

7

The justification of induction

§7.1. INTRODUCTION

Our discussion of *a priori* justification so far has been in the main relentlessly abstract, with only a few of the most obvious examples to enliven the way. While this seems to me appropriate where it is the very existence of non-tautological *a priori* justification that is at issue, it does leave the issue of the scope of *a priori* justification almost entirely unilluminated. For all that has been argued so far, it would be possible that a priori justification of the rationalist kind, though genuinely existent, is confined entirely to the general kinds of examples discussed in §4.2. And if this were so, then such justification, though perhaps important in these limited areas, would have little significance for human knowledge in general and would in particular do almost nothing to solve the problem of observation-transcendent inference raised in §1.1. Radical empiricists would indeed be mistaken in their central claim, but their error would be of little consequence; their general epistemological position would still be closer to the truth than that of the rationalist in the ways that matter most.

My conviction is that, on the contrary, rationalistic *a priori* justification is of crucial importance for epistemology and indeed for philosophy generally. While a full defense of this claim would

be as large as philosophy itself and would greatly transcend the scope of this or any reasonable book, the aim of this final chapter is to make a start in this direction. In this chapter, I will offer a more specific and detailed (though still far from complete) discussion of one central epistemological problem, to the solution of which an *a priori* appeal is arguably essential: the classical Humean problem of induction.

I choose this problem for extended treatment because it is obviously central to the general issue of observation-transcendent inference. Induction is the intuitively simplest example of an inference that transcends direct observation, and inductively arrived at conclusions also provide the essential basis for many inferences of more complicated sorts, including, I would argue, the inference to the external world. Thus it is plausible to suppose that any adequate non-skeptical epistemology must be able to offer a justification of induction. I will argue in this chapter, first, that only an *a priori* justification of induction has any chance of success and, second, that the prospects for such an *a priori* justification, contrary to widespread belief (or prejudice), are quite good.

§7.2. THE SHAPE OF THE PROBLEM

In a lecture on Bacon delivered in 1926, CD. Broad describes the failure of philosophers to solve the problem of justifying inductive reasoning as “the scandal of Philosophy.”¹ Broad’s choice of terms is noteworthy. The failure to solve a serious intellectual problem would not in itself be scandalous: perhaps there simply is no

¹ “The Philosophy of Francis Bacon,” reprinted in Broad (1952), pp. 117-43; the passage quoted is from p. 143.

solution or only one so difficult and obscure that no stigma would attach to the failure to find it. What might make the situation with regard to induction seem a scandal is such a failure together with the overwhelming intuitive conviction that there must be a solution and indeed a fairly obvious one, that thoroughgoing inductive skepticism is obviously an unreasonable position. (Broad suggests such a view by describing inductive reasoning as “the glory of Science,” as well as “the scandal of Philosophy”)

One purpose of the present chapter is to suggest that the scandal of which Broad speaks (for I agree that it is a scandal) is still very much with us, despite the best efforts of recent analytic philosophy. Indeed, I shall argue, the typical analytic approaches to the problem of induction not only do not succeed in removing the scandal, but never had any chance of such success in the first place: rather than solving the central problem, they in effect concede that it cannot be solved, and then proceed to offer one or another sort of palliative.

I begin with a schematic account of the problem of induction as I shall understand it here. Suppose that there is some reasonably definite observational or experimental situation *A*, and that out of a large number of observed instances of *A*, some fraction m/n have also possessed some further, logically independent observable property or characteristic *B*; in brief, m/n of observed *As* have also been *Bs*. Suppose further that the locations and times of observation, the identity of the observers themselves, the conditions of observation, and any further background circumstances not specified in the description of *A* have been varied to a substantial degree; and also that there is no relevant background information available concerning either the incidence

of *Bs* in the class of *As* or the connection, if any, between being *A* and being *B*.²

In the situation as described, a standard (enumerative) inductive inference would move from the premise that m/n of observed *As* are *Bs* to the conclusion that (probably), within some reasonable measure of approximation, m/n of all *As* (observed or unobserved, past, present, or future, even hypothetical as well as actual) are (or will or would be) *Bs*. In the special case in which the fraction in question reduces to 1, the conclusion would be that probably all *As* are *Bs*. ...

