Wesley Salmon

The Problem of Induction

I. The Problem of Induction

We all believe that we have knowledge of facts extending far beyond those we directly perceive. The scope of our senses is severely limited in space and time; our immediate perceptual knowledge does not reach to events that happened before we were born to events that are happening now in certain other places or to any future events. We believe, nevertheless, that we have some kind of indirect knowledge of such facts. We know that a glacier once covered a large part of North America, that the sun continues to exist at night, and that the tides will rise and fall tomorrow. Science and common sense have at least this one thing in common: Each embodies knowledge of matters of fact that are not open to our direct inspection. Indeed, science purports to establish general laws or theories that apply to all parts of space and time without restriction. A "science" that consisted of no more than a mere summary of the results of direct observation would not deserve the name.

Hume's profound critique of induction begins with a simple and apparently innocent question: How do we acquire knowledge of the unobserved?[1] This question, as posed, may seem to call for an empirical answer. We observe that human beings utilize what may be roughly characterized as inductive or scientific methods of extending knowledge from the observed to the unobserved. The sciences, in fact, embody the most powerful and highly developed methods known, and we may make an empirical investigation of scientific methods much as we might for any other sort of human behavior. We may consider the historical development of science. We may study the psychological, sociological, and political factors relevant to the pursuit of science. We may try to give an exact characterization of the behavior of scientists. In doing all these things, however, important and interesting as they are, we will have ignored the *philosophical* aspect of the problem Hume raised. Putting the matter very simply, these empirical investigations may enable us to describe the ways in which people arrive at *beliefs* about unobserved facts, but they leave open the question of whether beliefs arrived at in this way actually constitute *knowledge.* It is one thing to describe how people go about seeking to extend their knowledge; it is quite another to claim that the methods employed actually do yield knowledge.

One of the basic differences between knowledge and belief is that knowledge must be founded upon evidence-i.e., it must be belief founded upon some rational justification. To say that certain methods yield knowledge of the unobserved is to make a cognitive claim for them. Hume called into question the justification of such cognitive claims. The answer cannot be found entirely within an empirical study of human behavior, for a *logical* problem has been raised. It is the problem of understanding the logical relationship between evidence and conclusion in logically correct inferences. It is the problem of determining whether the inferences by which we attempt to make the transition from knowledge of the observed to knowledge of the unobserved are logically correct. The fact that people do or do not use a certain type of inference is irrelevant to its justifiability. Whether people have confidence in the correctness of a certain type of inference has nothing to do with whether such confidence is justified. If we should adopt a logically incorrect method for inferring one fact from others, these facts would not actually constitute evidence for the conclusion we have drawn. The problem of induction is the problem of explicating the very concept of *inductive evidence.*

There is another possibly misleading feature of the question as I have formulated it. When we ask how we can *acquire* knowledge of the unobserved, it sounds very much as if we are asking for a method for the *discovery* of new knowledge. This is, of course, a vital problem, but it is not the fundamental problem Hume raised. Whether there is or can be any sort of inductive logic of discovery is a controversial question I shall discuss in detail in a later section.[2] Leaving this question aside for now, there remains the problem of *justification* of conclusions concerning unobserved matters of fact. Given some conclusion, however arrived at, regarding unobserved facts, and given some alleged evidence to support that conclusion, the question remains whether that conclusion is, indeed, supported by the evidence offered in support of it.

Consider a simple and highly artificial situation. Suppose a number of balls have been drawn from an urn, and that all of the black ones that have been drawn are licorice-flavored. I am not now concerned with such psychological questions as what makes the observer note the color of these balls, what leads him to taste the black ones, what makes him take note of the fact that licorice flavor is associated with black color in his sample, or what makes him suppose that the black balls not yet drawn will also be licorice- flavored. The problem-Hume's basic *philosophical* problem is this: Given that all of the observed black balls have been licorice-flavored, and given that somehow the conclusion has been entertained that the unobserved black balls in the urn are also licorice-flavored, do the observed facts constitute sound *evidence* for that conclusion? Would we be *justified* in accepting that conclusion on the basis of the facts alleged to be evidence for it?

As a first answer to this question we may point out that the inference does conform to an accepted inductive principle, a principle saying roughly that observed instances conforming to a generalization constitute evidence for it. It is, however, a very small step to the next question: What grounds have we for accepting this or any other inductive principle? Is there any reason or justification for placing confidence in the conclusions of inferences of this type? Given that the premises of this inference are true, and given that the inference conforms to a certain rule, can we provide any rational justification for accepting its conclusion rather than, for instance, the conclusion that black balls yet to be drawn will taste like quinine?

It is well known that Hume's answer to this problem was essentially skeptical. It was his great merit to have shown that a justification of induction, if possible at all, is by no means easy to provide. In order to appreciate the force of his argument it is first necessary to clarify some terminological points. This is particularly important because the word *induction* has been used in a wide variety of ways.

For purposes of systematic discussion one distinction is fundamental, namely, the distinction between demonstrative and nondemonstrative inference. A *demonstrative* inference is one whose premises necessitate its conclusion; the conclusion cannot be false if the premises are true. All valid deductions are demonstrative inferences. A *nondemonstrative* inference is simply one that fails to be demonstrative. Its conclusion is not necessitated by its premises; the conclusion could be false even if the premises are true. A demonstrative inference is *necessarily truth-preserving,* a nondemonstrative inference is not.

The category of nondemonstrative inferences, as I have characterized it, contains, among other things perhaps, all kinds of fallacious inferences. If, however, there is any kind of inference whose premises, although not necessitating the conclusion, do lend it weight, support it, or make it probable, then such inferences possess a certain kind of logical rectitude. It is not deductive validity, but it is important anyway. Inferences possessing it are *correct inductive inferences.*

Since demonstrative inferences have been characterized in terms of their basic property of necessary truth preservation, it is natural to ask how they achieve this very desirable trait. For a large group of demonstrative inferences, including those discussed under "valid deduction" in most logic texts, the answer is rather easy. Inferences of this type purchase necessary truth preservation by sacrificing any extension of content. The conclusion of such an inference says no more than do the premises-often less.[3] The conclusion cannot be false if the premises are true *because* the conclusion says nothing that was not already stated in the premises. The conclusion is a mere reformulation of all or part of the content of the premises. In some cases the reformulation is unanticipated and therefore psychologically surprising, but the conclusion cannot augment the content of the premises. Such inferences are *nonampliative;* an ampliative inference, then, has a conclusion with content not present either explicitly or implicitly in the premises.

While it is easy to understand why nonampliative inferences are necessarily truth-preserving, the further question arises whether there are any necessarily truth-preserving inferences that are also ampliative. Is there any type of inference whose conclusion must, of necessity, be true if the premises are true, but whose conclusion says something not stated by the premises? Hume believed that the answer is negative and so do I, but it is not easy to produce an adequate defense of this answer. Let us see, however, what an affirmative answer would amount to.

Suppose there were an ampliative inference that is also necessarily truth-preserving. Consider the implication from its premises, $\boldsymbol{P}_1 \ldots \mathrm{P}_k$, to its conclusion $\boldsymbol{C}$. If the inference were an ordinary nonampliative deduction, this implication would be analytic and empty; but since the argument is supposed to be ampliative, the implication must be synthetic. At the same time, because the argument is supposed to be necessarily truth-preserving, this implication must be not only true but necessarily true. Thus, to maintain that there are inferences that are both ampliative and necessarily truth-preserving is tantamount to asserting that there are synthetic *a priori* truths.[4] This may be seen in another way. Any ampliative inference can be made into a nonampliative one by adding a premise. In particular, if we add to the foregoing ampliative inference the synthetic *a priori* premise, "If $\boldsymbol{P}_1$ and $\boldsymbol{P}_2$ and . . . and $\boldsymbol{P}_k$. then $\boldsymbol{C}$," the resulting inference will be an ordinary valid nonampliative deduction. Consider our example once more; this time let us set it out more formally:

1. Some black balls from this urn have been observed.
All observed black balls from this urn are licorice-flavored.

All black balls in this urn are licorice-flavored.

