stop forming them. Any such attempt would be in vain. But Hume's failure to recommend that we try to resist this irresistible psychological tendency should not lead us to conclude (with Garrett [1997]) that Hume believes our expectations to be justified or that Hume is uninterested in evaluating their epistemic standing.

Hume sometimes uses normative-sounding language in giving his naturalistic, psychological account of how we come by our expectations:

> [N]one but a fool or a madman will ever pretend to dispute the authority of experience, or to reject that great guide of human life. . . (*E*, p. 23)

But by "authority" here, he presumably means nothing normative, but merely the control or influence that experience in fact exercises over our expectations — experience's "hold over us".[16] As Hume goes on to explain on the same page, he wants "to examine the principle of human nature, which gives this mighty authority to experience. . . " That "principle of human nature" has no capacity to justify our expectations; it merely explains them. (And since it is irresistible, none can reject it and none but a fool or a madman will *pretend* to reject it.) Hume often terms this principle of the association of ideas "custom" or "habit":

> 'Tis not, therefore, reason which is the guide of life, but custom. That alone determines the mind, in all instances, to suppose the future conformable to the past. However easy this step may seem, reason would never, to all eternity, be able to make it. (*T*, p. 652)

[16]That "authority" here should be interpreted as brute power to bring about rather than entitlement ("rightful authority") to bring about is evident from other passages: "If the mind be not engaged by argument to make this step, it must be induced by some other principle of equal weight and authority [namely, custom]. . . " (*E*, p. 27). An interpretation of "authority" as normative has led some to regard Hume not as an inductive skeptic, but instead as offering a *reductio* of some particular conception of knowledge on the grounds that it would deem what we know inductively not to be knowledge: "Far from being a skeptical challenge to induction, Hume's 'critique' is little more than a prolonged argument for the general position that Newton's inductive method must replace the rationalistic model of science" according to which *a priori* reasoning is "capable of deriving sweeping factual conclusions." [Beauchamp and Rosenberg, 1981, p. 43] See also [Smith, 1941; Stroud, 1977].

> For wherever the repetition of any particular act or operation produces a propensity to renew the same act or operation, without being impelled by any reasoning or process of the understanding; we always say, that this propensity is the effect of *Custom*. By employing that word, we pretend not to have given the ultimate reason of such a propensity. We only point out a principle of human nature, which is universally acknowledged, and which is well known by its effects. ... [A]fter the constant conjunction of two objects, heat and flame, for instance, weight and solidity, we are determined by custom alone to expect the one from the appearance of the other. (*E*, p. 28)

Hume is thus offering a theory of how our minds work.[17] His "arguments" for this scientific theory are, of course, inductive. What else could they be? For instance, Hume points out that in the cases we have seen, a correlation in some observer's past observations (such as between seeing fire and feeling heat) is usually associated with that observer's forming a certain expectation in future cases (e.g., expecting heat, on the next occasion of seeing fire). Having noted this association between an observer's expectations and the correlations in her past observations, Hume extrapolates the association; Hume is thereby led to form certain expectations regarding what he has not yet observed, and so to believe in a general "principle of human nature".[18] Some interpreters have suggested that since Hume is here using induction, he must believe that he is (and we are) entitled to do so — and so that induction is justified.[19] However, in my view, Hume's use of induction shows no such thing. Hume says that we cannot help but form expectations in an inductive way — under the sway ("authority") of certain mental instincts. Hume's behavior in forming his own expectations regarding the expectations of others is just one more example of these mental instincts in action. So Hume's own belief in his theory of human nature is accounted for by that very theory. Like his expectations regarding the next slice of bread he will eat, his expectations regarding human belief-formation fail to suggest that Hume regards our expectations regarding what we have not observed to be justified.

By the same token, Hume notes that as we become familiar with cases where someone's expectations regarding what had not yet been observed turn out to be accurate and other cases where those expectations turn out not to be met, we notice the features associated with these two sorts of cases. We then tend to be

---

[17] As we have seen, the "principle of the uniformity of nature" must be applied selectively, on pain of leading to contradiction. So insofar as Hume's theory of the mind incorporates such a principle as governing the association of ideas, it does not suffice to account for our expectations.

[18] By analogous means, Hume arrives at other parts of his theory, such as that every simple idea is a copy of a prior impression, that there are various other principles of association among ideas, etc.

[19] See, for instance, [Garrett, 1997], according to which Hume's main point is not to give "an evaluation of the epistemic worth of inductive inferences" (p. 94) but rather to do cognitive psychology — to identify the component of the mind that is responsible for those opinions (imagination rather than reason). For more discussion of this line of interpretation, see [Read and Richman, 2000].

guided by these associations in forming future expectations. This is the origin of the "Rules by which to judge of causes and effects" that Hume elaborates (*T*, p. 173): "reflexion on *general rules* keeps us from augmenting our belief upon every encrease of the force and vivacity of our ideas" (*T*, p. 632). We arrive at these "rules" by way of the same inductive instincts that lead us to form our other expectations regarding what we have not observed. Some have argued that Hume's offering these rules shows that Hume is not an inductive skeptic, since if there are rules distinguishing stronger from weaker inductive arguments, then such arguments cannot all be bad.[20] But as I have explained, Hume's endorsement of these rules does not mean that he believes that expectations formed in accordance with them are justified. The rules do not distinguish stronger from weaker inductive arguments. Rather, they result from our instinctively forming expectations regarding the expectations we tend to form under various conditions.

Today, in the wake of Darwin's theory of evolution by natural selection, we might argue that natural selection has equipped us with various innate belief-forming instincts. Creatures with these instincts stood at an advantage in the struggle for existence, since these instincts gave them accurate expectations and so enabled them to reproduce more prolifically. But once again, this theory does not solve Hume's problem; it does not justify induction. To begin with, we have used induction of some kind to arrive at this scientific explanation of the instincts's origin. Furthermore, even if the possession of a certain belief-forming instinct was advantageous to past creatures because it tended to lead to accurate predictions, we would need to use induction to justify regarding the instinct's past predictive success as confirming its future success.

7 BONJOUR'S *A PRIORI* JUSTIFICATION OF INDUCTION

Laurence BonJour [1986; 1998, pp. 203–216] has maintained that philosophers have been too hasty in accepting Hume's argument that there is no *a priori* means of proceeding from the premises of an inductive argument to its conclusion. BonJour accepts that there is no contradiction in an inductive argument's premises being true and its conclusion false. But BonJour rejects Hume's view that the only truths that can be established *a priori* are truths that hold on pain of contradiction. BonJour is inclined to think that there is something right in the view that only necessary truths can be known *a priori*. But again, he believes that there are necessary truths that are not analytic (i.e., necessary truths the negations of which are not contradictions).

BonJour concedes that we cannot know *a priori* that anything like the "principle of the uniformity of nature" holds. However, he thinks that a good inductive argument has as its premise not merely that the fraction of Gs among examined Fs is m/n, but something considerably less likely to be a coincidence: that the fraction converged to m/n and has since remained approximately steady as more

[20] See prior note.

(and more diverse) Fs have been examined. BonJour suggests that (when there is no relevant background information on the connection between being F and being G or on the incidence of Gs among Fs) we know by *a priori* insight (for certain properties F and G) that when there has been substantial variation in the locations and times at which our observations were made, the character of the observers, and other background conditions, the fraction m/n is unlikely to remain steady merely as a brute fact (i.e., a contingent fact having no explanation) or just by chance — e.g., by there being a law that approximately r/n of all Fs are G, but "by chance" (analogous to a fair coin coming up much more often heads than tails in a long run of tosses), the Fs we observed were such that the relative frequency of Gs among them converged to a value quite different from r/n and has since remained about there. We recognize *a priori* that it is highly unlikely that any such coincidence is at work.