In its most basic form, the problem of induction is the problem of why inferences that satisfy this schema should be expected to lead or at least to be likely to lead to the truth about the world. Is there any sort of rationale that can be offered for thinking that conclusions reached in this way are likely to be true if the inductive premise is true – or even that the chance that such a conclusion is true is enhanced to any degree at all by the truth of such a premise? If we understand epistemic justification in the way discussed earlier in this book, that is, as justification that increases to some degree the likelihood that the justified belief is true and that is thus conducive to finding the truth, the issue is whether inductive reasoning confers any degree of epistemic justification, however small, on its conclusion.

² In this chapter, I will simply stipulate that the predicates involved in such arguments are not of the sort (such as “grue” and “bleen”) that are involved in Goodman’s “new riddle of induction.” See Goodman (1955). Contrary to the views of many philosophers, I cannot see that the issues involved in the “new riddle” have any major bearing on the classical problem of induction; but a consideration of them would in any case take more space than is available in the present chapter.

Hume's original elaboration of this problem continues, as we shall see, to have a major and not altogether salutary influence on contemporary views, and it will be helpful to look briefly at his discussion.³ Hume's focus is narrower in two ways: he is concerned only with cases (i) in which *all* observed cases of *A* are cases of *B*, and (ii) in which the intended conclusion is that being *A* is the *cause* of being *B*. But the nub of the problem is the same. Having argued that causal knowledge always depends on repeated experience of the putative causal sequence, Hume proceeds to ask how such repeated experience warrants or justifies the causal conclusion. What sort of *reasoning* moves from the observation of particular cases in which *A* has been followed by *B* to the general conclusion that *A* will always be followed by *B*? His initially startling thesis is that there is no such reasoning, that the conclusion in question is not based on reasoning at all but is rather the result of an ultimately arational process: *custom* or *habit*.

Besides the challenge to supply such reasoning, Hume offers an argument, specifically a dilemma, to show that no possible line of reasoning could justify the inductive conclusion. Such reasoning, he argues, would have to be either *a priori* demonstrative reasoning concerning relations of ideas or "experimental" (i.e., empirical) reasoning concerning matters of fact and existence. It cannot be the former, because all demonstrative reasoning relies on the avoidance of contradiction, and it is not a contradiction to suggest that "the course of nature may change," that sequences of events which occurred regularly in the past may not be repeated in the future. But the reasoning also cannot be based on experience since the justifiability of experimental reasoning, of generalizing

³ See David Hume, *An Inquiry Concerning Human Understanding* (Hume 1748), section IV. I will not consider here the similar but more complicated account in Hume's *Treatise of Human Nature* (Hume 1739-40).

from experience, is precisely what is at issue and cannot be assumed without begging the question. Hence, he concludes, there can be no such reasoning.

An alternative formulation of Hume's dilemma, in some ways clearer, may be obtained by formulating it with reference to a principle that he mentions but never focuses on very directly: the Principle of Induction, which says roughly that the future will resemble the past (or, better, that unobserved cases will resemble observed cases). The suggestion is that inductive arguments should be construed as enthymematic, with some such principle serving as the suppressed premise. Hume's argument is then that there is no way in which the Principle of Induction can itself be epistemically justified: it cannot be justified *a priori* because its denial is not a contradiction; and it cannot be justified by appeal to experience without reasoning in a circle, since an experiential argument will presumably be based on the fact that the principle has been (generally) true in the observed past and hence will ultimately depend on the very same principle. Thus inductive reasoning, being dependent on an unjustifiable principle, is itself unjustifiable.