This argument is clearly ampliative, for the premise makes a statement about observed balls only, while the conclusion makes a statement about the unobserved as well as the observed balls. It appears to be nondemonstrative as well, for it seems perfectly possible for the conclusion to be false even if the premises are true. We see no reason why someone might not have dropped a black marble in the urn which, when it is drawn, will be found to be tasteless. We could, however, rule out this sort of possibility by adding another premise:

2. Some black balls from this urn have been observed.
All observed black balls in this urn are licorice-flavored.
Any two balls in this urn that have the same color also have the same flavor.

All black balls in this urn are licorice-flavored.

The additional premise has transformed the former nondemonstrative inference into a demonstrative inference, but we must also admit that we have transformed it into a nonampliative inference. If, however, the third premise of 2 were a synthetic *a priori* truth, the original inference, although ampliative, would have been necessarily truth-preserving and, hence, demonstrative. If the premise that transformed inference 1 into inference 2 were necessarily true, then it would be impossible for the conclusion of inference 1 to be false if the premises were true, for that would contradict the third premise of inference 2.

Hardly anyone would be tempted to say that the statement, "Any two balls in this urn that have the same color also have the same flavor," expresses a synthetic *a priori* truth. Other propositions have, however, been taken to be synthetic *a priori*. Hume and many of his successors noticed that typical inductive inferences, such as our example concerning licorice-flavored black balls, would seem perfectly sound if we could have recourse to some sort of principle of uniformity of nature. If we could only prove that the course of nature is uniform, that the future will be like the past, or that uniformities that have existed thus far will continue to hold in the future, then we would seem to be justified in generalizing from past cases to future cases--from the observed to the unobserved. Indeed, Hume suggests that we presuppose in our inductive reasoning a principle from which the third premise of 2 would follow as a special case: "We always presume, when we see like sensible qualities, that they have like secret powers, and expect that effects, similar to those which we have experienced, will follow from them."[5] Again, "From causes which appear *similar* we expect similar effects. This is the sum of all our experimental conclusions."[6]

Hume's searching examination of the principle of uniformity of nature revealed no ground on which it could be taken as a synthetic *a priori* principle. For all we can know *a priori*, Hume argued, the course of nature might change, the future might be radically unlike the past, and regularities that have obtained in respect to observed events might prove completely inapplicable to unobserved cases. We have found by experience, of course, that nature has exhibited a high degree of uniformity and regularity so far, and we infer inductively that this will continue, but to use an inductively inferred generalization as a justification for induction, as Hume emphasized, would be flagrantly circular. He concluded, in fact, that there are no synthetic *a priori* principles in virtue of which we could have demonstrative inferences that are ampliative. Hume recognized two kinds of reasoning: reasoning concerning relations of ideas and reasoning concerning matters of fact and existence. The former is demonstrative but nonampliative while the latter is ampliative but not necessarily truth-preserving.

If we agree that there are no synthetic *a priori* truths, then we must identify necessarily truth-preserving inference with nonampliative inference. All ampliative inference is nondemonstrative. This leads to an exhaustive trichotomy of inferences: valid deductive inference, correct inductive inference, and assorted fallacies. The first question is, however, whether the second category is empty or whether there are such things as correct inductive inferences. This is Hume's problem of induction. Can we show that any particular type of ampliative inference can be justified in any way? If so, it will qualify as correct induction.

Consider, then, any ampliative inference whatever. The example of the licorice-flavored black balls illustrates the point. We cannot show *deductively* that this inference will have a true conclusion given true premises. If we could, we would have proved that the conclusion must be true if the premises are. That would make it necessarily truth-preserving, hence, demonstrative. This, in turn, would mean that it was nonampliative, contrary to our hypothesis. Thus, if an ampliative inference could be justified deductively it would not be ampliative. It follows that ampliative inference cannot be justified deductively.

At the same time, we cannot justify any sort of ampliative inference *inductively*. To do so would require the use of some sort of nondemonstrative inference. But the question at issue is the justification of nondemonstrative inference, so the procedure would be question begging. Before we can properly employ a nondemonstrative inference in a justifying argument, we must already have justified that nondemonstrative inference.

Hume's position can be summarized succinctly: We cannot justify any kind of ampliative inference. If it could be justified deductively it would not be ampliative. It cannot be justified nondemonstratively because that would be viciously circular. It seems, then, that there is no way in which we can extend our knowledge to the unobserved. We have, to be sure, many beliefs about the unobserved, and in some of them we place great confidence. Nevertheless, they are without rational justification of any kind!

This is a harsh conclusion, yet it seems to be supported by impeccable arguments. It might be called "Hume's paradox," for the conclusion, although ingeniously argued, is utterly repugnant to common sense and our deepest convictions. We *know* ("in our hearts") that we have knowledge of unobserved fact. The challenge is to show how this is possible.

II. Attempted Solutions

It hardly needs remarking that philosophers have attempted to meet Hume's intriguing challenge in a wide variety of ways. There have been direct attacks upon some of Hume's arguments. Attempts to provide inductive arguments to support induction and attempts to supply a synthetic *a priori* principle of uniformity of nature belong in this category. Some authors have claimed that the whole problem arises out of linguistic confusion, and that careful analysis shows it to be a pseudo-problem. Some have even denied that inductive inference is needed, either in science or in everyday affairs. In this section I shall survey what seem to me to be the most important efforts to deal with the problem.

1. Inductive Justification. If Hume's arguments had never been propounded and we were asked why we accept the methods of science, the most natural answer would be, I think, that these methods have proved themselves by their results. We can point to astonishing technological advances, to vastly increased comprehension, and to impressive predictions. Science has provided us with foresight, control, and understanding. No other method can claim a comparable record of successful accomplishment. If methods are to be judged by their fruits, there is no doubt that the scientific method will come out on top.

Unfortunately, Hume examined this argument and showed that it is viciously circular. It is an example of an attempt to justify inductive methods inductively. From the premise that science has had considerable predictive success in the past, we conclude that it will continue to have substantial predictive success in the future. Observed cases of the application of scientific method have yielded successful prediction; therefore, as yet unobserved cases of the application of scientific method will yield successful predictions. This argument has the same structure as our black-balls-in-the-urn example; it is precisely the sort of ampliative inference from the observed to the unobserved whose justifiability is in question.

Consider the parallel case for a radically different sort of method. A crystal gazer claims that his method is the appropriate method for making predictions. When we question his claim he says, "Wait a moment; I will find out whether the method of crystal gazing is the best method for making predictions." He looks into his crystal ball and announces that future cases of crystal gazing will yield predictive success. If we should protest that his method has not been especially successful in the past, he might well make certain remarks about parity of reasoning. "Since you have used your method to justify your method, why shouldn't I use my method to justify my method? If you insist upon judging my method by using your method, why shouldn't I use my method to evaluate your method? By the way, I note by gazing into my crystal ball that the scientific method is now in for a very bad run of luck."

The trouble with circular arguments is obvious: with an appropriate circular argument you can prove anything. In recent years, nevertheless, there have been several notable attempts to show how inductive rules can be supported inductively. The authors of such attempts try to show, of course, that their arguments are not circular. Although they argue persuasively, it seems to me that they do not succeed in escaping circularity.