Moreover, as the observed Fs become ever more diverse, it eventually becomes *a priori* highly unlikely that the explanation for the steady m/n fraction of Gs among them is that although it is not the case that there is a law demanding that approximately m/n of all Fs are G, the Fs that we have observed have all been Cs and there is a law demanding that approximately m/n of all FCs are G. As the pool of observed Fs becomes larger and more diverse, it becomes increasingly *a priori* unlikely that our observations of Fs are confined only to Cs where the natural laws demand that FCs behave differently from other Fs. For that to happen would require an increasingly unlikely coincidence: a coordination between the range of our observations and the natural laws. (Analogous arguments apply to other sorts of possible explanations of the steady m/n fraction of Gs among the observed Fs, such as that the Fs we observed in each interval happened to consist of about the same fraction of Cs, and the laws assign a different likelihood to FCs being G than to $F \sim C$s being G.) Thus, in the case of a good inductive argument, it is *a priori* likely (if the act of observing an F is not itself responsible for its G-hood) that our evidence holds only if it is a law that approximately m/n of all Fs are G. In other words, we know *a priori* that the most likely explanation of our evidence is the "straight inductive explanation".

BonJour does not say much about the likelihoods that figure in these truths that we know *a priori*. They seem best understood as "logical probabilities" like those posited by John Maynard Keynes [1921] and Rudolf Carnap [1950], among others — logical, necessary, probabilistic relations obtaining among propositions just in virtue of their content.[21] Just as we know by rational insight that the premises of a deductive argument can be true only if the conclusion is true, so likewise (BonJour seems inclined to say) we know by rational insight that the premises of a good inductive argument make its conclusion highly likely. As Keynes wrote:

> Inasmuch as it is always assumed that we can sometimes judge directly that a conclusion *follows from* a premises, it is no great extension of this assumption to suppose that we can sometimes recognize that a

[21] BonJour [personal communication] is sympathetic to the logical interpretation of probability.

> conclusion *partially follows* from, or stands in a relation of probability to a premiss. [Keynes, 1921, p. 52]

Presumably, part of what we grasp in recognizing these probabilistic relations is that (in the absence of other relevant information) we should have high confidence in any proposition that stands to our evidence in a relation of logically high probability. But I wonder *why* we should. If a conclusion logically follows from our evidence, then we should believe the conclusion because it is impossible for the premises to be true without the conclusion being true. But we cannot say the same in the case of a conclusion that is merely "highly logically probabilified" by our evidence. To say that the conclusion is made likely, and so we should have great confidence that it is true, is to risk punning on the word "likely". There is *some* sort of logical relation between the premises and the conclusion, and we call this relation "high logical probabilification" presumably because it obeys the axioms of probability and we think that (in the absence of other relevant information) we ought to align our subjective degrees of probability (i.e., our degrees of confidence) with it. But then this relation needs to do something to *deserve* being characterized as "high logical probabilification". What has it done to merit this characterization?

The problem of justifying induction then boils down to the problem of justifying the policy of being highly confident in those claims that stand in a certain logical relation to our evidence. Calling that relation "high logical probabilification" or claiming rational insight into that relation's relevance to our assignment of subjective probability does not reveal what that relation does to merit our placing such great weight upon it. Why should my personal probability distribution be one of the "logical" probability functions? How do we know that my predictions would then tend to be more accurate? (To say that they are then "likely" to be more accurate, in the sense of "logical" probability, is to beg the question.)

Let's turn to a different point. BonJour recognizes that no matter how many Fs we observe, there will always be various respects C in which they are unrepresentative of the wider population of Fs. (All Fs so far observed existed sometime before tomorrow, to select a cheap example.) Of course, we can never show *conclusively* that the steady m/n frequency of Gs among the observed Fs does not result from a causal mechanism responsible for the G-ness of FCs but not for the G-ness of other Fs. That concession is no threat to the justification of induction, since strong inductions are not supposed to be proofs. But how can we be entitled even to place high confidence in the claim that there is no such C? BonJour writes (limiting himself to spatial Cs for the sake of the example):

> [Our data] might be skewed in relation to some relevant factor C ... because C holds in the limited area in which all the observations are in fact made, but not elsewhere. It is obviously a quite stubborn empirical fact that all of our observations are made on or near the surface of the earth, or, allowing for the movement of the earth, in the general region of the solar system, or at least in our little corner of the galaxy, and

> it is possible that C obtains there but not in the rest of the universe, in which case our standard inductive conclusion on the basis of those observations would presumably be false in relation to the universe as a whole, that is, false simpliciter. ... The best that can be done, I think, is to point out that unless the spatio-temporal region in which the relevant C holds is quite large, it will still be an unlikely coincidence that our observations continue in the long run to be confined to that region. And if it is quite large, then the inductive conclusion in question is in effect true within this large region in which we live, move, and have our cognitive being. [BonJour, 1998, p. 215]

BonJour seems to be saying that we know *a priori* that it would be very unlikely for all of our observations so far to have been in one spatiotemporal region (or to have been made under one set of physical conditions C) but for the laws of nature to treat that region (or those conditions) differently from the region in which (or conditions under which) our next observation will be made. But this does not seem very much different from purporting to know *a priori* that considering the diversity of the Fs that we have already examined and the steady rate at which Gs have arisen among them, the next case to be examined will probably be like the cases that we have already examined. Inevitably, there will be infinitely many differences between the Fs that we have already examined and the Fs that we will shortly examine (if we are willing to resort to "gruesome" respects of similarity and difference). Is the likely irrelevance of these differences (considering the irrelevance of so many other factors, as manifested in the past steadiness of the m/n fraction) really something that we could know *a priori*? Isn't this tantamount to our knowing *a priori* (at least for certain properties F and G) that if it has been the case at every moment from some long past date until now that later Fs were found to be Gs at the same rate as earlier Fs, then (in the absence of other relevant information) it is likely that the Fs soon to be examined will be Gs at the same rate as earlier Fs? That seems like helping ourselves directly to induction *a priori*.

Consider these two hypotheses: (h) A law requires that approximately m/n of all Fs are G; (k) A law requires that approximately m/n of all Fs before today are G but that no Fs after today are G. On either of these hypotheses, it would be very likely that approximately m/n of any large, randomly-selected sample of Fs before today will be G.[22] Do we really have *a priori* insight into which of these hypotheses provides the most likely explanation of this fact? BonJour apparently thinks that we do, at least for certain Fs and Gs, since (k) would require an *a priori* unlikely coordination between the present range of our observations and the discriminations made by the natural laws. Moreover, insofar as (k) is changed so that the critical date is pushed back from today to the year 3000 (or 30,000, or 300,000...), (k)'s truth would require less of an *a priori* unlikely coordination between the laws and the present range of our observations — but, BonJour seems

[22]See section 10.

to be saying, it then becomes increasingly the case that "the inductive conclusion in question is in effect true within this large region in which we live, move, and have our cognitive being."

8 REICHENBACH'S PRAGMATIC JUSTIFICATION OF INDUCTION

Hans Reichenbach [1938, pp. 339–363; 1949a, 469–482; 1968, 245–246] has proposed an intriguing strategy for justifying induction. (My discussion is indebted to [Salmon, 1963].) Reichenbach accepts Hume's argument that there is no way to show (without begging the question) that an inductive argument from true premises is likely to lead us to place high confidence in the truth. However, Reichenbach believes that we can nevertheless justify induction by using pragmatic (i.e., instrumental, means-ends) reasoning to justify the policy of forming our expectations in accordance with an inductive rule. Of course, Reichenbach's argument cannot justify this policy by showing that it is likely to lead us to place high confidence in the truth — since that approach is blocked by Hume's argument. Reichenbach believes that to justify an inductive policy, it is not necessary to show that induction will probably be successful, or that induction is more likely to succeed than to fail, or even that the claims on which induction leads us to place high confidence are at least 10% likely to be true. It suffices, Reichenbach thinks, to show that *if* any policy for belief-formation will do well in leading us to the truth, then induction will do at least as well. That is, Reichenbach believes that to justify forming our opinions by using induction, it suffices to show that no other method can do better than induction — even if we have not shown anything about how well induction will do. Induction is (at least tied for) our best hope, according to Reichenbach, though we have no grounds for being at all hopeful.[23]

