CHAPTER FOUR

The Problem of Induction

Having obtained, with Descartes's help, an initial overview of the epistemological landscape and having explored the general concept of knowledge, we now turn in this and the next several chapters to a consideration of some more specific epistemological problems or issues, all of them having to do with the justification of specific kinds of beliefs. We will begin with a problem that has the merit of being tightly focused in a way that makes the epistemological issue exceptionally clear.[1] The problem of induction has to do with what reasons or justification there are for accepting *general* conclusions on the basis of observations of *particular* instances falling under them, for example, for accepting the general conclusion that a cube of sugar will always dissolve in a large glass of water at room temperature on the basis of many observations of such cubes dissolving under those conditions and none where they fail to dissolve. Clearly we often reason in this general way, but what may not be immediately apparent is how utterly central such reasoning is to most of our supposed knowledge of the world (a point that will be further discussed below).

What exactly is the *problem* about such reasoning? We commonly regard the observation of many particular instances as providing a good reason for the corresponding general conclusion, but are we in fact justified or rational in reasoning in this way? And if so, why? What specific form do the reasons or the justification take? How, that is, could you explain to someone who somehow failed to see the point just *why* a conclusion of this sort genuinely follows from the corresponding observational premises?

Are there perhaps intermediate steps of some sort that could be filled in to make the reasoning clearer, or is there some other way to do this? Without some more specific account of this sort, the claim that reasoning of this sort is in fact good reasoning remains open to challenge—and it has in fact been very seriously challenged.

As mentioned earlier, the problem of induction is a relatively recent addition to the catalog of fundamental epistemological problems, having almost entirely escaped the notice of Descartes, all of his predecessors, and his immediate successors. It was first explicitly formulated by David Hume,[2] who advocated the skeptical thesis that observations of particular instances provide *no good reason at all* for the corresponding general conclusions, that such inductive reasoning is reasoning in name only and is in fact quite unjustified. From the standpoint of common sense, this is quite a startling conclusion, and you may be even more surprised to learn that a very substantial majority of recent philosophers have agreed that Hume is essentially right—though many of them, as we shall see, have tried to find ways to put a cosmetically better face on this intuitively unappealing result. We will have to try to decide whether or not this quite skeptical conclusion is in fact correct.

Inductive Reasoning: Two Examples and a General Characterization

We will begin by trying to get a clearer and more detailed idea of the precise sort of (supposed) reasoning whose justification is in question. I begin with two relatively simple examples, one of them already briefly mentioned, on the basis of which we can arrive at a more general characterization of the essential features of inductive reasoning.

Example 1. I put a small cube of ordinary white sugar (sucrose) into a large (approximately twelve-ounce) glass of tap water at room temperature (which I will specify broadly as the range from 60 to 80 degrees Fahrenheit), and in a fairly short time the cube dissolves completely. It occurs to me to wonder whether white sugar always behaves in this way, and so I proceed to do a series of tests. I purchase as many different brands and configurations of white sugar as I can find (beet sugar and cane sugar, sugar from different regions and countries, cubes and tablets, large bags and small packets) and put approximately the same quantity of each sample of sugar into the corresponding number of separate glasses of water, all of approximately the same volume and again at approximately room temperature. Though the time required varies somewhat, the sugar always dissolves. I do the same thing over and over, as I travel around the country and to different parts of the world. I also ask all of my friends and acquaintances in this country and abroad and perhaps even on the space shuttle to do likewise, and to report the results to me. In this series of tests, the source and other specific details concerning the sugar vary widely, as do the source of the water, the shape of the container, the time of day, the day of the week, the season of the year, and so forth. Under all of these varied conditions, the sugar still *always* dissolves, as long as both it and the water are reasonably pure. Assume also that I have no relevant background information of any kind, that my information relevant to the behavior of sugar in this respect is entirely confined to these specific tests. Eventually I conclude on the basis of my information that small quantities of sugar (approximately one teaspoon in size) always dissolve when placed in twelve ounces of water at room temperature. Am I *epistemically justified* in accepting this conclusion? That is, do I have a good reason to think that the conclusion in question is highly likely to be *true*? And if so, what exactly is that reason? Does it depend only on the set of observations or is there perhaps some sort of further premise or principle involved? (Stop and think about this question for yourself before proceeding further. Do you think that this conclusion is justified, and if so, why?)

Example 2. I own a new air gun (essentially a fancier version of a BB gun), and I become curious about how consistently it shoots, that is, about how much variation there is in where the pellet hits that is due to the gun itself and not to the steadiness of aim of the person using it or to outside conditions that affect either the gun itself or the flight of the pellet. I therefore decide to perform the following experiment. I carefully and firmly fasten the gun to a fixed support (thereby avoiding the problem of steadiness of aim) so that it is aimed in a horizontal direction at a blank target. I proceed to fire an extended series of shots, using one specific kind of pellet and being careful to avoid gusts of wind and variations in the temperature of the gun itself (such as might be produced by sunlight). The result is that the shots cluster, with most of them in a very small area and the rest distributed fairly symmetrically around that area in different directions. Measuring and counting carefully, I determine that approximately 90 percent of the shots fall within a two-inch-diameter circle centered on the area of greatest concentration. I repeat the experiment and have it repeated by others in many different locations, continuing to use the same pellets and new samples of the same brand and model of gun, having the target in a horizontal direction at the same approximate distance, keeping the temperature range fairly constant and avoiding windy conditions, but varying the other circumstances as much as possible. The results are always the same, within a close measure of approximation: the percentage of pellets within a two-inch-diameter circle centered on the area of greatest concentration is always between 88 percent and 92 percent.