One of the most widely discussed attempts to show that self-supporting inductive inferences are possible without circularity is due to Max Black.[7] Black correctly observes that the traditional fallacy of circular argument *(petitio principii)* involves assuming as a premise, often unwittingly, the conclusion that is to be proved. Thus, for example, a variety of "proofs" of Euclid's fifth postulate offered by mathematicians for about two millennia before the discovery of non-Euclidean geometry are circular in the standard fashion. They fail to show that the fifth postulate follows from the first four postulates alone; instead, they require in addition the assumption of a proposition equivalent to the proposition being demonstrated. The situation is quite different for self-supporting inductive arguments. The conclusion to be proved does not appear as one of the premises. Consider one of Black's examples: [8]

> 3. In most instances of the use of $\boldsymbol{R_2}$ in arguments with true premises examined in a wide variety of conditions, $\boldsymbol{R_2}$ has usually been successful.
>
> *Hence (probably):*
>
> In the next instance to be encountered of the use of $\boldsymbol{R_2}$ in an argument with true premises, $\boldsymbol{R_2}$ will be successful.

To say that an argument with true premises is successful is merely to say that it has a true conclusion. The rule $\boldsymbol{R_2}$ is

> To argue from ***Most instances of A's examined in a wide variety of conditions have been B to (probably) The next A to be encountered will be B.***

Inference 3 can be paraphrased suggestively, although somewhat inaccurately, as:

> 4. $\boldsymbol{R_2}$ has usually been successful in the past.
>
> *Hence (probably):*
>
> $\boldsymbol{R_2}$ will be successful in the next instance.

Inference 3 is governed by $\boldsymbol{R_2}$ that is, it conforms to the stipulation laid down by $\boldsymbol{R_2}$. $\boldsymbol{R_2}$ is *not* a premise, however, nor is any statement to the effect that all, some, or any future instances of $\boldsymbol{R_2}$ will be successful. As Lewis Carroll showed decisively, there is a fundamental distinction between premises and rules of inference.[9] Any inference, inductive or deductive, must conform to some rule, but neither the rule nor any statement about the rule is to be incorporated into the inference as an additional premise. If such additional premises were required, inference would be impossible. Thus, inference 3 is not a standard *petitio principii.*

What, then, are the requirements for a self-supporting argument? At least three are immediately apparent: (1) The argument must have true premises. (2) The argument must conform to a certain rule. (3) The conclusion of that argument must say something about the success or reliability of that rule in unexamined instances of its application. Inference 3 has these characteristics.

It is not difficult to find examples of deductive inferences with the foregoing characteristics.

> 5. If snow is white, then *modus*

ponens is valid. Snow is white.

Modus ponens is valid.

Inference 5 may seem innocuous enough, but the same cannot be said for the following inference:

6. If affinning the consequent is valid, then coal is black.
Coal is black.

Affirming the consequent is valid.

Like inference 5, inference 6 has true premises, it conforms to a certain rule, and its conclusion asserts the validity of that rule. Inference 5 did nothing to enhance our confidence in the validity of *modus ponens,* for we have far better grounds for believing it to be valid. Inference 6 does nothing to convince us that affirming the consequent is valid, for we know on other grounds that it is invalid. Arguments like 5 and 6 are, nevertheless, instructive. Both are circular in some sense, though neither assumes *as a premise* the conclusion it purports to establish. In deductive logic the situation is quite straightforward. A deductive inference establishes its conclusion if it has true premises and has a valid form. If either of these features is lacking the conclusion is not established by that argument. If the argument is valid but the premises are not true we need not accept the conclusion. If the premises are true but the argument is invalid we need not accept the conclusion. One way in which an argument can be circular is by adopting as a premise the very conclusion that is to be proved; this is the fallacy of *petitio principii* which I shall call "premise-circularity." Another way in which an argument can be circular is by exhibiting a form whose validity is asserted by the very conclusion that is to be proved; let us call this type of circularity "rule- circularity." Neither type of circular argument establishes its conclusion in any interesting fashion, for in each case the conclusiveness of the argument depends upon the assumption of the conclusion of that argument. Inferences 5 and 6 are not premise-circular; each is rule-circular. They are, nevertheless, completely question begging.

The situation in induction is somewhat more complicated, but basically the same.[10] Consider the following argument:

7. In most instances of the use of $\boldsymbol{R_3}$ in arguments with true premises examined in a wide variety of conditions, $\boldsymbol{R_3}$ has usually been *unsuccessful.*

Hence (probably):

In the next instance to be encountered of the use of $\boldsymbol{R_3}$ in an argument with true premises, $\boldsymbol{R_3}$ will be successful.

The rule $\boldsymbol{R_3}$ is

To argue from **Most instances of A's examined in a**

wide variety of conditions have been non-B to
(probably) The next A to be encountered will be B.

Inference 7 can be paraphrased as follows:

8. $\boldsymbol{R}_3$ has usually been unsuccessful in the past.
Hence (probably):
$\boldsymbol{R}_3$ will be successful in the next instance.

Notice that there is a perfect parallel between $\boldsymbol{R}_2$, 3, 4 on the one hand and $\boldsymbol{R}_3$, 7, 8 on the other. Since those instances in which $\boldsymbol{R}_2$ would be successful are those in which $\boldsymbol{R}_3$ would be unsuccessful, the premises of 3 and 4 describe the same state of affairs as do the premises of 7 and 8. Thus, the use of $\boldsymbol{R}_3$ in the next instance seems to be supported in the same manner and to the same extent as the use of $\boldsymbol{R}_2$ in the next instance. However, $\boldsymbol{R}_2$ and $\boldsymbol{R}_3$ conflict directly with each other. On the evidence that most Italians examined in a wide variety of conditions have been dark-eyed, $\boldsymbol{R}_2$ allows us to infer that the next Italian to be encountered will be dark-eyed, while $\boldsymbol{R}_3$ permits us to infer from the same evidence that he will have light-colored eyes. It appears then that we can construct self-supporting arguments for correct and incorrect inductive rules just as we can for valid and invalid deductive rules.

Black would reject self-supporting arguments for the fallacy of affirming the consequent and for a counterinductive rule like $\boldsymbol{R}_2$, because we know on independent grounds that such rules are faulty. Affirming the consequent is known to be fallacious, and the counterinductive method can be shown to be self-defeating, an additional requirement for a self-supporting argument is that the rule thus supported be one we have no independent reason to reject. Nevertheless, the fact that we can construct self-supporting arguments for such rules should give us pause. What if we had never realized that affirming the consequent is fallacious? What if we had never noticed anything wrong with the counterinductive method? Would arguments like 6, 7, and 8 have to be considered cogent? What about the standard inductive method? Is it as incorrect as the counterinductive method, but for reasons most of us have not yet realized?

It sounds as if a self-supporting argument is applicable only to rules we already know to be correct; as a matter of fact, this is the view Black holds. He has argued in various places that induction is in no need of a general justification.[11] He holds that calling into question of all inductive methods simultaneously results in a hopelessly skeptical position. He is careful to state explicitly at the outset of his discussion of self-supporting inductive arguments that he is not dealing with the view "that *no* inductive argument ought to be regarded as correct until a philosophical justification of induction has been provided."[12] At the conclusion he acknowledges, moreover, that "anybody who thinks he has good grounds for condemning all inductive arguments will also condemn inductive arguments in support of inductive rules."[13] Black is careful to state explicitly that self-supporting inductive arguments provide no answer to the problem of justification of induction as raised by Hume. What good, then, are self-supporting inductive arguments?

In deductive logic, correctness is an all-or-nothing affair. Deductive inferences are either totally valid or totally invalid; there cannot be such a thing as degree of validity. In inductive logic the situation is quite different. Inductive correctness does admit of degrees; one inductive conclusion may be more strongly supported than another. In this situation it is possible, Black claims, to have an inductive rule we know to be correct to some degree, but whose status can be enhanced by self-supporting arguments. We might think a rather standard inductive rule akin to Black's R_2 is pretty good, but through inductive investigation of its application we might find that it is extremely good--much better than we originally thought. Moreover, the inductive inferences we use to draw that conclusion might be governed by precisely the sort of rule we are investigating. It is also possible, of course, to find by inductive investigation that the rule is not as good as we believed beforehand.