BonJour [1998, pp. 194—196] has objected that Reichenbach's argument cannot justify induction because it does not purport to present us with good grounds for believing that induction will probably succeed. It does not justify our believing in the (likely) truth of the claims that receive great inductive support. It purports to give us *pragmatic* rather than *epistemic* grounds for forming expectations in accordance with induction. Surely, BonJour says, if we are not entitled to believe that a hypothesis that has received great inductive support is likely to be true, then we have not really solved Hume's problem. Reichenbach writes:

> A blind man who has lost his way in the mountains feels a trail with his stick. He does not know where the path will lead him, or whether it may take him so close to the edge of a precipice that he will be plunged into the abyss. Yet he follows the path, groping his way step by step; for if there is any possibility of getting out of the wilderness, it is by

[23] Feigl [1950] distinguishes "validating" a policy (i.e., deriving it from more basic policies) from "vindicating" it (i.e., showing that it is the right policy to pursue in view of our goal). A fundamental policy cannot be validated; it can only be vindicated. Accordingly, Reichenbach is often interpreted as purporting to "vindicate" induction.

> feeling his way along the path. As blind men we face the future; but we feel a path. And we know: if we can find a way through the future it is by feeling our way along this path. [Reichenbach, 1949a, p. 482]

BonJour replies:

> We can all agree that the blind man should follow the path and that he is, in an appropriate sense, acting in a justified or rational manner in doing so. But is there any plausibility at all to the suggestion that when we reason inductively, or accept the myriad scientific and commonsensical results that ultimately depend on such inference, we have no more justification for thinking that our beliefs are likely to be true than the blind man has for thinking that he has found the way out of the wilderness? [BonJour, 1998, pp. 195—196]

BonJour's objection illustrates how much Reichenbach is prepared to concede to Hume. Reichenbach's point is precisely that an agent "makes his posits because they are means to his end, not because he has any reason to believe in them." [Reichenbach, 1949b, p. 548][24]

The policy for forming our expectations that Reichenbach aims to justify is the "straight rule": If n Fs have been examined and m have been found to be G, then take m/n to equal (within a certain degree of approximation) the actual fraction of Gs among Fs — or, if there are infinitely many Fs and Gs (and so the fraction is infinity divided by infinity), take m/n to approximate the limiting relative frequency of Gs among Fs. Reichenbach presents his policy as yielding beliefs about limiting relative frequencies, rather than as yielding degrees of confidence, because Reichenbach identifies limiting relative frequencies with objective chances. Accordingly, the "straight rule" is sometimes understood as follows: If n Fs have been examined and m have been found to be G, then take m/n to equal (within a certain degree of approximation) an Fs objective chance of being G.

Reichenbach then argues that if there is a successful policy for forming beliefs about limiting relative frequencies, then the straight rule will also succeed. For instance, suppose that some clairvoyant can predict the outcome of our next experiment with perfect accuracy. Then let F be that the clairvoyant makes a certain prediction regarding the outcome and G be that the clairvoyant is correct. Since m/n (from our past observations of Fs and Gs) equals 1, the straight rule endorses our taking 1 to be the limiting relative frequency of truths among the clairvoyant's predictions (or chance that a given prediction by the clairvoyant will come to pass). In short, Reichenbach argues that if the world is non-uniform (i.e., if there is no successful policy for forming beliefs about limiting relative frequencies), then the straight rule will fail but so will any other policy, whereas if the world is uniform (i.e., if there is a successful policy), then the straight rule will seize upon the uniformity (or policy). Hence, the straight rule can do no worse than

[24]However, Reichenbach [1968, p. 246] says "in my theory good grounds are given to treat a posit as true".

any other policy for making predictions. Under any circumstances, it is at least tied for best policy. Therefore, its use is pragmatically justified.

Even this rough statement of Reichenbach's argument suffices to reveal several important difficulties it faces. First, as we saw in connection with the principle of the uniformity of nature, the straight rule licenses logically inconsistent predictions. For instance, if the Fs are the emeralds, then "G" could be "green" or "grue." In either case, m/n is 1, but we cannot apply the straight rule to both hypotheses on pain of believing that emeralds after 3000 are all green and all blue. For the rule to license logically consistent predictions, it must consist of the straight rule along with some principle selecting the hypotheses to which the straight rule should be applied. However, a straight rule equipped with a principle of selection is no longer guaranteed to succeed if any rule will. If all emeralds are grue (and there are emeralds after 3000), then a straight rule equipped with a principle of selection favoring the grue hypothesis over the green hypothesis will succeed whereas a straight rule favoring green over grue will fail (as long as all of our emeralds are observed before the year 3000).[25]

Here is a related point. The straight rule does not tell us what properties to take as F and G. It merely specifies, given F and G, what relative frequency (or chance) to assign to unexamined Fs being G. But then there could be an F and a G to which the straight rule would lead us to assign an accurate relative frequency, but as it happens, we fail to think of that F and G. (For instance, we might simply not think of tallying the rate at which the clairvoyant's predictions have been accurate in the past, or of taking "grue" as our G. [Putnam, 1994, p. 144]) In other words, the straight rule is concerned with justifying hypotheses once they have been thought up — not with thinking them up in the first place. (It is not a "method of discovery"; it is a "method of justification.") So the sense in which the straight rule is guaranteed to lead us to the genuine limiting relative frequency, if any rule could, is somewhat limited.

Let's set aside this difficulty to focus on another. Reichenbach compares the straight rule to other policies for making predictions. But what about the policy of making no predictions at all? Of course, the straight rule is more likely (or, at least, not less likely) than this policy to arrive at accurate predictions. But it is also more likely than this policy to arrive at inaccurate predictions. If our goal is to make accurate predictions and we incur no penalty for making inaccurate ones, then the straight rule is obviously better than the no-prediction policy. This seems to be Reichenbach's view:

> We may compare our situation to that of a man who wants to fish in

[25] Reichenbach responds, "The rule of induction ... leads only to posits that are justified asymptotically." [1949a, p. 448] In the long run, we observe emeralds after 3000. So although "applying the rule of induction to [grue], we shall first make bad posits, but while going on will soon discover that [emeralds after 3000 are not grue]. We shall thus turn to positing [green] and have success." This response is vulnerable to the reply that we make all of our actual predictions in the short run rather than the long run (as I will discuss momentarily). Moreover, if we "apply the rule of induction" to grue as well as to green, then we make predictions that are contradictory, not merely inaccurate.

> an unexplored part of the sea. There is no one to tell him whether or not there are fish in this place. Shall he cast his net? Well, if he wants to fish in that place I should advise him to cast the net, to take the chance at least. It is preferable to try even in uncertainty than not to try and be certain of getting nothing. [Reichenbach, 1938, pp. 362—363][26]

This argument presumes that there is no cost to trying — that the value of a strategy is given by the number of fish that would be caught by following it, so that if a strategy leads us to try and fail, then its value is zero, which is the same as the value of the strategy of not trying at all. So the fisherman has everything to gain and nothing to lose by casting his net. But doesn't casting a net come with some cost? (It depletes the fisherman's energy, for instance.)

In other words, the straight rule offers us some prospect of making accurate predictions, whereas the policy of making no predictions offers us no such prospect, so (Reichenbach concludes) the straight rule is guaranteed to do no worse than the no-prediction rule. But why shouldn't our goal be "the truth, the whole truth, and *nothing but the truth*", so that we favor making accurate predictions over making inaccurate predictions or no predictions, but we favor making no predictions over making inaccurate predictions? Reichenbach then cannot guarantee that the straight rule will do at least as well as any other policy, since if the straight rule fails, then the policy of making no predictions does better. Thus, Reichenbach's argument may favor the straight rule over alternative methods of making predictions, but it does not justify making some predictions over none at all.