Assume again that I have no relevant background information (though I have made the untested conjectures that wide temperature variation might affect the behavior of the gun and that wind might affect the flight of the pellets). I eventually conclude that under the specified conditions, approximately 90 percent of the shots from a new gun of that brand and model using those pellets will fall within such a two-inch circle. Do the experiments described give me a good reason for thinking that this conclusion is highly likely to be true? (Though these are quite simple examples, notice how carefully I have had to describe them in trying to make sure that nothing relevant has been left out—something that in fact took repeated additions and corrections when I was writing this section. Have I succeeded in this, or can you think of further things that should have been mentioned?)

How might we give a more general characterization of the structure of these and similar examples and of the reasoning involved? First, we have two observationally determinable features or conditions (which may be as complicated as we choose): first, the feature or condition that fixes the general sort of case being investigated (call this A); and, second, a further independently observable feature or condition that may or may not result from or be associated with a particular instance of condition A (call this B). Thus in example 1, condition A would be the specified quantity of sugar being placed in the specified quantity of water with the temperature falling in the indicated range; and condition B would be the subsequent dissolving of the sugar. And in example 2, condition A would be an air gun of a certain specific type being fixed in place and fired with a certain specific kind of pellet at a target a certain horizontal distance away under the further conditions specified; and condition B would be the clustering of the shots within the specified area to the degree indicated. Second, we have many observed instances of A, with the observers and other circumstances (those not specified in the description of A) being varied as widely as possible, out of which some fraction that we may formulate as m/n are also observed to be instances of B. (In the first example, this fraction is just all or 100 percent, while in the second example it is 90 percent.) A full description of all of these observations for such a case is what I will call a *standard inductive premise*. Third, on the basis of this premise and with no other relevant information, the conclusion is drawn that *approximately* m/n of *all* instances of A will also be instances of B (a *standard inductive conclusion*). This is intended to be understood as claiming not only that this will be true of past, present, and future instances, but also that it would have been true of *possible* instances that never became actual: sugar that was never put in water or even sugar that might have been produced but never was; guns that could have been so tested but weren't and even guns that might have been manufactured but weren't. *Inductive reasoning* or an *inductive inference* is just reasoning from a standard inductive premise to the corresponding standard inductive conclusion, that is, concluding on the basis of this kind of premise (and no other information) that this kind of conclusion is highly likely to be true.

The problem of induction then is just whether or not reasoning of this sort is rationally cogent (and of course why): whether and why such a premise does indeed provide a good epistemic reason or strong epistemic justification for the resulting conclusion; that is, whether and why the truth of a standard inductive premise makes it highly likely—or even, for that matter, likely to any degree at all—that the corresponding standard inductive conclusion is true.[3]

Sometimes a version of this issue is formulated in terms of an envisaged further premise that could be added to the argument so as to make the reasoning more obviously cogent. Such a premise, often labeled the Principle of Induction, is perhaps most often formulated as the claim that the future will resemble the past, but this is not really adequate to justify the full scope of the standard inductive conclusion. A somewhat better version would say that unobserved and merely possible instances are likely to resemble observed instances. But if such a premise were added, this would merely shift the issue to that of how this new premise is itself justified. Thus adding such a further premise really does nothing to advance the main issue.

We will first look at Hume's argument for a skeptical response to this problem, an argument that is interesting on its own and has also had an enormous impact on subsequent discussions of the issue.

Hume's Dilemma

Hume begins by raising a challenge to those who think that inductive inference is rationally cogent. A standard inductive premise and the corresponding standard inductive conclusion are, he points out, two quite distinct propositions. The transition from the one to the other thus requires some inferential "process of thought" that needs to be spelled out and explained:

> If you insist that the inference is made by a chain of reasoning, I desire you to produce that reasoning. The connection between these propositions is not intuitive [that is, not just self-evident]. There is required a medium which may enable the mind to draw such an inference, if indeed it may be drawn by reasoning and argument.[4]

Hume's point here is that the supposed inferential connection between a standard inductive premise and a standard inductive conclusion is certainly not so straightforward and obvious as not to require any sort of explanation.

One way to see this point more clearly is to notice that the conclusion of such an inference goes *very far* beyond the information contained in the premise, making claims about indefinitely many unobserved instances and even about merely possible ones. Why then should the relatively narrow information in the premise be regarded as a reason for thinking that this much wider and more sweeping conclusion is true? This is not something that can be just taken for granted or assumed without question. If the conclusion is to be reasonable at all, Hume is suggesting, then some further account must be possible of the inferential process of thought or the steps of reasoning whereby it is reached. Hume confesses that he is unable to arrive at any satisfactory account of this reasoning, and suggests that others will do no better. (This is an exceptionally clear example of the general form that epistemological problems typically take: we have some sort of evidence or basis E upon the strength of which some sort of further claim or conclusion C is accepted, and the question is whether and why the transition from E to C is rationally cogent.)

While the force of a challenge of the sort that Hume is raising obviously increases as philosophers over the years try and apparently fail to meet it, this failure alone obviously cannot establish conclusively that, as Hume claims, there is no such account to be given because the so-called inference is in fact not rationally cogent at all. As he himself suggests [48–49], perhaps the reasoning in question is very subtle or very difficult, and this accounts for the repeated failure to give a clear account of it. Hume's response to this suggestion is that the reasoning cannot be as difficult as that, indeed cannot be very difficult at all, since it is apparently familiar to young children and even animals, who generalize from experience in more or less the same way. This, however, is inconclusive: it is certainly possible that a cogent line of reasoning of the sort in question genuinely exists, even though animals, children, and even unsophisticated adults arrive at their conclusions in some other way—perhaps by a process of conditioned habit formation that is not really governed by reason (this being Hume's eventual suggestion for a general account of how such conclusions are arrived at—but *not* justified).