Perhaps the best way to appreciate the destructiveness of this conclusion is to consider the skeptical view that is its apparent corollary. As Hume, along with many others, points out, the conclusion that inductive reasoning is unjustifiable appears to decisively undermine the rational credentials of both the scientific and the commonsense views of the world. Not only does it render epistemically unjustified all inductively supported beliefs in laws or regularities in the world, but since even the beliefs in a world of enduring objects and, via memory, in one's own past history seem to rely ultimately on such regularities, the unjustifiability of induction arguably leads to perhaps the most radical form of skepticism imaginable: a solipsism in which my epistemically

justified beliefs are restricted entirely to my own present experience. Such an extreme version of skepticism is obviously enormously implausible from an intuitive standpoint, thus providing an equally strong intuitive reason for thinking that a satisfactory justification for inductive reasoning must be available and making it seem intellectually scandalous if none can be found.

What is the contemporary response to this problem? Though there has been little explicit discussion of late, the generally received view seems to go something like this: Hume's dilemma, it is claimed, demonstrates decisively that induction cannot be epistemically justified if epistemic justification is understood in the way discussed earlier, that is, demonstrates that it is impossible to give any non-question-begging argument or reason to show that the conclusion of an inductive argument which fits the schema set out above is likely to be true or even that its chances of truth are thereby enhanced to some degree. For such an argument would have to be either deductive or inductive in character: a deductive argument could not succeed because there is no contradiction in supposing that any or all such inductive conclusions (whose truth has not been independently established) are false; while an inductive argument would beg the question. But this result, the received view continues, does not show that induction is unjustified or rationally unacceptable, so that the skeptic would prevail. Instead, it is claimed (and here the received view divides into two main versions) either:

- (a) that induction can be adequately justified in a different, "pragmatic," way, roughly by showing that it is nonetheless our best hope for finding the truth; or
- (b) that the problem of induction can be "dissolved" by showing, through linguistic or conceptual analysis, that the demand for

a non-trivial justification of inductive reasoning ultimately makes no sense.

And in either case, it seems to be suggested (though often not very explicitly), the skeptical challenge is adequately dealt with, even if not exactly refuted.⁴

This response to the problem—one which, as we will see more clearly below, flows more or less directly from a repudiation of the rationalist view of *a priori* justification—seems to me deeply unsatisfactory. My conviction is that neither of these distinctively analytic “solutions” to the problem of induction is adequate to meet the problem or to lessen at all the force of the threatened skeptical conclusion. ...

§7.3. THE PRAGMATIC JUSTIFICATION OF INDUCTION

Here BonJour considers Hans Reichenbach’s pragmatic justification of induction, which was developed further by Wesley Salmon.⁵ I have cut almost all of this section, leaving only a

⁴ I will not consider in this chapter a third contemporary approach to the problem of induction, that of Karl Popper. See, e.g., Popper, “Conjectural Knowledge: My Solution to the Problem of Induction,” reprinted in Popper (1972), pp. 1-31. Though Popper describes his view as a solution to the problem, it seems to amount mainly to the insistence that the problem as posed here cannot be solved, i.e., that inductive evidence provides no reason at all to think that the corresponding inductive conclusions are true, thus endorsing inductive skepticism rather than even attempting to answer it. More generally, Popper’s overall epistemological view is devastatingly skeptical in its implications, implications that are only lightly disguised by his use of the term ‘corroboration’ in a highly misleading way that departs strongly from its ordinary meaning.

⁵ See Reichenbach (1938), pp. 339-63; and Reichenbach (1949), pp. 469-82. References in the text are to the pages of Reichenbach (1938).

summary of Reichenbach's approach and Bonjour's main criticism of it.

Reichenbach's basic move is to treat induction, not as a form of inference, but rather as a *method* for arriving at "*posits*." A posit is not a statement or belief, not something asserted or maintained as true. Instead, it is analogous to a bet made in a gambling situation. Just as a gambler who wagers on red while playing roulette is not thereby asserting and need not believe that red will be the actual result (though he may of course also have such a belief), so also the scientist in the standard inductive situation who adopts the posit that the proportion of *A*'s that are *B*s is m/n is not thereby asserting and need not believe that this is even likely to be the true value in reality. His posit is an intellectual wager, nothing more. Construed in these terms, the inductive method says roughly that one should posit the observed proportion as the true proportion and then correct and continue to correct that initial posit as new information comes in. ...