It is actually irrelevant to the present discussion to attempt to evaluate Black's view concerning the possibility of increasing the justification of inductive rules by self-supporting arguments. The important point is to emphasize, because of the possibility of constructing self-supporting arguments for counterinductive rules, that the attempt to provide inductive support of inductive rules cannot, without vicious circularity, be applied to the problem of justifying induction from scratch. If there is any way of providing the beginnings of a justification, or if we could show that some inductive rule stands in no need of justification in the first instance, then it would be suitable to return to Black's argument concerning the increase of support. I am not convinced, however, that Black has successfully shown that there is a satisfactory starting place.

I have treated the problem of inductive justification of induction at some length, partly because other authors have not been as cautious as Black in circumscribing the limits of inductive justification of induction.[14] More important, perhaps, is the fact that it is extremely difficult, psychologically speaking, to shake the view that past success of the inductive method constitutes a genuine justification of induction. Nevertheless, the basic fact remains: Hume showed that inductive justifications of induction are fallacious, and no one has since proved him wrong.

2. The Complexity of Scientific Inference.

The idea of a philosopher discussing inductive inference in science is apt to arouse grotesque images in many minds. People are likely to imagine someone earnestly attempting to explain why it is reasonable to conclude that the sun will rise tomorrow morning because it always has done so in the past. There may have been a time when primitive man anticipated the dawn with assurance based only upon the fact that he had seen dawn follow the blackness of night as long as he could remember, but this primitive state of knowledge, if it ever existed, was unquestionably prescientific. This kind of reasoning bears no resemblance to science; in fact, the crude induction exhibits a complete absence of scientific understanding. Our scientific reasons for believing that the sun will rise tomorrow are of an entirely different kind. We understand the functioning of the solar system in terms of the laws of physics. We predict particular astronomical occurrences by means of these laws in conjunction with knowledge of particular initial conditions that prevail. Scientific laws and theories have the logical form of general statements, but they are seldom, if ever, simple generalizations from experience.

Consider Newton's gravitational theory: Any two bodies are mutually attracted by a force proportional to the product of their masses and inversely proportional to the square of the distance between their centers. Although general in form, this kind of statement is not established by generalization from instances. We do not go around saying, "Here are two bodies--the force between them is such and such; here are two more bodies--the force between them is such and such; etc. " Scientific theories are taken quite literally as hypotheses. They are entertained in order that their consequences may be drawn and examined. Their acceptability is judged in terms of these consequences. The consequences are extremely diverse--the greater the variety the better. For Newtonian theory, we look to such consequences as the behavior of Mars, the tides, falling bodies, the pendulum, and the torsion balance. These consequences have no apparent unity among themselves; they do not constitute a basis for inductive generalization. They achieve a kind of unity only by virtue of the fact that they are consequences of a single physical theory.

The type of inference I have been characterizing is very familiar; it is known as the *hypothetico-deductive method*.[15] It stands in sharp contrast to *induction by enumeration*, which consists in simple inductive generalization from instances. Schematically, the hypothetico-deductive method works as follows: From a general hypothesis and particular statements of initial conditions, a particular predictive statement is deduced. The statements of initial conditions, at least for the time, are accepted as true; the hypothesis is the statement whose truth is at issue. By observation we determine whether the predictive statement turned out to be true. If the predictive consequence is false, the hypothesis is disconfirmed. If observation reveals that the predictive statement is true, we say that the hypothesis is confirmed to some extent. A hypothesis is not, of course, conclusively proved by anyone or more positively confirming instances, but it may become highly confirmed. A hypothesis that is sufficiently confirmed is accepted, at least tentatively.

It seems undeniable that science uses a type of inference at least loosely akin to the hypothetico-deductive method.[16] This has led some people to conclude that the logic of science is thoroughly deductive in character. According to this view, the only nondeductive aspect of the situation consists in thinking up hypotheses, but this is not a matter of logic and therefore requires no justification. It is a matter of psychological ingenuity of discovery. Once the hypothesis has been discovered, by some entirely nonlogical process, it remains only to *deduce* consequences and check them against observation.

It is, of course, a fallacy to conclude that the premises of an argument must be true if its conclusion is true. This fact seems to be the basis for the quip that a logic text is a book that consists of two parts; in the first part (on deduction) the fallacies are explained, in the second part (on induction) they are committed. The whole trouble with saying that the hypothetico-deductive method renders the logic of science entirely deductive is that we are attempting to establish a *premise* of the deduction, not the conclusion. Deduction is an indispensable part of the logic of the hypothetico-deductive method, but it is not the only part. There is a fundamental and important sense in which the hypothesis must be regarded as a conclusion instead of a premise. Hypotheses (later perhaps called "theories" or "laws") are among the *results* of scientific investigation; science aims at establishing general statements about the world. Scientific prediction and explanation require such generalizations. While we are concerned with the status of the general hypothesis--whether we should accept it or reject it--the hypothesis must be treated as a conclusion to be supported by evidence, not as a premise lending support to other conclusions. The inference *from* observational evidence *to* hypothesis is surely not deductive. If this point is not already obvious it becomes clear the moment we recall that for any given body of observational data there is, in general, more than one hypothesis compatible with it. These alternative hypotheses differ in factual content and are incompatible with each other. Therefore, they cannot be deductive consequences of any consistent body of observational evidence.

We must grant, then, that science embodies a type of inference resembling the hypothetico-deductive method and fundamentally different from induction by enumeration. Hume, on the other hand, has sometimes been charged with a conception of science according to which the only kind of reasoning is induction by enumeration. His typical examples are cases of simple generalization of observed regularities, something like our example of the licorice-flavored black balls. In the past, water has quenched thirst; in the future, it will as well. In the past, fires have been hot; in the future, they will be hot. In the past, bread has nourished; in the future, it will do so likewise. It might be said that Hume, in failing to see the essential role of the hypothetico-deductive method, was unable to appreciate the complexity of the theoretical science of his own time, to say nothing of subsequent developments. This is typical, some might say, of the misunderstandings engendered by philosophers who undertake to discuss the logic of science without being thoroughly conversant with mathematics and natural science.

This charge against Hume (and other philosophers of induction) is ill-founded. It was part of Hume's genius to have recognized that the arguments he applied to simple enumerative induction apply equally to any kind of ampliative or nondemonstrative inference whatever. Consider the most complex kind of scientific reasoning--the most elaborate example of hypothetico-deductive inference you can imagine. Regardless of subtle features or complications, it is ampliative overall. The conclusion is a statement whose content exceeds the observational evidence upon which it is based. A scientific theory that merely summarized what had already been observed would not deserve to be called a theory. If scientific inference were not ampliative, science would be useless for prediction, postdiction, and explanation. The highly general results that are the pride of theoretical science would be impossible if scientific inference were not ampliative.

In presenting Hume's argument, I was careful to set it up so that it would apply to any kind of ampliative or nondemonstrative inference, no matter how simple or how complex. Furthermore, the distinction between valid deduction and nondemonstrative inference is completely exhaustive. Take any inference whatsoever. It must be deductive or nondemonstrative. Suppose it is nondemonstratwe. If we could justify it deductively it would cease to be nondemonstrative. To justify it nondemonstratively would presuppose an already justified type of nondemonstrative inference, which is precisely the problem at issue. Hume's argument does *not* break down when we consider forms more complex than simple enumeration. Although the word "induction" is sometimes used as a synonym for "induction by simple enumeration," I am not using it in that way. Any type of logically correct ampliative inference is induction; the problem of induction is to show that some particular form of ampliative inference is justifiable. It is in this sense that we are concerned with the problem of the justification of inductive inference.