But Hume also offers a much more powerful line of argument, one that purports to show conclusively that no cogent reasoning of the sort in question is even possible. The argument in question is what logicians call a *dilemma*: that is, it argues (a) that there are only two relevant possibilities (in this case two possibilities for the sort of reasoning that might justify induction), and (b) that each of these possibilities leads to the same conclusion, which must therefore be correct (in this case the conclusion that no possible reasoning could genuinely justify an inductive inference). More specifically, Hume claims that there are two and only two general kinds of reasoning: what he calls "demonstrative reasoning," which proceeds a priori (by thought or reason alone, without reliance on experience), and what he calls "moral reasoning" (or, later and much more clearly for a modern reader, "experimental reasoning" [51]), which relies on experience [49]. His claim is that neither of these two fundamental sorts of reasoning can do the job of establishing that a standard inductive conclusion genuinely follows from the corresponding standard inductive premise.

Consider first demonstrative or a priori reasoning. Hume here advances (though without very much in the way of argument) the claim that all demonstrative or a priori reasoning (i) pertains only to "relations of ideas," that is, to relations among our concepts, and (ii) relies essentially on the avoidance of contradiction.[5] It is part (ii) of this claim that is most immediately relevant to the present discussion. Here it is important to bear in mind that what we are concerned with is the justification of the *inference* or *transition* from the premise to the conclusion of an inductive argument; the standard inductive premise itself is, of course, justified by experience, specifically the experiences involved in making the various observations, but that fact has no direct bearing on how the inference *from* that premise *to* the inductive conclusion is justified. Hume is claiming that the *only* way to be justified on a purely demonstrative, a priori basis in inferring from a premise to a conclusion, that is, without relying in any way on experience to justify this transition, is if accepting that premise and rejecting that conclusion leads to a contradiction. Thus, to take a very simple example, the reason that *it is raining* follows demonstratively from the claim that *today is Monday, and it is raining* is that it would be contradictory to deny the former claim (by saying that it is *not* raining) while accepting the latter (that today is Monday, and it is raining): this would amount to saying that it is both false and true that it is raining, which is an explicit contradiction (the simultaneous assertion and denial of the very same proposition).[6]

But this does not work for the inference we are interested in, for: "it implies no contradiction [to say] 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" [49]. Applied to our examples, the point is that it is in no way *contradictory* to say that all of the actually observed quantities of sugar have dissolved under the conditions indicated earlier, but that others have not or will not or would not; and similarly that there is no *contradiction* involved in saying that in the observed instances 90 percent of the pellets fired from air guns under the conditions indicated have hit in the indicated area, but that this is not true for other actual or possible instances where no such tally has been made. Therefore, Hume concludes, these alleged inferences cannot be justified by demonstrative or a priori reasoning.

Hume's discussion of the second alternative, namely reasoning that relies on experience, is a bit less straightforward and introduces some irrelevant complications having to do with causality. But the essential point can be seen by asking how an appeal to experience could possibly justify an inference from a particular standard inductive premise to the corresponding standard inductive conclusion. Clearly the correctness of such a conclusion is not itself a matter of direct observation, since it makes a claim about unobserved instances (and otherwise, the inductive inference would of course not be necessary). Thus the only apparent way that experience could play a justificatory role would be by (i) appealing to particular observed instances in the past where standard inductive premises were observed to be true and where the corresponding standard inductive conclusion turned out to be true also (and the absence or rarity of contrary instances), (ii) concluding on that basis to the general thesis that whenever a standard inductive premise is true, the corresponding standard inductive conclusion is highly likely to be true also, and then (iii) using this general thesis to justify the particular inductive inference in question.

There are, however, two obvious difficulties with an attempted justification of this sort (only the second of which is mentioned explicitly by Hume). The first is that the truth of even past inductive conclusions is not in fact something that can be simply observed: because such conclusions apply to indefinitely many future, possible, and unobserved past instances, the most that can be known by observation is that they have never (or probably rather only very rarely) been subsequently refuted. And this does not seem to be a strong enough result to do the justificatory work that is needed: that past arguments of this sort have usually led to true conclusions *as far as we can tell* does not show that the conclusion of the argument we are interested in will be true without this qualification. (This is a tricky point that you should think about carefully.)

The second, much more fundamental and obvious difficulty is that the inference from the observations in step (i) of the proposed justification to the general thesis in step (ii) is *itself* just another instance of inductive reasoning (think carefully about just how this is so), whose rational cogency is thus just as much in question as that of any other instance of such reasoning. Thus, Hume argues, to attempt to justify inductive inferences in general by appeal to this particular instance of such an inference "must be evidently going in a circle and taking that for granted which is the very point in question" [50]. As long as it is inductive reasoning in general whose justification is in question, the cogency of the argument from step (i) to step (ii) is just as much in doubt as that of any other case of such reasoning and so cannot help to remove that doubt.[7]

Thus, according to Hume, inductive inferences cannot be justified by either of the two possible kinds of reasoning, and so *cannot be justified at all*! In thinking about the significance of this claim, it is important to be clear about an aspect of the point that has in fact been mentioned above, but not emphasized: Hume's claim is not merely that such inferences are not *conclusive*, that we cannot be *completely certain* that the conclusions are true, a claim that would be at most mildly unsettling. It is rather that inductive inferences yield *no justification at all* for their conclusions: that is, that they fail to increase or enhance *to even the smallest degree* the likelihood that their conclusions are true. If he is right, then what we call "inductive reasoning" does not really deserve that label, for it is in fact of no value at all for supporting its supposed conclusions.

The Implications of Hume's Conclusion

Before turning to a consideration of the ways in which subsequent philosophers have responded to Hume's argument, we should pause to reflect on the consequences that would follow if he were correct—consequences that seem[8] from an epistemological standpoint to amount to almost total catastrophe. It is hard to develop this point fully at this stage in our discussion, but we can get some idea of it by considering briefly how some major kinds of belief and putative knowledge apparently depend, directly or indirectly, on inductive reasoning for their justification.