Let us now look at Reichenbach's more rigorous formulation of his argument. Consider a sequence of Fs and whether or not each is G. Perhaps the sequence is $G, \sim G, \sim G, G, \sim G, \sim G, \sim G \ldots$, and so at each stage, the relative frequency of Gs among the Fs is 1/1, 1/2, 1/3, 1/2, 2/5, 1/3, 2/7... . Either this sequence converges to a limiting relative frequency or it does not. (According to Reichenbach, this is equivalent to: either there is a corresponding objective chance or there is not.) For instance, if 1/4 of the first 100 Fs are G, 3/4 of next 1000 are G, 1/4 of the next 10,000 are G, and so forth, then the relative frequency of Gs among Fs never converges and so there is no limiting relative frequency. No method can succeed in arriving at the limit if there is no limit, so in that event, the straight rule does no worse than any other method. On the other hand, if there is a limiting relative frequency, then the straight rule is guaranteed to converge to it in the long run: the rule's prediction is guaranteed eventually to come within any given degree of approximation to the limiting relative frequency, and thenceforth to remain within that range. That is because by definition, L is the limit of the sequence a_1, a_2, a_3,... exactly when for any small positive number ε, there is an integer N such that for any $n > N, a_n$ is within ε of L. So if L is the limit of

[26] Salmon [1991, p. 100] takes himself to be following Reichenbach in arguing that the policy of making no predictions fails whether nature is uniform or not, so it cannot be better than using induction, since the worst that induction can do is to fail. But although the no-prediction rule fails in making successful predictions, it succeeds in not making unsuccessful predictions.

the sequence $a_1, a_2, a_3, \ldots$ where a_n is the fraction (m/n) of Gs among the first n Fs, then at some point in the sequence, its members come and thenceforth remain within ε of L, and since at any point the straight rule's prediction of the limit is just the current member m/n of the sequence, the straight rule's prediction of the limit is guaranteed eventually to come and thenceforth to remain within ε of L. If instead our goal is the accurate estimation of the objective chance of an F's being G, then if there is such a chance, the straight rule is 100% likely to arrive at it — to within any specified degree of approximation — in the long run. (That is, although a fair coin *might* land heads repeatedly, the likelihood of its landing heads about half of the time becomes arbitrarily high as the number of tosses becomes arbitrarily large.)

So the straight rule is "asymptotic": its prediction is guaranteed to converge to the truth in the long run, if any rule will. Surely, Reichenbach seems to be suggesting, it would be irrational to knowingly employ a rule that is not asymptotic if an asymptotic rule is available. Plenty of rules are not asymptotic. For instance, consider the "counterinductive rule": If n Fs have been examined and m have been found to be G, then take $(1-m/n)$ to approximate the actual fraction of Gs among Fs. The counterinductive rule is not asymptotic since, unless the limit is 50%, its prediction is guaranteed to diverge from the straight rule's in the long run, and the straight rule's is guaranteed to converge to the truth (if there is a truth to converge to) in the long run. In this way, Reichenbach's argument purports to justify following the straight rule rather than the counterinductive rule. (Notice how this argument aims to walk a very fine line: to justify induction without saying anything about induction's likelihood of leading us to the truth!)

This argument does not rule out another rule's working better than the straight rule *in the short run* — that is, converging more quickly to the relative limiting frequency than the straight rule does. For instance, the rule that would have us (even before we have ever tossed a coin!) guess 50% as the approximate limiting relative frequency of heads among the coin-toss outcomes might happen to lead to the truth right away. So in this respect, it might do better than the straight rule. But it cannot do better than the straight rule in the long run, since the straight rule is guaranteed to lead to the truth in the long run. (The 50% rule is not *guaranteed* to lead to the truth in the long run.)

However, in the long run (as Keynes famously quipped), we are all dead! All of our predictions are made in the short run — after a finite number of observations have been made. Why should a rule's success in the long run, no matter how strongly guaranteed, do anything to justify our employing it in the short run? We cannot know how many cases we need to accumulate before the straight rule's prediction comes and remains within a certain degree of approximation of the genuine limit (if there is one). So the fact that the straight rule's prediction is guaranteed to converge *eventually* to the limit (if there is one) seems to do little to justify our being guided by the straight rule in the short run. Why should the straight rule's success under conditions that we have no reason to believe we currently (or will ever!) occupy give us a good reason to use the straight rule?

(This seems to me closely related to BonJour's objection to Reichenbach.)

A final problem for Reichenbach's argument is perhaps the most serious. A nondenumerably infinite number of rules are entitled to make the same boast as the straight rule: each is guaranteed to converge in the long run to the limiting relative frequency if one exists. Here are a few examples:

> If n Fs have been examined and m have been found to be G, then take $m/n + k/n$ for some constant k (or, if this quantity exceeds 1, then take 1) to equal (within a certain degree of approximation) the actual fraction of Gs among Fs.
>
> If n Fs have been examined and m have been found to be G, then take the actual fraction of Gs among Fs to equal 23.83%, if $n < 1,000,000$, or m/n, otherwise.
>
> If n Fs have been examined and m have been found to be G, and among the first 100 Fs examined, M were found to be G, then take the actual fraction of Gs among Fs to equal $(m+M)/(n+ \inf\{n, 100\})$. (In other words, "double count" the first 100 Fs.)

In the long run, each of these rules converges to the straight rule and so must converge to the limit (if there is one). The straight rule cannot be shown to converge faster than any other asymptotic rule. Moreover, the asymptotic rules disagree to the greatest extent possible in their predictions: for any evidence and for any prediction, there is an asymptotic rule that endorses making that prediction on the basis of that evidence.[27]

The first of these three rivals to the straight rule violates the constraint that if G and H are mutually exclusive characteristics and a rule endorses taking p as Gs relative frequency and q as Hs relative frequency among Fs, then the rule should endorse taking $(p + q)$ as the relative frequency of (G or H). Furthermore, the second and third of these three rivals to the straight rule violate the constraint that the rule endorse taking the same quantity as the limiting relative frequency for any sequence of Gs and $\sim G$s with a given fraction of Gs and $\sim G$s, no matter how long the sequence or in what order the Gs and $\sim G$s appear in it. Constraints like these have been shown to narrow down the asymptotic rules to the straight rule alone. [Salmon, 1967, pp. 85–89, 97–108; Hacking, 1968, pp. 57–58] However, it is difficult to see how to justify such constraints without begging the question. For instance, if the sequence consists of the Fs in the order in which we have observed them, then to require making the same prediction regardless of the order (as long as the total fraction of Gs is the same) is tantamount to assuming that

[27] Reichenbach [1938, pp. 353–354] notes this problem. In [Reichenbach, 1949a, p. 447], he favors the straight rule on grounds of "descriptive simplicity". But although Reichenbach regards "descriptive simplicity" as relevant for selecting among empirically equivalent theories, rules are not theories. In any case, the rival rules do not endorse all of the same predictions in the short run, so they are not "equivalent" there. Of course, in the long run, they are "equivalent", but why is that fact relevant?

later Fs are no different from earlier Fs — that the future is like the past, that each F is like an independent flip of the same coin as every other F (i.e., that each F had the same objective chance of being G).

9 BAYESIAN APPROACHES

Suppose it could be shown — perhaps by a Dutch Book argument or an argument from calibration [Lange, 1999] — that rationality obliges us (in typical cases) to update our opinions by Bayesian conditionalization (or some straightforward generalization thereof, such as Jeffrey's rule). This would be the kind of argument that Hume fails to rule out (as I explained in section 3): an argument that is *a priori* (despite not turning solely on semantic relations, since a degree of belief is not capable of being true or false) and that proceeds from the opinions that constitute a given inductive argument's premises to the degree of belief that constitutes its conclusion.