A further misunderstanding is often involved in this type of criticism of Hume. There is a strong inclination to suppose that induction is regarded as the method by which scientific results are discovered.[17] Hume and other philosophers of induction are charged with the view that science has developed historically through patient collection of facts and generalization from them. I know of no philosopher--not even Francis Bacon!--who has held this view, although it is frequently at tacked in the contemporary literature.[18] The term "generalization" has an unfortunate ambiguity which fosters the confusion. In one meaning, "generalization" refers to an inferential process in which one makes a sort of mental transition from particulars to a universal proposition; in this sense, generalization is an act of generalizing--a process that yields general results. In another meaning, "generalization" simply refers to a universal type or proposition, without any reference to its source or how it was thought of. It is entirely possible for science to contain many generalizations (in the latter sense) without embodying any generalizations (in the former sense). As I said explicitly at the outset, the problem of induction I am discussing is a problem concerning justification, not discovery. The thesis I am defending--that science does embody induction in a logically indispensable fashion--has nothing to do with the history of science or the psychology of particular scientists. It is simply the claim that scientific inference is ampliative.

3. Deductivism.

One of the most interesting and controversial contemporary attempts to provide an account of the logic of science is Karl Popper's deductivism.[19] In the preceding section I discussed the view that the presence of the hypothetico-deductive method in the logic of science makes it possible to dispense with induction in science and, thereby, to avoid the problem of induction. I argued that the hypothetico-deductive method, since it is ampliative and nondemonstrative, is not strictly deductive; it is, in fact, inductive in the relevant sense. As long as the hypothetico-deductive method is regarded as a method for supporting scientific hypotheses, it cannot succeed in making science thoroughly deductive. Popper realizes this, so in arguing that deduction is the sole mode of inference in science he rejects the hypothetico-deductive method as a means for confirming scientific hypotheses. He asserts that induction plays no role whatever in science; indeed, he maintains that there is no such thing as correct inductive inference. Inductive logic is, according to Popper, a complete delusion. He admits the psychological fact that people (including himself) have faith in the uniformity of nature, but he holds, with Hume, that this can be no more than a matter of psychological fact. He holds, with Hume, that there can be no rational justification of induction, and he thinks Hume proved this point conclusively.

Popper's fundamental thesis is that falsifiability is the mark by which statements of empirical science are distinguished from metaphysical statements and from tautologies. The choice of falsifiability over verifiability as the criterion of demarcation is motivated by a long familiar fact--namely, it is possible to falsify a universal generalization by means of one negative instance, while it is impossible to verify a universal generalization by any limited number of positive instances. This, incidentally, is the meaning of the old saw which is so often made into complete nonsense: "The exception proves the rule." In this context, a rule is a universal generalization, and the term "to prove" means archaically "to test." The exception (i.e., the negative instance) proves (i.e., tests) the rule (i.e., the universal generalization), not by showing it to be true, but by showing it to be false. There is no kind of positive instance to prove (i.e., test) the rule, for positive instances are completely indecisive. Scientific hypotheses, as already noted, are general in form, so they are amenable to falsification but not verification.

Popper thus holds that falsifiability is the hallmark of empirical science. The aim of empirical science is to set forth theories to stand the test of every possible serious attempt at falsification. Scientific theories are hypotheses or conjectures; they are general statements designed to explain the world and make it intelligible, but they are never to be regarded as final truths. Their status is always that of tentative conjecture, and they must continually face the severest possible criticism. The function of the theoretician is to propose scientific conjectures; the function of the experimentalist is to devise every possible way of falsifying these theoretical hypotheses. The attempt to confirm hypotheses is no part of the aim of science.[20]

General hypotheses by themselves do not entail any predictions of particular events, but they do in conjunction with statements of initial conditions. The laws of planetary motion in conjunction with statements about the relative positions and velocities of the earth, sun, moon, and planets enable us to predict a solar eclipse. The mode of inference is deduction. We have a high degree of intersubjective agreement concerning the initial conditions, and we likewise can obtain intersubjective agreement as to whether the sun's disc was obscured at the predicted time and place. If the predicted fact fails to occur, the theory has suffered falsification. Again, the mode of inference is deduction. If the theory were true, then, given the truth of the statements of initial conditions, the prediction would have to be true. The prediction, as it happens, is false; therefore, the theory is false. This is the familiar principle of *modus tollens;* it is, according to Popper, the only kind of inference available for the acceptance or rejection of hypotheses, and it is clearly suitable for rejection only.

Hypothetico-deductive theorists maintain that we have a confirming instance for the theory if the eclipse occurs as predicted. Confirming instances, they claim, tend to enhance the probability of the hypothesis or give it inductive support. With enough confirming instances of appropriate kinds, the probability of the hypothesis becomes great enough to warrant accepting it as true--not, of course, with finality and certainty, but provisionally. With sufficient inductive support of this kind we are justified in regarding it as well established. Popper, however, rejects the positive account, involving as it does the notion of inductive support. If a hypothesis is tested and the result is negative, we can reject it. If the test is positive, all we can say is that we have failed to falsify it. We cannot say that it has been confirmed or that it is, because of the positive test result, more probable. Popper does admit a notion of *corroboration* of hypotheses, but that is quite distinct from confirmation. We shall come to corroboration presently. For the moment, all we have are successful or unsuccessful attempts at falsification; all we can say about our hypotheses is that they are falsified or unfalsified. This is as far as inference takes us; according to Popper, this is the limit of logic. Popper therefore rejects the hypothetico-deductive method as it is usually characterized and accepts only the completely deductive *modus tollens.*

Popper--quite correctly I believe--denies that there are absolutely basic and incorrigible protocol statements that provide the empirical foundation for all of science. He does believe that there are relatively basic observation statements about macroscopic physical occurrences concerning which we have a high degree of intersubjective agreement. Normally, we can accept as unproblematic such statements as, "There is a wooden table in this room," "The pointer on this meter stands between 325 and 350," and "The rope just broke and the weight fell to the floor." Relatively basic statements of this kind provide the observation base for empirical science. This is the stuff of which empirical tests of scientific theories are made.

Although Popper's basic statements must in the last analysis be considered hypotheses, falsifiable and subject to test like other scientific hypotheses, it is obvious that the kinds of hypotheses that constitute theoretical science are far more general than the basic statements. But now we must face the grim fact that valid deductive inference, although necessarily truth-preserving, is nonampliative.[21] It is impossible to deduce from accepted basic statements any conclusion whose content exceeds that of the basic statements themselves. Observation statements and deductive inference yield nothing that was not stated by the observation statements themselves. If science consists solely of observation statements and deductive inferences, then talk about theories, their falsifiability, and their tests is empty. The content of science is coextensive with the content of the statements used to describe what we directly observe. There are no general theories, there is no predictive content, there are no inferences to the remote past. Science is barren.

Consider a few simple time-honored examples. Suppose that the statement "All ravens are black" has been entertained critically and subjected to every attempt at falsification we can think of. Suppose it has survived all attempts at falsification. What is the scientific content of all this? We can say that "All ravens are black" has not been falsified, which is equivalent to saying that we have not observed a nonblack raven. This statement is even poorer in content than a simple recital of our color observations of ravens. To say that the hypothesis has not been falsified is to say less than is given in a list of our relevant observation statements. Or, consider the generalization, "All swans are white." What have we said when we say that this hypothesis has been falsified? We have said only that a nonwhite swan has been found. Again, the information conveyed by this remark is less than we would get from a simple account of our observations of swans.