First, consider beliefs about the properties of various kinds of material objects and material substances. What justification do I have for the belief that the wooden floor I am walking on will support my weight, that various kinds of food will nourish rather than poison me, that the detergent I use to wash dishes will clean them rather than exploding in my face, and so on? The issues here are complicated by the presence of background knowledge and many levels of reasoning, but it is nonetheless impossible to see how beliefs of this kind could be justified without relying on inductive generalizations about the behavior of the objects and substances in question. (In fact, the first of our earlier examples was a simple case of this kind of reasoning.) One particularly important category of belief included under this general heading is beliefs about the persistence of various kinds of objects and substances through time—what reason do I have for thinking that an object, such as a tree or a building, that I observe at one time will continue to exist at later times (unless disturbed or destroyed in some definite way)? How do I know that such objects don't just vanish or pop into and out of existence for no reason?

Second, consider scientific beliefs about causal laws and also about various kinds of unobservable entities and processes (electrons, radioactivity, and the like) alleged by science to exist. Contrary to what Hume suggests,[9] there is almost certainly more to causality than just the regular succession of the events in question, but it is still impossible to see how we could have any justification for beliefs that one specific kind of event causes another without relying on inductive generalizations about the sequences in which events of those kinds or perhaps similar kinds occur. And though theories about unobservable entities and processes obviously cannot be directly justified by inductive inferences based on observation, the main arguments for the truth of such theories is that they provide the best *explanations* for patterns or regularities pertaining to things that are observable in various sorts of experimental situations, with the existence of these patterns or regularities being itself established by induction. Thus if inductive inference is unjustified, so also apparently are all such scientific beliefs.

Third, consider my beliefs about past events that are not based on memory of my own direct observations, for instance the belief that some particular historical event, such as the Battle of Waterloo or the adoption of the American Constitution, occurred. Any justification that I have for such beliefs must clearly rely on *evidence* falling into various general categories: written reports of various kinds, reported memories of other people, photographs, artifacts of various kinds, and so forth. But how could I be justified in thinking that any of these sorts of alleged evidence are genuinely reliable indications of the sorts of events that they are alleged to be evidence for without relying in some way on inductively established generalizations pertaining to the relation between such events and the production of the corresponding evidence, for example, between the occurrence of major political or social events and the production of written accounts of them that are at least roughly accurate? (None of these last three points is particularly obvious, and a full explanation of them would take us too far afield. For the time being you should take them as challenges to think about how beliefs of the various kinds might be justified and see whether you can think of any sort of justification that does not rely at some point on inductive conclusions.)

Later on, we will see how the justification of beliefs about the material world in general and also of beliefs about the mental states of other people also apparently rests in large part on generalizations that are in turn established by inductive reasoning. In fact, it has been argued pretty convincingly (see if you can see how the argument would go) that without inductive reasoning, I would be justified only in beliefs about my own existence and subjective experience at the present moment. For now, we need not decide in any firm way whether this extreme conclusion is really correct; but we can at least see that the skeptical consequences of accepting Hume's conclusion that there is no justification for inductive inferences would be very severe indeed.

How, then, have other philosophers responded to Hume's arguments, especially the dilemma argument? Although there have been a number of attempts through the years to show that Hume is mistaken, none of these has ever been very widely accepted. On the contrary, as noted earlier, the prevailing response, especially in recent times, has been that Hume is basically correct: that his argument succeeds in showing that inductive inferences cannot be justified *if* that means showing that such inferences establish that their conclusions are *to any degree* likely to be true on the basis of the truth of their premises. Indeed, the main recent responses to Hume have been to concede completely this central point and then try to soften or mitigate its significance by arguing that inductive reasoning can still be said to be justified or rationally acceptable in other ways that do not conflict with Hume's conclusion. We will consider next the two main versions of this sort of attempt: first, the pragmatic "vindication" of induction; and, second, the "ordinary language" justification of induction.

The Pragmatic "Vindication" of Induction[10]

The main idea of the pragmatic approach to induction is that while *inductive reasoning* cannot, as the pragmatist agrees that Hume showed, be justified by showing that it is likely to lead to true conclusions, the *inductive method* of arriving at general statements about the world can nonetheless be "vindicated" by showing that in the long run it is *guaranteed* to find the truth, *if* there is a truth of the relevant sort to be found. As thus summarily formulated, the pragmatic thesis should seem extremely puzzling (how, you should be asking, can there be such a guarantee if Hume is right?), and it will take some careful work to arrive at a clear understanding of it.

We must first try to understand what the pragmatist has in mind by the inductive *method*, as contrasted with inductive reasoning. Consider again the air gun example given earlier, but think now of the investigation as an extended process in time during which I gradually acquire more and more evidence, from my own experiments and those of others, concerning the behavior in the relevant respect of air guns of the kind in question. Suppose that I am at a relatively early stage in this process: I have done a few trials and perhaps received a relatively small amount of similar information from others. Then what the inductive method instructs me to do is to (i) *tentatively* accept the claim that the proportion of pellets falling in the indicated sort of two-inch circle so far is the true or correct proportion in general, and then (ii) revise this claim when and if the overall observed proportion changes as new trials are performed. (Of course, if the proportion remains approximately constant, then no actual revision may be necessary. And this would be the case if in particular, as in the other example, the observed proportion was and remained simply all or 100 percent.)