The fundamental problem, however, is that ... the significance of Reichenbach's pragmatic justification in relation to the original problem remains obscure. As he himself indeed insists, that justification still yields no reason at all for thinking that inductive conclusions, or any of the myriad further beliefs that are epistemically dependent on them, are to any degree likely to be true. The sort of justification in question is thus not epistemic justification, as that concept was construed above; to show that beliefs are justified in this alternative way does not answer, or even purport to answer, the basic skeptical worry about induction, and is indeed quite compatible with the deepest degree of skepticism. It is thus hard to see why it should be regarded as any sort of solution to the classical problem of induction.

It is clear what the response of the proponents of the pragmatic justification to this criticism would be: they would argue, following Hume, that this is the best justification that is possible for induction, with the implication being that the best we can do must be good enough. But of course the fact, if it is a fact, that the best we can do is quite compatible with extreme skepticism tells in favor of the skeptical view, not against it. The point I want to insist on, however, is the extreme intuitive implausibility of such a result, according to which the most carefully derived results of science are epistemically no better, indeed worse, than a gambler's bets. ...

Here in especially clear-cut form is the intellectual scandal of which Broad spoke. I find it hard to believe that anyone who is at all familiar with the spectacular successes of modern science or its even more conspicuous technological by-products can believe this for even a moment, and perhaps even harder to understand how such vigorous proponents of science and scientific method as Reichenbach and Salmon can accept it with apparent equanimity.

§7.4. THE ORDINARY LANGUAGE JUSTIFICATION OF INDUCTION

This attempt to dissolve the problem of induction is summarised by Robert Martin as: "Using one's past experience as a guide to the future is, after all, *exactly what we mean* by 'rationality'." (*Epistemology*, Chapter 7, in the section "The Problem of Induction".)

BonJour focuses here on Peter Strawson's version of this approach, but also mentions A. J. Ayer and Paul Edwards.⁶ BonJour summarises Strawson's view roughly as follows:

(1) Believing in accordance with strong evidence is reasonable.

(2) Believing in accordance with inductive standards is, by definition, believing in accordance with strong evidence.

∴ Believing in accordance with inductive standards is reasonable.

Here's part of BonJour's criticism:

The central problem with Strawson's argument may perhaps be made clearer by considering an analogous case. Imagine a religiously oriented community in which judgments on a wide variety of factual issues are made by appeal to a body of sacred literature that is generally accepted as authoritative. If a skeptic were to question whether believing in accordance with evidence of this sort yields beliefs that are epistemically justified, that is, likely to be true, we could imagine a member of the community replying as follows:

Of course believing in accordance with scripture results in justified beliefs! Beliefs arrived at in this way are what we mean by "justified beliefs" in this community. It is an analytic truth that beliefs supported by strong evidence are justified; and it is also an analytic truth that being highly in accord with scripture constitutes strong evidence.

But such a reply to the skeptic is irrelevant to the skeptic's challenge if "justified" does not mean epistemically justified; and

⁶ See Ayer (1946), pp. 49-50; Edwards (1949); and Strawson (1952), chapter 9. References in the text in this section are to the pages of Strawson (1952).

either question-begging or guilty of equivocation otherwise. Here too, the basic issue is whether what the community in question accepts as strong evidence really is strong evidence in the epistemically interesting sense. And on this question, the argument just offered, like Strawson's argument concerning induction, sheds no light at all. Nor can any argument that appeals only to generally accepted standards (or to the reflection of such standards in ordinary usage) do any better. ...

§7.5. THE INDUCTIVE JUSTIFICATION OF INDUCTION

The previous two sections disposed of option (a) that induction can be justified pragmatically, and (b) induction is reasonable by definition. Bonjour then says that there are only two possible remaining ways out of this "intellectual coal pit":

- (i) Induction can be justified empirically, as Mill claimed, or
- (ii) Induction is justified by *a priori* knowledge.

BonJour concludes this section as follows:

As we have seen, the basic objection to an empirical justification of induction is that it is inevitably circular and question-begging: obviously no set of particular experiential claims can by themselves constitute such a justification, and any attempt to generalize beyond such particular claims will employ the very mode of reasoning whose acceptability is at issue.