Popper has never claimed that falsification by itself can establish scientific hypotheses. When one particular hypothesis has been falsified, many alternative hypotheses remain unfalsified. Likewise, there is nothing unique about a hypothesis that survives without being falsified. Many other unfalsified hypotheses remain to explain the same facts. Popper readily admits all of this. If science is to amount to more than a mere collection of our observations and various reformulation thereof, it must embody some other methods besides observation and deduction. Popper supplies that additional factor: *corroboration.*[22]

When a hypothesis has been falsified, it is discarded and replaced by another hypothesis which has not yet experienced falsification. Not all unfalsified hypotheses are on a par. There are principles of selection among unfalsified hypotheses. Again, falsifiability is the key. Hypotheses differ from one another with respect to the ease with which they can be falsified, and we can often compare them with respect to degree of falsifiability. Popper directs us to seek hypotheses that are as highly falsifiable as possible. Science, he says, is interested in bold conjectures. These conjectures must be consistent with the known facts, but they must run as great a risk as possible of being controverted by the facts still to be accumulated. Furthermore, the search for additional facts should be guided by the effort to find facts that will falsify the hypothesis.

As Popper characterizes falsifiability, the greater the degree of falsifiability of a hypothesis, the greater its content. Tautologies lack empirical content because they do not exclude any possible state of affairs; they are compatible with any possible world. Empirical statements are not compatible with every possible state of affairs; they are compatible with some and incompatible with others. The greater the number of possible states of affairs excluded by a statement, the greater its content, for the more it does to pin down our actual world by ruling out possible but nonactual states of affairs. At the same time, the greater the range of facts excluded by a statement--the greater the number of situations with which the statement is incompatible--the greater the risk it runs of being false. A statement with high content has more *potential falsifiers* than a statement with low content. For this reason, high content means high falsifiability. At the same time, content varies inversely with probability. The logical probability of a hypothesis is defined in terms of its range--that is, the possible states of affairs with which it is compatible. The greater the logical probability of a hypothesis, the fewer are its potential falsifiers. Thus, high probability means low falsifiability.

Hypothetico-deductive theorists usually recommend selecting, from among those hypotheses that are compatible with the available facts, the most probable hypothesis. Popper recommends the opposite; he suggests selecting the most falsifiable hypothesis. Thus, he recommends selecting a hypothesis with low probability. According to Popper, a highly falsifiable hypothesis which is severely tested becomes highly corroborated. The greater the severity of the tests--the greater their number and variety--the greater the corroboration of the hypothesis that survives them.

Popper makes it very clear that hypotheses are not regarded as true because they are highly corroborated. Hypotheses cannot be firmly and finally established in this or any other way. Furthermore, because of the inverse relation between falsifiability and probability, we cannot regard highly corroborated hypotheses as probable. To be sure, a serious attempt to falsify a hypothesis which fails does add to the corroboration of this hypothesis, so there is some similarity between corroboration and confirmation as hypothetico-deductive theorists think of it, but it would be a misinterpretation to suppose that increasing corroboration is a process of accumulating positive instances to increase the probability of the hypothesis.[23]

Nevertheless, Popper does acknowledge the need for a method of selecting among unfalsified hypotheses. He has been unequivocal in his emphasis upon the indispensability of far-reaching theory in science. Empirical science is not an activity of merely accumulating experiences; it is theoretical through and through. Although we do not regard any hypotheses as certainly true, we do accept them tentatively and provisionally. Highly corroborated hypotheses are required for prediction and explanation. From among the everpresent multiplicity of hypotheses compatible with the available evidence, we select and accept.

There is just one point I wish to make here regarding Popper's theory. It is not properly characterized as *deductivism.* Popper has not succeeded in purging the logic of science of all inductive elements. My reason for saying this is very simple. Popper furnishes a method for selecting hypotheses whose content exceeds that of the relevant available basic statements. Demonstrative inference cannot accomplish this task alone, for valid deductions are nonampliative and their conclusions cannot exceed their premises in content. Furthermore, Popper's theory does not pretend that basic statements plus deduction can give us scientific theory; instead, corroboration is introduced. Corroboration is a nondemonstrative form of inference. It is a way of providing for the acceptance of hypotheses even though the content of these hypotheses goes beyond the content of the basic statements. *Modus tollens* without corroboration is empty; *modus tollens* with corroboration is induction.

When we ask, "Why should we reject a hypothesis when we have accepted one of its potential falsifiers?" the answer is easy. The potential falsifier contradicts the hypothesis, so the hypothesis is false if the potential falsifier holds. That is simple deduction. When we ask, "Why should we accept from among all the unfalsified hypotheses one that is highly corroborated?" we have a right to expect an answer. The answer is some kind of justification for the methodological rule--for the method of corroboration. Popper attempts to answer this question.

Popper makes it clear that his conception of scientific method differs in important respects from the conceptions of many inductivists. I do not want to quibble over a word in claiming that Popper is, himself, a kind of inductivist. The point is not a trivial verbal one. Popper has claimed that scientific inference is exclusively deductive. We have seen, however, that demonstrative inference is not sufficient to the task of providing a reconstruction of the logic of the acceptance--albeit tentative and provisional--of hypotheses. Popper himself realizes this and introduces a mode of nondemonstrative inference. It does not matter whether we call this kind of inference "induction"; whatever we call it, it is ampliative and not necessarily truth preserving. Using the same force and logic with which Hume raised problems about the justification of induction, we may raise problems about the justification of any kind of nondemonstrative inference. As I argued in the preceding section, Hume's arguments are not peculiar to induction by enumeration or any other special kind of inductive inference; they apply with equal force to any inference whose conclusion can be false, even though it has true premises. Thus, it will not do to dismiss induction by enumeration on grounds of Hume's argument and then accept some other mode of nondemonstrative inference without even considering how Hume' s argument might apply to it. I am not arguing that Popper's method is incorrect.[24] I am not even arguing that Popper has failed in his attempt to justify this method. I do claim that Popper is engaged in the same task as many inductivists--namely, the task of providing some sort of justification for a mode of nondemonstrative inference. This enterprise, if successful, *is* a justification of induction.

. . . .

5. The Principle of Uniformity of Nature.

A substantial part of Hume's critique of induction rested upon his attack on the principle of the uniformity of nature. He argued definitively that the customary forms of inductive inference cannot be expected to yield correct predictions if nature fails to be uniform--if the future is not like the past--if like sensible qualities are not accompanied by like results.

> All inferences from experience suppose, as their foundation, that the future will resemble the past, and that similar powers will be conjoined with similar sensible qualities. If there be any suspicion that the course of nature may change, and that the past may be no rule for the future, all experience becomes useless, and can give rise to no inference or conclusion.[25]

He argued, moreover, that there is no logical contradiction in the supposition that nature is not uniform--that the regularities we have observed up to the present will fail in wholesale fashion in the future.

> It implies no contradiction that the course of nature may change, and that an

> object, seemingly like those which we have experienced, may be attended with different or contrary effects. May I not clearly and distinctly conceive that a body, falling from the clouds, and which, in all other respects resembles snow, has yet the taste of salt or feeling of fire? Is there any more intelligible proposition than to affirm, that all the trees will flourish in December and January, and decay in May and June? Now whatever is intelligible, and can be distinctly conceived, implies no contradiction, and can never be proved false by any demonstrative argument. . . .[26]

He argues, in addition, that the principle of uniformity of nature cannot be established by an inference from experience: "It is impossible, therefore, that any arguments from experience can prove this resemblance of the past to the future; since all these arguments are founded on the sup- position of that resemblance. "[27] Throughout Hume's discussion there is, however, a strong suggestion that we might have full confidence in the customary inductive methods if nature were known to be uniform.