Such an argument would still be far from a justification of induction. Whether Bayesian conditionalization yields induction or counterinduction, whether it underwrites our ascribing high probability to "All emeralds are green" or to "All emeralds are grue," whether it leads us to regard a relatively small sample of observed emeralds as having any bearing at all on unexamined emeralds — all depend on the "prior probabilities" plugged into Bayesian conditionalization along with our observations. So in order to explain why we ought to reason *inductively* (as an adequate justification of induction must do — see section 2 above), the rationality of Bayesian conditionalization would have to be supplemented with some constraints on acceptable priors.

This argument has been challenged in several ways. Colin Howson [2000] argues that a justification of induction should explain why *we* ought to reason inductively. Thus, it can appeal to *our* prior probabilities; Bayesian conditionalization, acting on these particular priors, underwrites recognizably *inductive* updating. The "initial assignments of positive probability ... cannot themselves be justified in any absolute sense". [Howson, 2000, p. 239] But never mind, Howson says. Inductive arguments are in this respect

> like sound deductive arguments, they don't give you something for nothing: you must put synthetic judgements in to get synthetic judgements out. But get them out you do, and in a demonstrably consistent way that satisfies certainly the majority of those intuitive criteria for inductive reasoning which themselves stand up to critical examination. [Howson, 2000, p. 239, see also p. 171]

All we really want from a justification of induction is a justification for *updating* our beliefs in a certain way, and that is supplied by arguments showing Bayesian conditionalization to be rationally compulsory. As Frank Ramsey says,

> We do not regard it as belonging to formal logic to say what should be a man's expectation of drawing a white or black ball from an urn; his

> original expectations may within the limits of consistency be any he likes, all we have to point out is that if he has certain expectations, he is bound in consistency to have certain others. This is simply bringing probability into line with ordinary formal logic, which does not criticize premisses but merely declares that certain conclusions are the only ones consistent with them. [Ramsey, 1931, p. 189]

Ian Hacking puts the argument thus:

> At any point in our grown-up lives (let's leave babies out of this) we have a lot of opinions and various degrees of belief about our opinions. The question is not whether these opinions are 'rational'. The question is whether we are reasonable in modifying these opinions in light of new experience, new evidence. [Hacking, 2001, p. 256]

But the traditional problem of induction is whether by reasoning inductively, we arrive at knowledge. If knowledge involves justified true belief, then the question is whether true beliefs arrived at inductively are thereby justified. And if an inductive argument, to justify its conclusion, must proceed from a prior state of opinion that we are entitled to occupy, then the question becomes whether we are entitled to those prior opinions, and if so, how come.

I said that Bayesian conditionalization can underwrite reasoning that is intuitively inductive, but with other priors plugged into it, Bayesian conditionalization underwrites reasoning that is counterinductive or even reasoning that involves the confirmation of no claims at all regarding unexamined cases. However, it might be objected that if hypothesis h (given background beliefs b) logically entails evidence e, then as long as $pr(h|b)$ and $pr(e|b)$ are both non-zero, it follows that $pr(e|h\&b) = 1$, and so by Bayes's theorem, we have $pr(h|e\&b) = pr(h|b)pr(e|h\&b)/pr(e|b) = pr(h|b)/pr(e|b) > pr(h|b)$, so by Bayesian conditionalization, e confirms h. On this objection, then, Bayesian conditionalization automatically yields induction.

However, this confirmation of h (of "All emeralds are green," for example) by e ("The emerald currently under examination is green") need not involve any *inductive* confirmation of h — roughly, any confirmation of h's *predictive accuracy*. For example, it need not involve any confirmation of g: "The next emerald I examine will turn out to be green." Since g (given b) does not logically entail e, $pr(e|g\&b)$ is not automatically 1, and so $pr(g|e\&b)$ is not necessarily greater than $pr(g|b)$.

Howson insists that it is no part of the justification of induction to justify the choice of priors, just as deductive logic does not concern itself with justifying the premises of deductive arguments. [Howson, 2000, p. 2, see also pp. 164, 171, 239; cf. Howson and Urbach, 1989, pp. 189–190] To my mind, this parallel between deduction and induction is inapt. It presupposes that prior probabilities are the premises of inductive arguments — are, in other words, the neutral input or substrate to which is applied Bayesian conditionalization, an inductive rule of inference. But, as Howson rightly emphasizes, it is Bayesian conditionalization that is neutral; anything distinctively "inductive" about an episode of Bayesian updating must come from the priors. Consequently, a justification of induction must say something about how we are entitled to those priors.

Some personalists about probability have argued that if anything distinctively "inductive" about Bayesian updating must come from the priors, then so much the better for resolving the problem of induction, since we are automatically entitled to adopt *any* coherent probability distribution as our priors. This view is prompted by the notorious difficulties (associated with Bertrand's paradox of the chord) attending any principle of indifference for adjudicating among rival priors. For example, Samir Okasha writes:

> 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 guess work is involved in all empirical enquiry.[Okasha, 2001, p. 322]

Okasha's argument seems to be that a prior state of opinion embodies no unjustified information about the world since any prior opinion embodies *some* information. But the inductive sceptic should reply by turning this argument around: Since any prior opinion strong enough to support an inductive inference embodies some information, no prior opinion capable of supporting an inductive inference is justified.

In other words, Okasha's argument seems to be that there are no objectively neutral priors, so if the inductive sceptic accuses our priors of being unjustified,

> we need only ask the sceptic 'What prior probability do you recommend?' [...] 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. [Okasha, 2001, p. 323]

But there *is* an alternative to operating from a prior opinion strong enough to support the confirmation of predictions. If the sceptic is asked to recommend a prior probability, she should suggest a distribution that makes no probability assignment at all to any prediction about the world that concerns logically contingent matters of fact. By this, I do not mean the extremal assignment of *zero* subjective probability to such a claim. That would be to assign it a probability: zero. Nor do I mean assigning it a *vague* probability value. I mean making no assignment at all to any such claim. According to the inductive sceptic, there is *no* degree of confidence to which we are entitled regarding predictions regarding unexamined cases.[28]

[28]Admittedly, the sceptic's prior distribution violates the requirement that the domain of a

Though an observation's direct result may be to assign some probability to e, the sceptic's prior distribution fails to support inductive inferences from our observations (since it omits some of the probabilities required by Bayesian conditionalization or any generalization of it). But that is precisely the inductive sceptic's point. There is no alternative to operating with a prior distribution that embodies information about the world, as Okasha says, *if* we are going to use our observations to confirm predictions. But to presuppose that we are justified in using our observations to confirm predictions is obviously to beg the question against the inductive sceptic.

10 WILLIAMS' COMBINATORIAL JUSTIFICATION OF INDUCTION

In 1947, Donald Williams offered an *a priori* justification of induction that continues to receive attention [Williams, 1947; Stove, 1986, pp. 55–75]. The first ingredient in Williams' argument is a combinatorial fact that can be proved *a priori*: if there is a large (but finite) number of Fs, then in most collections of Fs beyond a certain size (far smaller than the total population of Fs), the fraction of Gs is close to the fraction of Gs in the total F population. For example, consider a population of 24 marbles, of which 16 (i.e., 2/3) are white. The number of 6-member sets that are also 2/3 white (i.e., 4 white, 2 non-white) is [(16 x 15 x 14 x 13)/(4 x 3 x 2)] x [(8 x 7)/2] = 40,960. The number of sets containing 5 white and 1 non-white marbles is 34,944, and the number containing 3 white and 3 non-white marbles is 31,360. So among the 134,596 6-member sets, about 80% contain a fraction of white marbles within 1 marble (16.7%) of the fraction of white marbles in the overall population. For a marble population of any size exceeding one million, more than 90% of the possible 3000-member samples have a fraction of white marbles within 3% of the overall population's fraction, no matter what that fraction is (even if no sample has a fraction exactly matching the overall population's). Notice that this is true no matter how small the sample may be as a fraction of the total population. [Williams, 1947, p. 96; Stove, 1986, p. 70]

The second ingredient in Williams' argument is the rationality of what he calls the "statistical [or proportional] syllogism": if you know that the fraction of As that are B is r and that individual a is A, then if you know nothing more about a that is relevant to whether it is B, then r is the rational degree of confidence for you to have that a is B. Williams regards this principle as an *a priori* logical truth: "the native wit of mankind ... has found the principle self-evident." [Williams,

probability function be a sigma algebra. For example, it may violate the additivity axiom [$pr(q$ or $\sim q) = pr(q) + pr(\sim q)$] by assigning to ($q$ or $\sim q$) a probability of 1 but making no probability assignment to q and none to $\sim q$. Some Bayesians would conclude that the sceptic's "pr" does not qualify as a probability function. However, the sceptic is not thereby made vulnerable to a Dutch Book. She is not thereby irrational or incoherent. What is the worst thing that can be said of her? That she shows a certain lack of commitment. That characterization will hardly bother the sceptic! It may well be overly restrictive to require that the domain of a probability function be a sigma algebra. (See, for instance, [Fine, 1973, p. 62] or any paper discussing the failure of logical omniscience.)