The pragmatist uses a special term to refer to these *tentative* claims about the proportion of A's that are B's that we are instructed by the inductive method to adopt in such a case: he calls them *posits*. A posit is a statement or claim that is not *asserted* or *believed* to be true or even probable, but is rather temporarily adopted and *treated*, for some further purpose, *as though it were true*—in the case of induction so that it can gradually be modified, hopefully in the direction of something closer to the real truth. (But whether there is any real basis for such a hope is, of course, the central issue.) Such a posit is described by the pragmatist as a kind of intellectual *wager*: it is analogous, according to him, to a bet made in a gambling situation, for example, to betting that the ball in the roulette wheel will land on red. But whereas in at least many gambling situations, such a bet is an "appraised posit," in that the gambler knows the odds that it will be correct (slightly less than 50 percent in the example just given), an inductive posit, according to the pragmatist, is a "blind posit," for which it is impossible to know the chances of success—or even that there is any chance at all![11]

To see why this is so, we need to consider more explicitly what "success" in the inductive case, according to the pragmatist, would amount to. What we are seeking is a statement of the *true* proportion of A's that are B's, as opposed to the proportion that has been observed so far. But what exactly would such truth amount to? In fact, there is a problem here that is easy to miss. We understand what a true proportion would be where the instances of A constitute a fixed, definitely delimited set. Thus if A's are just the people in a certain classroom at a certain time and B's are the females, then the true proportion of A's that are B's is just the ratio of the number of females in the room to the total number of people in the room. But this simple account does not apply straightforwardly to the issue we are mainly concerned with, which concerns not only *indefinitely* many actual instances of A (in the unobserved past and in the future), but also merely *possible* instances of A (ones that might have occurred even though they actually did not). In this sort of issue, it is much less clear what the "true" proportion even amounts to.

The pragmatist approaches this issue by drawing an analogy with mathematics (though it must be clearly borne in mind that this is *only* an analogy). Mathematicians speak of what they call *limits*: for example, the limit of the value of the mathematical expression 1/x as the value of x increases indefinitely, that is, gets larger and larger without ever stopping, is zero. Of course for any specific value of x, however large, 1/x does not equal zero, but rather is equal to some small positive value. But as the values of x get larger and larger, the value of 1/x gets closer and closer to zero, *converges* on zero, so that the *difference* between that value and zero can be made smaller than any fixed value just by making the value of x large enough; and this is precisely what it means, according to the standard mathematical definition, to say that the limit is zero. Analogously, according to the pragmatist, the true proportion of A's that are B's is what he calls "the limit of the frequency": the value, *if any*, on which the observed proportion of A's that are B's converges, in approximately the same sense,[12] as the number of observed instances increases indefinitely. Thus if there is such a limit in, for example, the air gun example, then the difference between the observed proportion of pellets hitting inside the specified circle and the limit value can be made smaller than any small fixed value—and made to remain smaller for all further total observations—simply by making the number of observed pellets sufficiently large.[13]

It should be obvious that there is no guarantee or even likelihood that the method of induction will arrive at such a limit value in any relatively short number of trials. In the short run, chance variation could always yield values that are quite different from the limit value, in principle different to *any* specified degree. For the pragmatist, this point is really just an application of Hume's original argument: there is no contradiction between (1) the claim that the limit of the frequency has one value and (2) the claim that the observed proportion in any finite number of instances has some different value; and any appeal to experience would again be circular.

Moreover, much more seriously, there is no guarantee or even likelihood, according to the pragmatist, that the inductive method will find such a limit even if pursued in the long run, *even the infinitely long run*, for the very simple reason that there is no guarantee or even likelihood in general that such a limit even exists. Again apply Hume's argument, which the pragmatist endorses: there is no contradiction in saying that such-and-such proportions are observed at various times, but that the series of observed proportions never converges on a definite limit value; and again there is no noncircular way to argue for the existence of such a limit by appeal to experience.

This is perhaps a surprising result, but it is in fact fairly easy to think of examples where it is not at all implausible to doubt the existence of such a limit. Consider the proportion of people who are left-handed: perhaps left-handedness results from cultural and/or environmental factors that vary enough over time to prevent the proportion from ever converging on a limit. Or for an even clearer example, consider the proportion of people who wear pink shirts on Tuesdays: here it is very plausible that this proportion varies according to fads and fashions, the discovery and availability of dyes, religious or cultural values, changes in the work week, and the like, so as to vary widely over time and never converge on a limit.

The pragmatist's point is then that any case we happen to be interested in *might* turn out to be like these, with no limit value to be found. Thus if inductive success means finding such a limit value, there is no guarantee or even likelihood that induction will succeed in any given case. This is in fact the fundamental reason, according to the pragmatist, why inductive reasoning cannot be justified in the sense of showing that it is likely to lead to the truth.

What we *can* be sure of, according to the pragmatist, is that following the inductive method will succeed in finding the truth *if* such success is possible, that is, *if* there is a truth of the kind in question, a "limit of the frequency," to be found—or rather, strictly speaking, that the inductive method will yield a value for the proportion of A's that are B's that *approximates* the true or limiting proportion to any degree of closeness that we specify. And this is so simply because of the way that the limit in question was defined. *If* there is such a limit, then a large enough number of observations *must*, by definition, bring the observed value within any specified distance of the limit value. It would be a *contradiction* to deny that this is so, to say that there is a limit but that the observed proportion never approaches it, no matter how large the number of observed instances.

Thus this guarantee of success *if* success is possible is a demonstrative or a priori result that even Hume would accept. Moreover, the pragmatist claims, nothing better can be established for *any* other method of arriving at general conclusions on the basis of observation. Hence while the pragmatist's argument does not, for the reasons already discussed, constitute a *justification* of inductive reasoning, it does, according to him, constitute a *vindication* of the inductive method by showing that it is rational or reasonable to adopt it: what could be more reasonable than to adopt a method that is guaranteed to succeed if success is possible when there is nothing better than this to be had?