§7.6. IS IT POSSIBLE TO JUSTIFY INDUCTION *A PRIORI*?

Here Bonjour attempts (once again) to deflect the charge that "seeking an *a priori* justification of induction is knocking futilely on

doors that have long been firmly and irrevocably closed and boarded up.”

§7.7. TOWARD AN *A PRIORI* JUSTIFICATION OF INDUCTION

Finally BonJour is ready to give his account of how inductive reasoning can work, if some *a priori* insights are available to supplement our empirical knowledge. Recall that he is attempting to show that, in cases where m/n of observed *As* have also been *Bs*, it is reasonable to infer (with high probability) that about m/n of all *As* are *Bs*, including those – the vast majority – that have not been observed.

At the outset, BonJour realises that an important qualification is needed. It is not enough for the *Bs* to have a proportion m/n in the *As* that have been observed. It is also necessary that the proportion of *Bs* among the observed *As* should appear to be *converging* to m/n , as time goes by.

... To see this, imagine a case where the relation between *A* and *B* is entirely unlawful or random. Depending on what other factors are relevant to each of them, the observed proportion of *As* that are *Bs* might vary indefinitely over time, drifting from one value to another and assuming no stable proportion. In such a case, while at any particular moment there would still be a proportion of observed *As* that have been *Bs*, there would be no reason at all to think that this proportion reflects any objective regularity that can be justifiably extended to unobserved cases or future cases or hypothetical cases. Thus what needs to be added to our earlier specification of standard inductive evidence is the further requirement that the observed proportion of *As* that are *Bs*, rather than varying irregularly over the range of possible values, converges over time to the fraction m/n and thereafter remains at least approximately constant as significant numbers of new

observations come in. Subsequent references to standard inductive evidence or to a standard inductive premise will be understood to include the stipulation that this constancy condition is satisfied. (In the case where *all* observed *As* are *Bs*, this condition is of course automatically satisfied - which may be why its importance has usually been overlooked.)

What sort of an *a priori* reason might be offered, then, for thinking that a standard inductive conclusion is likely to be true when such a standard inductive premise is true? The intuitive idea behind the reason to be suggested here is that an objective regularity of a sort that would make the conclusion of a standard inductive argument true provides the best *explanation* for the truth of the premise of such an argument. This idea is not especially novel by itself: something like it has been suggested by a number of other recent discussions of induction, though usually without making clear what the epistemological status of the underlying premises is supposed to be and in particular without construing the resulting justification as *a priori*.⁷ I will first offer a sketch of the main line of argument and then consider briefly some further problems and refinements. ...

Consider again the situation described by standard inductive evidence, under our revised account: the proportion of observed *As* that are *Bs* has converged on some relatively constant value m/n and continues to closely approximate that value as significant numbers of new observations are added. From an intuitive standpoint, the overwhelmingly obvious question to ask is: what is the *explanation* for this situation? *Why* does the observed proportion continue to approximate m/n rather than fluctuating widely as new observations are made? This is not a situation that would obtain for just any choice of *A* and *B*, and some reason

⁷ BonJour doesn't list these authors, but one of them is David Armstrong, in *What is a Law of Nature?* (1983).


seems to be needed for its occurring in the case in question. Of course, it is logically possible that the results in question represent the operation of nothing more than mere random coincidence or chance, but it seems evident, and, as far as I can see, evident on a purely *a priori* basis, that it is highly unlikely that only coincidence is at work, an unlikelihood that increases rapidly as the number of observations is made larger. My suggestion is thus that the following thesis is justified *a priori*:

(I-1) In a situation in which a standard inductive premise obtains, it is highly likely that there is some explanation (other than mere coincidence or chance) for the convergence and constancy of the observed proportion (and the more likely, the larger the number of cases in question).

Indeed, once general prejudices about *a priori* knowledge have been defused, the *a priori* status of (I-1) seems sufficiently obvious to require little discussion.

...

In other words, BonJour (like Leibniz) sees a rational inference from A to B as proceeding indirectly, via a conclusion about the cause (or explanation) of A, as shown in the diagram below. From A, we infer the theory that best explains A, and then from the theory we infer B.