1947, p. 8] It appears to be far enough from the inductive arguments in question that no circularity is involved in justifying them by appealing to it. A statistical syllogism would be another argument of the kind that Hume fails to rule out (as I explained in section 3): an argument that is *a priori* (despite not turning solely on semantic relations, since a degree of belief is not capable of being true or false) and that proceeds from the opinions constituting a given inductive argument's premises to the degree of belief that constitutes its conclusion.

But how can we use a statistical syllogism to justify induction? Let the As be the large samples of Fs and let B be the property of having a fraction of Gs approximating (to whatever specified degree) the fraction in the overall F population. Since it is a combinatorial fact that the fraction of As that are B is large, it follows (by the statistical syllogism) that in the absence of any evidence to the contrary, we are entitled to have great confidence that the fraction of Gs in a given large sample of Fs approximates the fraction of Gs among all Fs.[29] So if f is the fraction of Gs in the observed large sample of Fs, then we are entitled (in the absence of any countervailing evidence) to have great confidence that f is the fraction of Gs among all Fs. As Williams says,

> Without knowing exactly the size and composition of the original population to begin with, we cannot calculate ... exactly what proportion of our "hyper-marbles" [i.e., large samples of marbles] have the quality of nearly-matching-the-population, but we do know *a priori* that most of them have it. Before we choose one of them, it is hence very probable that the one we choose will be one of those which match or nearly match; after we have chosen one, it remains highly probable that the population closely matches the one we have, so that we need only look at the one we have to read off what the population, probably and approximately, is like. [Williams, 1947, pp. 98–99]

Induction is thereby justified.

However, we might question whether the statistical syllogism is indeed a principle of good reasoning.[30] It presupposes that we assign every possible large sample the same subjective probability of being selected (so since there are overwhelmingly more large samples that are representative of the overall population, we are overwhelmingly confident that the actual sample is representative). Why is this equal-confidence assignment rationally obligatory? The intuition seems to be that

[29] This argument, as used to infer from a population's fraction of Gs to a sample's likely fraction, is often called "direct inference", and hence the statistical syllogism is termed "the principle of direct inference" [Carnap, 1950]. An inference in the other direction, from sample to population, is then termed "inverse reasoning". That probably a sample's fraction of Gs approximates the population's fraction can apparently take us in either direction.

[30] It cannot be justified purely by Bayesian considerations. If P is that the fraction of Gs in the population is within a certain range of f, and S is that f is the fraction of Gs in the large observed sample, then Bayes's theorem tells us that $pr(P|S) = pr(P)\ pr(S|P)/pr(S)$, where all of these probabilities are implicitly conditional on the size of the sample. It is unclear how to assign the priors $pr(S)$ and $pr(P)$. Even $pr(S|P)$ does not follow purely from combinatorial considerations.

when we have no reason to assign any of these samples greater subjective probability than any other, we ought to assign them equal subjective probabilities. To do otherwise would be irrational. But perhaps, in the absence of any relevant information about them, we have no reason to assign any subjective probabilities to any of them.

In other words, the motivation for the equal-confidence assignment seems to be that if we have no relevant information other than that most marbles in the urn (or most possible samples) are red (or representative), then it would be irrational to be confident that a non-red marble (or unrepresentative sample) will be selected. But this undoubted fact does not show that it would be rational to expect that a red (or representative) one will be selected. Perhaps we are not entitled to any expectation unless we have further information, such as that the marble (or sample) is selected randomly (i.e., that every one has an equal objective chance of being selected). Williams is quite correct in insisting that in the absence of any relevant information, we are not entitled to believe that the sample is selected randomly. So why are we entitled to other opinions in the absence of any relevant information?

Furthermore, we do have further information — for example, that samples with members that are remote from us in space or time will not be selected. This information does not suggest that the sample we select is unrepresentative — *if* we believe that Fs are uniform in space and time. But we cannot suppose so without begging the question.[31]

Since Williams's argument is purely formal, we could apparently just as well take G as green or as grue while taking the Fs as the emeralds. But we cannot do both on pain of being highly confident both that all emeralds are green and that all emeralds are grue. If we regard the fact that all of the emeralds are sampled before the year 3000 as further information suggesting that the sample may be unrepresentative, then neither hypothesis is supported.[32]

Finally, is there any reason to believe that statistical syllogisms will lead us to place high confidence in truths more often than in falsehoods (or, at least, that they have a high objective chance of doing so)? If, in fact, our samples are unrepresentative in most cases where we have no other relevant information, then statistical syllogisms will lead us to place high confidence in falsehoods more often than in truths. That we have no good reason to think that a sample is unrepresentative does not show that we are likely to reach the truth if we presume

[31] For defense of the statistical syllogism, see [Williams, 1947, pp. 66–73 and p. 176; Carnap, 1959, p. 494; McGrew, 2001, pp. 161–167]. Maher [1996] has argued that the fraction of Gs in the sample may suggest that the sample is unrepresentative and so undercut the statistical syllogism. How are we entitled *a priori* to have the opinion that a sample with a given fraction of Gs is no less likely if it is representative than if it is unrepresentative?

[32] Stove [1986, pp. 140–142] says that he is not committed to reasoning with green in the same way as with grue, but he identifies no *a priori* ground for privileging one over the other. Campbell [1990], who endorses Williams's argument, says that "a complex array of higher-order inductions, about natural kinds, about discontinuities in nature, and about the kinds of properties it is significant to investigate" justify privileging green over grue. But aren't these inductions also going to be subject to versions of the grue problem?

it to be representative unless we believe that if it were unrepresentative, we would probably have a good reason to suspect so. But why should we believe that we would be so well-informed?[33]

11 THE INDUCTIVE LEAP AS MYTHICAL

Hume's argument appears to show that our observations, apart from any theoretical background, are unable to confirm or to disconfirm any predictions. Supplemented by different background opinions (whether understood as prior probabilities or a uniformity principle), the same observations exert radically different confirmatory influences. But then to be warranted in making any inductive leap beyond the safety of our observations, we must be justified in holding some opinions regarding some prediction's relation to our observations. These opinions may rest on various scientific theories, which in turn have been confirmed by other observations. But when we pursue the regress far enough — or look at cases where we have very little relevant background knowledge — how can any such opinions be justified?

My view [Lange, 2004] is that especially in theoretically impoverished circumstances, the content of our observation reports may include certain expectations regarding the observations' relations to as yet undiscovered facts — expectations enabling those observations to confirm predictions regarding those facts. An observation report ("That is an F") classifies something as belonging to a certain category (F). That category may be believed to be a "natural kind" of a certain sort (e.g., a species of star, mineral, animal, disease, chemical...). Part of what it is for a category to be a natural kind of a given sort is for its members generally to be alike in certain respects. (These respects differ for different sorts of natural kind.) Members of the same species of star, for instance, are supposed to be generally similar in temperature, intrinsic luminosity, the mechanism by which they generate light, and so forth (and generally different in many of these respects from members of other star species). Therefore, to observe that certain stars are (for instance) Cepheid-type variables, an astronomer must be justly prepared (in the absence of any further information) to regard the examined Cepheids' possession of various properties (of the sorts characteristic of natural kinds of stars) as confirming that unexamined Cepheids possess those properties as well. In that case, Cepheid observations (in the absence of any other evidence) suffice to justify astronomers in expecting that unexamined Cepheids will exhibit, say, a certain simple period-luminosity relation that examined Cepheids display. No opinions *independent* of Cepheid observations must be added to them in order to give them the capacity to confirm certain predictions. Hence, there arises no regress-inducing problem of justifying some such independent opinions.