Is the pragmatic vindication of induction an adequate response to the problem of induction? In particular, is it plausible that this is really the best we can do? Though the account of why adopting the inductive method is reasonable sounds initially pretty good, there are in fact two large problems with it. These do not show that the pragmatist's claims are mistaken in themselves, but they do suggest strongly that their significance is much more limited, and the resulting skepticism much more dire, than it might at first appear.[14]

First, given the claim that induction will succeed in the long run *if* success is possible, we need to ask *how long* the long run in question must be. The answer, already implicit in our earlier discussion, is that no run of *any* finite length is guaranteed *or even likely* to be long enough. The pragmatist argument guarantees success eventually (if such success is possible), but not success by *any* particular point in the sequence of observations. (If you don't see *clearly* why this is so, you should reread the three paragraphs before the previous one.) This means that for any actual application of induction we are interested in, such as the two earlier examples, there is no number of trials, however large, for which we can have *any* degree of justified confidence that the observed proportion is a reasonable approximation of the limit even if such a limit does exist. At any given point, we *might* in fact have succeeded in approximating the limit, but we can *never* have any reason at all for thinking that this is so—or, accordingly, for being confident that we can safely act on our results in ways that depend for success on their being true. This seems to mean that induction is practically more or less useless if the pragmatic vindication is the best we can do. Indeed, while induction is guaranteed to succeed in the long run if success is possible, its likelihood of success in *any* short run is on the pragmatic view no better than that of a random guess—and guessing is of course a much less labor-intensive "method" than careful experimentation.

Second, if this is really the best we can do in justifying induction, the result is of course skepticism—and, as we saw briefly above, very probably a quite deep and severe version of skepticism that would leave little of our supposed knowledge standing. Here is a rather picturesque description of our resulting epistemic situation, given in fact by the leading advocate of the pragmatic approach, the German-American philosopher Hans Reichenbach:

> 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 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.[15]

And even this seems too optimistic: probably he should have said that it is only an apparent path that may in fact lead directly into the abyss. And moreover, we will also never be able to be justifiably confident to any degree that we have in fact emerged from the wilderness.

It is surely overwhelmingly implausible, as we look around at our orderly world and at the various scientific and technological marvels that it contains, that our epistemic situation is as dismal as this. This does not show conclusively that the pragmatist is wrong, since for all that we have seen so far, it is at least possible that there is no better justification for induction to be found. But it surely gives us a very strong motivation to seek a better solution and to anticipate with some confidence that one will turn out to be available.

The "Ordinary Language" Justification of Induction[16]

A second, quite different attempt to defend the rationality of induction while still conceding the correctness of Hume's basic argument has been advanced by adherents of the approach to philosophy known as "ordinary language philosophy." The basic claim of this once popular philosophical approach is that the traditional problems of philosophy, including the problem of induction and the other main problems of epistemology, are "pseudo-problems" that arise from misuse of language or inadequate attention to ordinary linguistic usage. Such supposed problems, it is claimed, need to be "dissolved" rather than solved: they evaporate under careful scrutiny.

In fact, as we will see, the appeal to linguistic usage is rather inessential, particularly in the case of induction, and the specific view in question could just as well be described as the common-sense justification of induction. The main claim is that inductive reasoning is reasonable or justified simply because reasoning in this way is what we commonsensically call "reasonable" in the kinds of cases in question. Consider again the examples described earlier in this chapter. Clearly, from a common-sense standpoint, a person who, on the basis of the evidence indicated in the first example, accepts the conclusion that small quantities of sugar always dissolve in the way indicated would be described as having drawn the reasonable conclusion; and someone who concludes instead that such quantities will sometimes fail to thus dissolve would be said to be unreasonable. Similarly, in the second example, drawing the indicated conclusion would be described as reasonable, and drawing any significantly different conclusion as unreasonable. Thus, the ordinary language philosopher claims, there is no meaningful issue to be raised about the reasonableness or justification of reasoning inductively—or at least none that cannot be easily and trivially dealt with.

According to the ordinary language philosopher, the very idea that there exists a significant "problem of induction" is therefore a mistake, a kind of intellectual illusion. One account[17] of how this illusion arises is the following. The basic mistake is to demand implicitly that *inductive* reasoning meet the standards of *deductive* reasoning if it is to be reasonable or justified. In a deductive argument, such as the ones that occur in areas like logic and mathematics, the conclusions follow *conclusively* from the premises, so that it is impossible to consistently accept the latter and deny the former. It is allegedly because they notice that this is not so for inductive arguments—that (as Hume pointed out) it is possible to consistently accept the premise of such an argument and deny the conclusion—that philosophers are led to think that there is a problem about whether and why induction is reasonable: one that might perhaps be solved by adding something like the Principle of Induction mentioned earlier, if only that principle could itself be somehow justified. But this whole approach, according to the ordinary language philosopher, is just a confusion. Deduction is one kind of reasoning, and induction is simply a distinct, fundamentally different kind of reasoning. Each of the two possesses its own autonomous standards of correctness or reasonableness, and there is no reason at all to expect one kind of reasoning to meet the standards of the other or for demanding that it do so. And if this mistake is not made, then it is obvious at once that induction is reasonable or justified by *inductive* standards, those reflected in ordinary usage and common sense, which are the only standards that are genuinely relevant in this sort of case. Thus the supposed problem allegedly disappears.