Kant attempted to deal with the problem of induction in just this way, by establishing a principle of uniformity of nature, in the form of the principle of universal causation, as a synthetic *a priori* truth. Kant claimed, in other words, that every occurrence is governed by causal regularities, and this general characteristic of the universe can be established by pure reason, without the aid of any empirical evidence. He did not try to show that the principle of universal causation is a principle of logic, for to do so would have been to show that it was analytic--not synthetic--and thus lacking in factual content. He did not reject Hume's claim that there is no logical contradiction in the statement that nature is not uniform; he did not try to prove his principle of universal causation by deducing a contradiction from its denial. He did believe, however, that this principle, while not a proposition of pure logic, is necessarily true nevertheless. Hume, of course, argued against this alternative as well. He maintained not only that the uniformity of nature is not a logical or analytic truth, but also that it cannot be any other kind of *a priori* truth either. Even before Kant had enunciated the doctrine of synthetic *a priori* principles, Hume had offered strong arguments against them:

> I shall venture to affirm, as a general proposition, which admits of no exception, that the knowledge of this relation [of cause and effect] is not, in any instance, attained by reasonings *a priori*.[28]
>
> Adam, though his rational faculties be supposed, at the very first, entirely perfect, could not have inferred from the fluidity and transparency of water that it would suffocate him, or from the light and warmth of fire that it would consume him.[29]
>
> When we reason *a priori*, and consider merely any object or cause, as it appears to the mind, independent of all observation, it never could suggest to us the notion of any distinct object, such as its effect; much less, show us the inseparable and inviolable connexion between them. A man must be very sagacious who could discover by reasoning that crystal is the effect of heat, and ice of cold, without being previously acquainted with the operation of

> these qualities.[30]
>
> Now whatever is intelligible, and can be distinctly conceived. . . can never be proved false by any. . . abstract reasoning a priori.[31]

Hume argues, by persuasive example and general principle, that nothing about the causal structure of reality can be established by pure reason. He poses an incisive challenge to those who would claim the ability to establish a priori knowledge of a particular causal relation or of the principle of universal causation. In the foregoing discussion of synthetic a priori statements, I have given reasons for believing that Kant failed to overcome Hume's previous objections.

There is, however, another interesting issue that arises in connection with the principle of uniformity of nature. Suppose it could be established--never mind how--prior to a justification of induction. Would it then provide an adequate basis for a justification of induction? The answer is, I think, negative.[32]

Even if nature is uniform to some extent, it is not absolutely uniform. The future is something like the past, but it is somewhat different as well. Total and complete uniformity would mean that the state of the universe at any given moment is the same as its state at any other moment. Such a universe would be a changeless, Parmenidean world. Change obviously does occur, so the future is not exactly like the past. There are some uniformities, it appears, but not a complete absence of change. The problem is how to ferret out the genuine uniformities. As a matter of actual fact, there are many uniformities *within experience* that we take to be mere coincidences, and there are others that seem to represent genuine causal regularities. For instance, in every election someone finds a precinct, say in Maryland, which has always voted in favor of the winning presidential candidate. Given enough precincts, one expects this sort of thing by sheer chance, and we classify such regularities as mere coincidences. By contrast, the fact that glass windowpanes break when bricks are hurled at them is more than mere coincidence. Causal regularities provide a foundation for inference from the observed to the unobserved; coincidences do not. We can predict with some confidence that the next glass window pane at which a brick is hurled will break; we take with a grain of salt the prediction of the outcome of a presidential election early on election night when returns from the above-mentioned precinct are in. The most that a principle of uniformity of nature could say is that there are some uniformities that persist into the future; if it stated that every regularity observed to hold within the scope of our experience also holds universally, it would be patently false. We are left with the problem of finding a sound basis for distinguishing between mere coincidence and genuine causal regularity.

Kant's principle of universal causation makes a rather weak and guarded statement. It asserts only that there exist causal regularities: "Everything that happens presupposes something from which it follows according to some rule." For each occurrence it claims only the existence of *some* prior cause and *some* causal regularity. It gives no hint as to how we are to find the prior cause or how we are to identify the causal regularity. It therefore provides no basis upon which to determine whether the inductive inferences we make are correct or incorrect. It would be entirely consistent with Kant's principle for us always to generalize on the basis of observed coincidences and always to fail to generalize on the basis of actual causal relations. It would be entirely consistent with Kant's principle, moreover, for us always to cite a coincidentally preceding event as the cause instead of the event that is the genuine cause. Kant's principle, even if it could be established, would not help us to justify the assertion that our inductive inferences would always or usually be correct. It would provide no criterion to distinguish sound from unsound inductions. Even if Kant's program had succeeded in establishing a synthetic *a priori* principle of universal causation, it would have failed to produce a justification of induction.

. . . .

7. A Probabilistic Approach.
It may seem strange in the extreme that this discussion of the problem of induction has proceeded at such great length without seriously bringing in the concept of probability. It is very tempting to react immediately to Hume's argument with the admission that we do not have *knowledge* of the unobserved. Scientific results are not established with absolute certainty. At best we can make probabilistic statements about unobserved matters of fact, and at best we can claim that scientific generalizations and theories are highly confirmed. We who live in an age of scientific empiricism can accept with perfect equanimity the fact that the quest for certainty is futile; indeed, our thanks go to Hume for helping to destroy false hopes for certainty in science.

Hume's search for a justification of induction, it might be continued, was fundamentally misconceived. He tried to find a way of proving that inductive inferences with true premises would have *true* conclusions. He properly failed to find any such justification precisely because it is the function of *deduction* to prove the truth of conclusions, given true premises. Induction has a different function. An inductive inference with true premises establishes its conclusions as *probable.* No wonder Hume failed to find a justification of induction. He was trying to make induction into deduction, and he succeeded only in proving the platitude that induction is not deduction.[33] If we want to justify induction, we must show that inductive inferences establish their conclusions as probable, not as true.

The foregoing sort of criticism of Hume's arguments is extremely appealing, and it has given rise to the most popular sort of attempt, currently, to deal with the problem.[34] In order to examine this approach, we must consider, at least superficially, the meaning of the concept of probability. Two basic meanings must be taken into account at present.

One leading probability concept identifies probability with frequency--roughly, the probable is that which happens often, and the improbable is that which happens seldom. Let us see what becomes of Hume's argument under this interpretation of probability. If we were to claim that inductive conclusions are probable in this sense, we would be claiming that inductive inferences with true premises often have true conclusions, although not always. Hume's argument shows, unhappily, that this claim cannot be substantiated. It was recognized long before Hume that inductive inferences cannot be expected always to lead to the truth. Hume's argument shows, not only that we cannot justify the claim that *every* inductive inference with true premises will have a true conclusion, but also, that we cannot justify the claim that *any* inductive inference with true premises will have a true conclusion. Hume's argument shows that, for all we can know, every inductive inference made from now on might have a false conclusion despite true premises.

Thus, Hume has proved, we can show neither that inductive inferences establish their conclusions as true nor that they establish their conclusions as probable in the frequency sense. The introduction of the frequency concept of probability gives no help whatever in circumventing the problem of induction, but this is no surprise, for we should not have expected it to be suitable for this purpose.

A more promising probability concept identifies probability with degree of rational belief. To say that a statement is probable in this sense means that one would be rationally justified in believing it; the degree of probability is the degree of assent a person would be rationally justified in giving. We are not, of course, referring to the degree to which anyone *actually* believes in the statement, but rather to the degree to which one could *rationally* believe it. Degree of actual belief is a purely psychological concept, but degree of rational belief is determined objectively by the evidence. To say that a statement is probable in this sense means that it is supported by evidence. But, so the argument goes, if a statement is the conclusion of an inductive inference with true premises, it *is* supported by evidence--by inductive evidence--this is part of what it *means* to be supported by evidence. The very concept of evidence depends upon the nature of induction, and it becomes incoherent if we try to divorce the two. Trivially, then, the conclusion of an inductive inference is probable under this concept of probability. To ask, with Hume, if we should accept inductive conclusions is tantamount to asking if we should fashion our beliefs in terms of the evidence, and this, in turn, is tantamount to asking whether we should be rational. In this way we arrive at an "ordinary language dissolution" of the problem of induction. Once we understand clearly the meanings of such key terms as "rational," "probable," and "evidence," we see that the problem arose out of linguistic confusion and evaporates into the question of whether it is rational to be rational. Such tautological questions, if meaningful at all, demand affirmative answers.