[33] I made an analogous point regarding BonJour's *a priori* justification of induction. McGrew [2001, pp. 167–170] responds to this criticism of Williams's argument. Kyburg [1956] argues that even if we knew that some method of inductive reasoning would more often lead to truth than to falsehoods in the long run, we could not justify using it except by a statistical syllogism.

To observe that certain stars are Cepheids, an astronomer must already have the resources for going beyond those observations. As Wilfrid Sellars says, in arguing for a similar point:

> The classical 'fiction' of an inductive leap which takes its point of departure from an observation base undefiled by any notion as to how things hang together is not a fiction but an absurdity. ... [T]here is no such thing as the problem of induction if one means by this a problem of how to justify the leap from the safe ground of the mere description of particular situations, to the problematic heights of asserting lawlike sentences and offering explanations.[Sellars, 1963b, p. 355]

I could not make observations categorizing things into putative natural kinds if I were not justified in expecting the members of those kinds to be alike in the respects characteristic of such kinds.

The observations that we most readily become entitled to make in theoretically impoverished contexts are (perhaps paradoxically) precisely those with inductive import — those purporting to classify things into various sorts of natural kinds. That is because although an observation report ("That is an *F*") has noninferential justification, I am justified in making it only if I can justly infer, from my track record of accuracy in making similar responses in other cases, that my report on this occasion is probably true [Sellars, 1963a, pp. 167–170]. If I have no reason to trust myself — if I am not in a position to infer the probable accuracy of my report — then on mature reflection, I ought to disavow the report, regarding it as nothing more than a knee-jerk reaction. But obviously, this inference from my past accuracy in making similar responses is inductive. I am entitled to regard my past successes at identifying *F*s as confirming that my latest *F* report is accurate only because I am entitled to expect that generally, unexamined *F*s look like the *F*s that I have accurately identified in the past and look different from various kinds of non-*F*s. (Only then is my past reliability in distinguishing *F*s good evidence — in the absence of countervailing information — for my future reliability in doing so.) That is just what we expect when the *F*s form a natural kind. My past reliability at identifying *F*s, where *F*s form a natural kind, confirms (in the absence of countervailing information) my future reliability without having to be supplemented by further regress-inducing background opinions. Accordingly, "taxonomic observations" (i.e., identifications of various things as members of various species) are among the sorts of observations that scientists are most apt to be in a position to make in a new field, where their background theory is impoverished. Of course, there is no logical guarantee that these putative observations are accurate. But one becomes qualified to make them precisely because of — rather than despite — the "thickness" of their content.

If observers in theoretically impoverished contexts had not taken the *F*'s as forming a certain sort of natural kind, then the range of cases in which those observers are justified in making "That is *F*" reports would have been different. For example, suppose astronomers had taken the Cepheid-type variable stars as consisting simply of all and only those variables having light curves shaped like those of the two prototype Cepheids (delta Cephei and epsilon Aquilae). Then astronomers would have classified as non-Cepheids certain stars that they actually deemed to be Cepheids (and would have deemed other stars to be "somewhat Cepheid" or "Cepheid-like", categories that are never used). Instead astronomers took the Cepheid category as extending from the prototypical Cepheids out to wherever the nearest significant gap appears in the distribution of stellar light curves. That is because they understood the Cepheids to be a natural kind of star sharply different from any other kind. A taxonomic observation report (such as "That is a Cepheid") embodies expectations regarding as yet undiscovered facts, and these expectations — which ground the most basic inductive inferences made from those observations — are inseparable from the reports' circumstances of application. The content of the observation reports cannot be "thinned down" so as to remove all inductive import without changing the circumstances in which the reports are properly made.

Hume's problem of induction depends on unobserved facts being "loose and separate" (*E*, p. 49) from our observational knowledge. But they are not.

12 CONCLUSION

Having surveyed some of most popular recent responses to Hume's argument (and having, in the previous section, bravely sketched the kind of response I favor), I give the final word to Hume:

> Most fortunately it happens, that since reason is incapable of dispelling these clouds, nature herself suffices to that purpose, and cures me of this philosophical melancholy and delirium, either by relaxing this bent of mind, or by some avocation, and lively impression of my senses, which obliterates all these chimeras. I dine, I play a game of backgammon, I converse, and am merry with my friends; and when after three or four hour's amusement, I wou'd return to these speculations, they appear so cold, and strain'd, and ridiculous, that I cannot find in my heart to enter into them any farther. Here then I find myself absolutely and necessarily determin'd to live, and talk, and act like other people in the common affairs of life. (*T*, p. 269)

Most fortunately for us, Hume did not act like other people in all affairs of life. Rather, he bequeathed to us an extraordinary problem from which generations of philosophers have derived more than three or four hours' amusement. I, for one, am very grateful to him.

BIBLIOGRAPHY

[Beauchamp and Rosenberg, 1981] T. Beauchamp and A. Rosenberg. *Hume and the Problem of Causation*. Oxford University Press, New York, 1981.

[Black, 1954] M.Black. Inductive support of inductive rules. In Black, *Problems of Analysis*. Cornell University Press, Ithaca, NY, pp. 191–208, 1954.

[BonJour, 1986] L. BonJour. A reconsideration of the problem of induction. *Philosophical Topics* 14, pp. 93—124, 1986.

[BonJour, 1998] Laurence BonJour. *In Defense of Pure Reason*. Cambridge University Press, Cambridge, 1998.

[Brandom, 1994] R. Brandom. *Making it Explicit*. Harvard University Press, Cambridge, MA, 1994.

[Broad, 1952] C.D. Broad. *Ethics and the History of Philosophy*. Routledge and Kegan Paul, London, 1952.

[Brueckner, 2001] A. Brueckner. BonJour's a priori justification of induction. *Pacific Philosophical Quarterly* 82, pp. 1–10, 2001.

[Butler, 1813] J. Butler. Analogy of religion, natural and revealed. In Butler, *The Works of Joseph Butler*, volume 1. William Whyte, Edinburgh, 1813.

[Campbell, 1990] K. Campbell. *Abstract Particulars*. Blackwell, Oxford, 1990.

[Carnap, 1950] R. Carnap. *Logical Foundations of Probability*. University of Chicago Press, Chicago 1950.

[Dummett, 1981] M. Dummett. *Frege: Philosophy of Language*, 2^{nd} ed. Harvard University Press, Cambridge, MA, 1981.

[Feigl, 1950] H. Feigl. De principiis non disputandum. . . ?. In M. Black (ed.), *Philosophical Analysis*. Cornell Unniversity Press, Ithaca, pp. 119—156, 1950.

[Fine, 1973] T. Fine. *Theories of Probability*. Academic, New York, 1973.

[Franklin, 1987] J. Franklin. Non-deductive logic in mathematics. *British Journal for the Philosophy of Science* 38, 1—18, 1987.

[Garrett, 1997] D. Garrett. *Cognition and Commitment in Hume's Philosophy*. Oxford University Press, New York, 1997.

[Goodman, 1954] N. Goodman. *Fact, Fiction and Forecast*. Harvard University Press, Cambridge, MA, 1954.

[Hacking, 1968] I. Hacking. One problem about induction. In I. Lakatos (ed.), *The Problem of Induction*. North-Holland, Amsterdam, pp. 44—59, 1968.

[Hacking, 2001] I. Hacking. *An Introduction to Probability and Inductive Logic*. Cambridge University Press, Cambridge, 2001.