But there is in fact very much less force to this supposed dissolution of the problem than there may at first seem to be. For the main concern underlying the problem of induction is *not* whether inductive reasoning is "reasonable" or "justified" when judged by the standards that are implicit in ordinary usage and common sense, something about which there is no serious doubt—and which Hume does not question. The issue is instead whether those standards are *themselves* correct or reasonable or justified in a deeper sense: whether reasoning in accordance with those standards is in fact likely (as common sense of course would say) to lead to conclusions that are *true*. And this is not a question that is in any way answered by pointing out that the standards in question are the ones that we commonsensically accept. Nor, in fact, does the proponent of the ordinary language solution in fact claim that it is. On the contrary, proponents of this approach commonly concede, as indeed they must, that the fact that inductive reasoning is "justified" or "reasonable" in the way that they have explained does not in any way establish that conclusions reached in this way are likely to be true.[18] Thus the real problem of induction has been neither shown to be senseless nor in any real way dissolved.

An analogy may help to bring out the point more clearly. Suppose that there is a religious community that accepts the practice of settling certain sorts of issues, including many issues that we would regard as factual or scientific in character, by appeal to a body of sacred texts. Imagine that a skeptic about this practice emerges in the community in question, someone who asks whether there is any good reason or justification for thinking that the answers yielded by the texts are in fact likely to be true. And imagine an ordinary language philosopher who attempts to meet this challenge by pointing out that accepting the answers that are indicated by the texts is just what being reasonable means in the kinds of cases in question (according to the common-sense standards of the community in question). Clearly this does not genuinely answer the skeptic's challenge, which is really a challenge to the very practice of appealing to the texts and so cannot be satisfactorily answered by simply invoking that practice. And the situation is no different with the analogous case of induction. As one critic has nicely put the point, the ordinary language defense of induction seems to amount to no more than this: "If you use inductive procedures you can call yourself 'reasonable' [by common-sense standards]—*and isn't that nice*!"[19]

Can Inductive Reasoning Be Justified A Priori?

Thus the two most prominent recent attempts to show that it is possible to accept Hume's conclusion while still defending inductive reasoning as in some way reasonable or justified seem to come to very little. In particular, the skeptical implications of Hume's argument remain as deep and troubling as ever. Since there is no other attempt in this direction that seems to do any better, it seems pretty clear that the only way to avoid these deeply skeptical results is to find some more direct answer to Hume's dilemma that allows us to avoid his conclusion.

There appears to be no hope of refuting Hume's argument that induction cannot be justified by appeal to experience. Though a few recent philosophers have made attempts in this direction, the circular or question-begging character of such a justification seems too clear to be denied. Thus any defense of induction will apparently have to be independent of experience—that is, a priori. It also seems undeniable that Hume is again right that there is no *contradiction* involved in accepting a standard inductive premise and rejecting the corresponding standard inductive conclusion, so that an a priori argument defending induction cannot be of the simple, straightforward type that is based on avoiding contradiction. What is much less clear, as will be suggested here for this specific issue and defended in a more general way and at much greater length in the next chapter, is that Hume is correct that all a priori reasoning must be based in this way on the avoidance of contradiction—where, to repeat, a contradiction is being understood here as the simultaneous assertion and denial of the very same proposition. (The possibility to be discussed could in fact be regarded either as a third alternative to Hume's two, a different kind of a priori reasoning, or as a challenge to Hume's construal of the demonstrative, a priori alternative, but it makes no ultimate difference in which of these ways it is put.)

Reflect again on the two examples of inductive reasoning offered above and on other examples of the same kind. It certainly *seems* intellectually compelling to reason in this way in such cases, and there seems to be no particular plausibility to holding that this seeming reasonableness is somehow based on experience or observation (beyond that required to establish the standard inductive premises), nor that it is (as in the religious community case) merely a reflection of communal standards that we just happen, for no good reason, to accept. On the contrary, that the likelihood that the conclusions in question are true is substantially increased or enhanced by the corresponding observational premises *seems* very obvious, indeed just as intellectually obvious as the conclusion in many cases of logical or mathematical reasoning (even though the *degree* of support is less than conclusive). All this could still, for all we have seen so far, be an illusion of some sort, but if so, it is an extremely powerful and persistent illusion, and it is time to see whether we can find some better way of making sense of it.

What sort of an a priori reason might there be, then, for thinking that a standard inductive conclusion is likely to be true if the corresponding standard inductive premise is true? Here there is an important lesson to be learned from our earlier discussion of the pragmatic approach. The pragmatist claimed that there is in fact no a priori guarantee of any sort that in a series of observations of the kind that is summarized in a standard inductive premise, the proportion of A's that are B's will in fact converge on a definite value rather than varying irregularly among very different values, and I think he is right about this. Consider, then, a series of nonconverging observations, one in which the observed value over time does not approach closer and closer to any particular value, but simply fluctuates through the range of possible values in a way that exhibits no discernible pattern. (Perhaps the series of observed values of the proportion of people who wear pink shirts on Tuesdays might behave like this.) At any particular point in such a series, there will of course be some definite value so far of the observed proportion of A's that are B's, one that merely summarizes the observations to that point, and this fact could of course be formulated in a standard inductive premise. But does such a premise constitute any reason at all in this kind of case for thinking that the corresponding standard inductive conclusion is true? My suggestion is that the answer to this question is plainly "no." Without any appearance of convergence, such a conclusion may reflect only one temporary stage in an irregular series of values, and there is no reason at all to ascribe to it any more significance than that.

Consider in contrast the sort of case in which there is apparent convergence, that is, in which the observed values seem on the whole to be approaching closer and closer to one particular value, albeit perhaps with small fluctuations along the way—and let us modify the idea of a standard inductive premise for the rest of this discussion so as to include the stipulation that such apparent convergence has taken place. Now we do seem intuitively to have a good reason to accept the corresponding standard inductive conclusion, which in effect states that the convergence value is (approximately) the true value of the proportion of A's that are B's. But why? Why do observations that apparently converge in this way provide a kind of justification for a corresponding conclusion that nonconverging observations cannot?