For example, suppose that we repeat an experiment A (e.g. flipping a certain coin) many times and the proportion of cases in which outcome B is observed to occur seems to be converging to

the general vicinity of 71.4%. Rather than just expecting this to continue, on the basis of instinct, we instead ask *why* there is such a stable proportion of *B*s, taking this particular value, in this experiment *A*. A generic explanation for this type of phenomenon is that each time *A* is performed, there is a fixed objective chance of *B* occurring, that is something close to 0.714, and that the outcomes of successive experiments are independent. (One can show mathematically that such “iid” – independent, and identically-distributed – random variables generate such stable proportions with a very high probability.) Once we have a justified belief about the objective chance of *B* in the experiment *A*, we can use it to predict the occurrence of *B* in further cases.

BonJour closes this chapter with by addressing four more criticisms, but these are not included here. I have instead copied below a passage from the *Stanford Encyclopedia of Philosophy* that criticises BonJour’s *a priori* account of induction.

The Problem of Induction

Leah Henderson, 2022

The “Nomological-explanatory” solution, which has been put forward by Armstrong, BonJour and Foster (Armstrong 1983; BonJour 1998; Foster 2004) appeals to the principle of Inference to the Best Explanation (IBE). According to IBE, we should infer that the hypothesis which provides the best explanation of the evidence is probably true. Proponents of the Nomological-Explanatory approach take Inference to the Best Explanation to be a mode of inference which is distinct

from the type of “extrapolative” inductive inference that Hume was trying to justify. They also regard it as a type of inference which although non-deductive, is justified *a priori*. For example, Armstrong says “To infer to the best explanation is part of what it is to be rational. If that is not rational, what is?” (Armstrong 1983: 59).

The *a priori* justification is taken to proceed in two steps. First, it is argued that we should recognize that certain observed regularities require an explanation in terms of some underlying law. For example, if a coin persistently lands heads on repeated tosses, then it becomes increasingly implausible that this occurred just because of “chance”. Rather, we should infer to the better explanation that the coin has a certain bias. Saying that the coin lands heads not only for the observed cases, but also for the unobserved cases, does not provide an explanation of the observed regularity. Thus, mere Humean constant conjunction is not sufficient. What is needed for an explanation is a “non-Humean, metaphysically robust conception of objective regularity” (BonJour 1998), which is thought of as involving actual natural necessity (Armstrong 1983; Foster 2004).

Once it has been established that there must be some metaphysically robust explanation of the observed regularity, the second step is to argue that out of all possible metaphysically robust explanations, the “straight” inductive explanation is the best one, where the straight explanation extrapolates the observed frequency to the wider population. For example, given that a coin has some objective chance of landing heads, the best explanation of the fact that m/n heads have been so far observed, is that the objective chance of the coin landing heads is m/n . And this objective chance determines what happens not only in observed cases but also in unobserved cases.

The Nomological-Explanatory solution relies on taking IBE as a rational, *a priori* form of inference which is distinct

from [Humean style] inductive inferences like inference *I*. However, one might alternatively view inductive inferences as a special case of IBE (Harman 1968), or take IBE to be merely an alternative way of characterizing inductive inference (Henderson 2014). If either of these views is right, IBE does not have the necessary independence from inductive inference to provide a non-circular justification of it.

One may also object to the Nomological-Explanatory approach on the grounds that regularities do not necessarily require an explanation in terms of necessary connections or robust metaphysical laws. The viability of the approach also depends on the tenability of a non-Humean conception of laws. There have been several serious attempts to develop such an account (Armstrong 1983; Tooley 1977; Dretske 1977), but also much criticism (see J. Carroll 2016).

Another critical objection is that the Nomological-Explanatory solution simply begs the question, even if it is taken to be legitimate to make use of IBE in the justification of induction. In the first step of the argument we infer to a law or regularity which extends beyond the spatio-temporal region in which observations have been thus far made, in order to predict what will happen in the future. But why could a law that only applies to the observed spatio-temporal region not be an equally good explanation? The main reply seems to be that we can see *a priori* that laws with temporal or spatial restrictions would be less good explanations. ...