[Harman, 1965] G. Harman. Inference to the best explanation. *Philosophical Review* 74, pp. 88–95, 1965.

[Horwich, 1982] P. Horwich. *Probability and Evidence*. Cambridge University Press, Cambridge, 1982.

[Howson, 2000] C. Howson. *Hume's Problem: Induction and the Justification of Belief*. Clarendon, Oxford, 2000.

[Howson and Urbach, 1989] C. Howson and P. Urbach. *Scientific Reasoning: The Bayesian Approach*. Open Court, La Salle, IL, 1989.

[Hume, 1977] D. Hume. *An Enquiry Concerning Human Understanding*, ed. Eric Steinberg. Hackett, Indianapolis, 1977.

[Hume, 1978] D. Hume. *A Treatise of Human Nature*, ed. L.A. Selby-Bigge and P.H. Nidditch, 2^{nd} ed. Clarendon, Oxford, 1978.

[Keynes, 1921] J. M. Keynes. *A Treatise on Probability*. Macmillan, London, 1921.

[Kornblith, 1993] H. Kornblith. *Inductive Inference and its Natural Ground*. MIT Press, Cambridge, MA, 1993.

[Kyburg, 1956] H. Kyburg. The justification of induction. *Journal of Philosophy* 53, pp. 394–400, 1956.

[Lange, 1999] M. Lange. Calibration and the epistemological role of Bayesian conditionalization. *Journal of Philosophy* 96, pp. 294–324.

[Lange, 2004] M. Lange. Would direct realism resolve the classical problem of induction? *Nous* 38, pp. 197–232, 2004.

[Lipton, 1991] P. Lipton. *Inference to the Best Explanation*. Routledge, London, 1991.

[Mackie, 1980] J.L. Mackie. *The Cement of the Universe*. Clarendon, Oxford, 1980.

[Maher, 1996] P. Maher. The hole in the ground of induction. *Australasian Journal of Philosophy* 74, 423–432, 1996.

[McGrew, 2001] T. McGrew. Direct inference and the problem of induction. *The Monist* 84, pp. 153-78, 2001.

[Mill, 1872] J.S. Mill. *A System of Logic*, 8^{th} ed. Longmans, London, 1872.
[Norton, 2003] J. Norton. A material theory of induction. *Philosophy of Science* 70, pp. 647–70, 2003.
[Okasha, 2001] S. Okasha. What did Hume really show about induction? *The Philosophical Quarterly* 51, pp. 307–327, 2001.
[Papineau, 1993] D. Papineau. *Philosophical Naturalism*. Blackwell, Oxford, 1993.
[Popper, 1959] K. Popper. *The Logic of Scientific Discovery*. Basic Books, New York, 1959.
[Popper, 1972] K. Popper. Conjectural knowledge: my solution to the problem of induction. In Popper, *Objective Knowledge*. Clarendon, Oxford, pp. 1—31, 1972.
[Putnam, 1994] H. Putnam . Reichenbach and the limits of vindication. In Putnam, *Words and Life*. Harvard University Press, Cambridge, MA, pp. 131—148, 1994.
[Ramsey, 1931] F. Ramsey. Truth and probability. In Ramsey, *The Foundations of Mathematics and Other Logical Essays*. Routledge and Kegan Paul, London, pp. 156—198, 1931.
[Read and Richman, 2000] R.J. Read and K.A. Richman. *The New Hume Debate*. Routledge, London, 2000.
[Reichenbach, 1938] H. Reichenbach. *Experience and Prediction*. University of Chicago Press, Chicago, 1938.
[Reichenbach, 1949a] H. Reichenbach. *The Theory of Probability*. University of California Press, Berkeley, 1949.
[Reichenbach, 1949b] H. Reichenbach. Comments and criticism. *Journal of Philosophy* 46, pp. 545–549, 1949.
[Reichenbach, 1968] H. Reichenbach. *The Rise of Scientific Philosophy*. University of California Press, Berkeley, 1968.
[Rhees and Phillips, 2003] R. Rhees and D.Z. Phillips. *Wittgenstein's On Certainty: There — Like our Life*. Blackwell, Oxford, 2003.
[Russell, 1919] B. Russell. *Introduction to Mathematical Philosophy*. George Allen & Unwin, London, 1919.
[Russell, 1948] B. Russell. *Human Knowledge: Its Scope and Limits*. Simon and Schuster, New York, 1948.
[Russell, 1959] B. Russell. *The Problems of Philosophy*. Oxford University Press, London, 1959.
[Salmon, 1957] W. Salmon. Should we attempt to justify induction? *Philosophical Studies* 8, pp. 33–48, 1957.
[Salmon, 1963] W. Salmon. Inductive inference. In B. Baumrin (ed.), *Philosophy of Science: The Delaware Seminar*, volume II. Interscience Publishers, New York and London, pp. 353-70, 1963.
[Salmon, 1967] W. Salmon. *The Foundations of Scientific Inference*. University of Pittsburgh Press, Pittsburgh, 1967.
[Salmon, 1981] W. Salmon. Rational prediction. *British Journal for the Philosophy of Science* 32, pp. 115—25, 1981.
[Salmon, 1984] W. Salmon. *Scientific Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, 1984.
[Salmon, 1991] W. Salmon. Hans Reichenbach's vindication of induction. *Erkenntnis* 35, pp. 99–122, 1991.
[Salmon, Barker, and Kyburg, 1965] W. Salmon, S. Barker, and H. Kyburg, Jr. Symposium on inductive evidence. *American Philosophical Quarterly* 2, pp. 265–80, 1965.
[Sankey, 1997] H. Sankey. Induction and natural kinds. *Principia* 1, pp. 239—54, 1997.
[Sellars, 1963a] W. Sellars. Empiricism and the philosophy of mind. In Sellars, *Science, Perception and Reality*. Routledge and Kegan Paul, London, pp. 127—196, 1963.
[Sellars, 1963b] W. Sellars. Some reflections on language games. In Sellars, *Science, Perception and Reality*. Routledge and Kegan Paul, London, pp. 321–358, 1963.
[Skyrms, 1986] B. Skyrms. *Choice and Chance*, 3^{rd} ed. Wadsworth, Belmont, CA, 1986.
[Smith, 1941] N. K. Smith. *The Philosophy of David Hume*. Macmillan, London.
[Sober, 1988] E. Sober. *Reconstructing the Past*. Bradford, Cambridge, MA, 1988.
[Stove, 1965] D. Stove. Hume, probability, and induction. *Philosophical Review* 74, pp. 160—77, 1965.
[Stove, 1973] D. Stove. *Probability and Hume's Inductive Scepticism*. Clarendon, Oxford, 1973.
[Stove, 1986] D. Stove. *The Rationality of Induction*. Oxford University Press, Oxford, 1986.
[Strawson, 1952] P.F. Strawson. *An Introduction to Logical Theory*. Methuen, London, 1952.

[Strawson, 1958] P.F. Strawson. On justifying induction. *Philosophical Studies* 9, pp. 20—21. 1958.

[Stroud, 1977] B. Stroud. *Hume*. Routledge, London and New York, 1977.

[Thagard, 1978] P. Thagard. The best explanation: criterion for theory choice. *Journal of Philosophy* 75, pp. 76–92, 1978.

[van Cleve, 1984] J. van Cleve. Reliability, justification, and the problem of induction. In P. French, T. Uehling, and H. Wettstein (eds.), *Midwest Studies in Philosophy* IX. University of Minnesota Press, Minneapolis, pp. 555–567, 1984.

[van Fraassen, 1981] B. van Fraassen. *The Scientific Image*. Clarendon, Oxford, 1981.

[van Fraassen, 1989] B. van Fraassen. *Laws and Symmetry*. Clarendon, Oxford, 1989.

[Williams, 1947] D.C. Williams. *The Ground of Induction*. Harvard University Press, Cambridge, MA, 1947.

[Wittgenstein, 1953] L. Wittgenstein. *Philosophical Investigations.* Blackwell, Oxford, 1953.