My suggestion is that we now have a fact, the fact of apparent convergence, that seems to demand some sort of *explanation*. To be sure, it is always logically possible that such apparent convergence results merely from chance, but this seems more and more unlikely the longer it persists. (Think very carefully about this point: If only chance is at work, then convergence of the sort in question represents a striking coincidence, one that is unlikely to occur just because there are so many other possibilities that are equally likely, so many other patterns in which A's that are and are not B's could occur—all of which would destroy the apparent convergence.)

How then might such an apparently convergent series of observations be explained? Here is a two-part explanation that seems obvious and straightforward: (i) there is an *objective regularity* in the world, due in some way to the natures of A and B and the way in which they relate to each other, as a result of which just that (approximate) proportion of A's tend to be B's, and (ii) a series of observations of the sort in question will naturally tend to reflect that regularity, once enough instances have been observed to cancel out the effects of chance variation with regard to just which A's happen to be observed. Thus in the case of the air gun example, the idea would be that there is something about the construction and materials and operation of the guns in question that is regularly correlated, and in this case no doubt causes, the pattern of pellet distribution that is observed. And *if* this sort of explanation is the right explanation in such a case, then the proportion reflected in the convergent observations is (approximately) the true proportion, and the standard inductive conclusion is true. (Notice that we have here an account of what it is for a given proportion to be the true proportion in such a case that is significantly different from the pragmatist's and rather more natural: it is for it to reflect such an objective regularity in nature.)

Is there any reasonably plausible competing explanation for such a convergent series of observations that might upset this conclusion? Once chance has been ruled out as extremely unlikely, the only other possibility seems to be that (i) there is indeed an objective regularity involving A's and B's, resulting in a proportion of A's that are B's that is objective and regular, but that (ii) the A's that are actually observed represent a sample that is in some way or other *skewed* or *unrepresentative* in relation to the total set of A's, in such a way as to produce an observed proportion to which the observations converge, but one that does not accurately reflect the true proportion overall. If this were so, the conclusion of the explanatory argument would be false.

But *why* would the skewing reflected in (ii) occur, that is, *why* would the sample of A's be unrepresentative in spite of the variation of conditions, observers, and so forth, that is part of what is claimed by a standard inductive premise? A skewing due merely to chance would be extremely unlikely to produce regular enough results to account for the observed convergence. Thus the skewing in question would itself have to be *systematic*, that is, would itself have to result from some regular process or mechanism, which in this case could apparently only be due to the fact of observation itself: it would apparently have to be the case that the act or process of observation itself affects or somehow selects A's whose tendency to be B's differs from that of the overall population of A's.

It is in fact very hard to be sure in a particular case that a possibility of this sort does not obtain. In the case of the air gun example, perhaps merely the proximity of an observer somehow affects the gun so as to alter the results in a systematic way. This might be due to heat from handling the gun to the degree necessary to fire it repeatedly or to quantum mechanical effects of some sort[20] or to some still further, perhaps unknown mechanism.

But, somewhat surprisingly, this possibility, when carefully considered, turns out to have *no bearing at all* on the justification of inductive reasoning. It is entirely obvious that in at least some cases observational results may be influenced by the fact that observation has taken place; and it is equally obvious that observational results involving such influence have no genuine value as evidence of what would take place if observation were not occurring, so that generalizing from them would clearly be a mistake. But the claim that inductive reasoning is justified should not be construed as denying these obvious facts, and the problem of induction is not concerned with this sort of possibility. Instead, the standard problem of induction should be understood as the problem of whether and why observational results of the sort summarized in a standard inductive premise (including the claim of convergence) provide good evidence for a standard inductive conclusion—that is, as we now see, for the existence of an objective regularity—*on the assumption that this sort of observational influence does not occur*, for this is the only genuine issue.

Thus it turns out that the only apparent competitor to the explanation which makes the standard inductive conclusion true turns out not to be a genuine competitor at all, but rather reflects a possible circumstance that would make inductive reasoning not correctly applicable to the case. And, therefore, in the cases where the assumption just indicated holds, where observation is *mere* observation and does not itself affect the results, we have good reason to think that the standard inductive conclusion, representing as it does the only nonchance explanation available of the fact of apparent convergence, is true.

There is one further potential objection to be considered. As quoted in our earlier discussion, one of the things that Hume says is that inductive conclusions cannot be shown to be likely to be true because "the course of nature may change." Why doesn't this possibility still defeat our attempted justification? Even if the convergent observations were due to an objective regularity of the sort indicated, couldn't that regularity simply change in the next instant, so that even if the standard inductive conclusion still correctly describes at least the observed part of the past, it no longer correctly describes future or possible instances (and perhaps also not unobserved past ones)? This objection raises metaphysical issues that we cannot go very far into here. But the simple answer, which I believe to be correct, is that the regularity in question is not supposed to be just an ungrounded, coincidental pattern, but rather something that results in some way from the very natures of A and B themselves. Thus as long as those natures persist, that is, as long as there are A's and B's at all, the regularity in question is also at least very likely to persist, which is enough to safeguard our conclusion.

The foregoing defense of induction at least appears to be purely a priori in character. At no time did any sort of observational or experiential evidence (beyond the standard inductive premise itself) need to be brought in and appealed to in order to show either (i) that the truth of a standard inductive premise (understood as including the appearance of convergence) requires some explanation or (ii) that the existence of an objective regularity that would make the corresponding standard inductive conclusion true is the best explanation for such a fact. Instead, both of these points were defended on what appear to be entirely a priori grounds. Nonetheless, there are many philosophers who doubt strongly whether an a priori argument of this sort can genuinely be cogent. Their reservations have mainly to do, not with the specific issues surrounding induction, but rather with general views of the possibility and nature of a priori reasoning generally, a topic we will turn to in the next chapter.