Unfortunately, the dissolution is not satisfactory.[35] Its inadequacy can be exhibited by focusing upon the concept of inductive evidence and seeing how it figures in the foregoing argument. The fundamental difficulty arises from the fact that the very notion of inductive evidence is determined by the rules of inductive inference. If a conclusion is to be supported by inductive evidence, it must be the conclusion of a correct inductive inference with true premises. Whether the inductive inference is correct depends upon whether the rule governing that inference is correct. The relation of inductive evidential support is, therefore, inseparably bound to the correctness of rules of inductive inference. In order to be able to say whether a given statement is supported by inductive evidence we must be able to say which inductive rules are correct.

For example, suppose that a die has been thrown a large number of times, and we have observed that the side two came up in one sixth of the tosses. This is our "evidence" ***e***. Let ***h*** be the conclusion that, "in the long run," side two will come up one sixth of the times. Consider the following three rules:

> 1. (Induction by enumeration.) Given ***m/n*** of observed ***A*** are ***B***, to infer that the "long run" relative frequency of ***B*** among ***A*** is **m/n**.
>
> 2. (*A priori* rule.) Regardless of observed frequencies, to infer that the "long run" relative frequency of ***B*** among ***A*** is ***1/k***, where ***k*** is the number of

possible outcomes-six in the case of the die.

3. (Counterinductive rule.) Given ***m/n*** of observed ***A*** are ***B***, to infer that the "long run" relative frequency of ***B*** among ***A*** is ***(n-m)/n***.

Under Rule 1, ***e*** is positive evidence for ***h;*** under Rule 2, ***e*** is irrelevant to ***h;*** and under Rule 3, ***e*** is negative evidence for ***h***. In order to say which conclusions are supported by what evidence, it is necessary to arrive at a decision as to what inductive rules are acceptable. If Rule 1 is correct, the evidence ***e*** supports the conclusion ***h***. If Rule 2 is correct, we are justified in drawing the conclusion ***h***, but this is entirely independent of the observational evidence ***e;*** the same conclusions would have been sanctioned by Rule 2 regardless of observational evidence. If Rule 3 is correct, we are not only prohibited from drawing the conclusion ***h***, but also we are permitted to draw a conclusion ***h'*** which is logically incompatible with***h***. Whether a given conclusion is *supported by evidence--whether* it would be *rational to believe* it on the basis of given evidence--whether it is *made probable* by virtue of its relation to given evidence--depends upon selection of the correct rule or rules from among the infinitely many rules we might conceivably adopt.

The problem of induction can now be reformulated as a problem about evidence. What rules ought we to adopt to determine the nature of inductive evidence? What rules provide suitable concepts of inductive evidence? If we take the customary inductive rules to define the concept of inductive evidence, have we adopted a proper concept of evidence? Would the adoption of some alternative inductive rules provide a more suitable concept of evidence? These are genuine questions which need to be answered.[36]

We find, moreover, that what appeared earlier as a pointless question now becomes significant and difficult. If we take the customary rules of inductive inference to provide a suitable definition of the relation of inductive evidential support, it makes considerable sense to ask whether it is rational to believe on the basis of evidence as thus defined rather than to believe on the basis of evidence as defined according to other rules. For instance, I believe that the *a priori* rule and the counterinductive rule mentioned above are demonstrably unsatisfactory, and hence, they demonstrably fail to provide a suitable concept of inductive evidence. The important point is that something concerning the selection from among possible rules needs demonstration and is amenable to demonstration.

There is danger of being taken in by an easy equivocation. One meaning we may assign to the concept of inductive evidence is, roughly, the basis on which we ought to fashion our beliefs. Another meaning results from the relation of evidential support determined by whatever rule of inductive inference we adopt. It is only by supposing that these two concepts are the same that we suppose the problem of induction to have vanished. The problem of induction is still there; it is the problem of providing adequate grounds for the selection of inductive rules. We want the relation of evidential support determined by these rules to yield a concept of inductive evidence which is, in fact, the basis on which we ought to fashion our beliefs.[37]

We began this initially promising approach to the problem of the justification of induction by introducing the notion of probability, but we end with a dilemma. If we take "probability" in the frequency sense, it is easy to see why it is advisable to accept probable conclusions in preference to improbable ones. In so doing we shall be right more often. Unfortunately, we cannot show that inferences conducted according to any particular rule establish conclusions that are probable in this sense. If we take "probability" in a nonfrequency sense it may be easy to show that inferences which conform to our accepted inductive rules establish their conclusions as probable. Unfortunately, we can find no reason to prefer conclusions which are probable in this sense to those that are improbable. As Hume has shown, we have no reason to suppose that probable conclusions will often be true and improbable ones will seldom be true. This dilemma is Hume's problem of induction all over again. We have been led to an interesting reformulation, but it is only a reformulation and not a solution.

8. Pragmatic Justification.

Of all the solutions and dissolutions proposed to deal with Hume's problem of induction, Hans Reichenbach's attempt to provide a pragmatic justification seems to me the most fruitful and promising.[38] This approach accepts Hume's arguments up to the point of agreeing that it is impossible to establish, either deductively or inductively, that any inductive inferences will ever again have true conclusions. Nevertheless, Reichenbach claims, the standard method of inductive generalization can be justified. Although its *success* as a method of prediction cannot be established in advance, it can be shown to be superior to any alternative method of prediction.

The argument can be put rather simply. Nature may be sufficiently uniform in suitable respects for us to make successful inductive inferences from the observed to the unobserved. On the other hand, for all we know, she may not. Hume has shown that we cannot prove in advance which case holds. All we can say is that nature mayor may not be uniform-if she is, induction works; if she is not, induction fails. Even in the face of our ignorance about the uniformity of nature, we can ask what would happen if we adopted some radically different method of inference. Consider, for instance, the method of the crystal gazer. Since we do not know whether nature is uniform or not, we must consider both possibilities. If nature is uniform, the method of crystal gazing might work successfully, or it might fail. We cannot prove *a priori* that it will not work. At the same time, we cannot prove *a priori* that it will work, even if nature exhibits a high degree of uniformity. Thus, in case nature is reasonably uniform, the standard inductive method *must* work while the alternative method of crystal gazing *mayor may not* work. In this case, the superiority of the standard inductive method is evident. Now, suppose nature lacks uniformity to such a degree that the standard inductive method is a complete failure. In this case, Reichenbach argues, the alternative method must likewise fail. Suppose it did not fail--suppose, for instance, that the method of crystal gazing worked consistently. This would constitute an important relevant uniformity that could be exploited inductively. If a crystal gazer had consistently predicted future occurrences, we could infer inductively that he has a method of prediction that will enjoy continued success. The inductive method would, in this way, share the success of the method of crystal gazing, and would therefore be, contrary to hypothesis, successful. Hence, Reichenbach concludes, the standard inductive method will be successful *if any other method could succeed.* As a result, we have everything to gain and nothing to lose by adopting the inductive method. If any method works, induction works. If we adopt the inductive method and it fails, we have lost nothing, for any other method we might have adopted would likewise have failed. Reichenbach does not claim to prove that nature is uniform, or that the standard inductive method will be successful. He does not postulate the uniformity of nature. He tries to show that the inductive method is the best method for ampliative inference, whether it turns out to be successful or not.

This ingenious argument, although extremely suggestive, is ultimately unsatisfactory. As I have just presented it, it is impossibly vague. I have not specified the nature of the standard inductive method. I have not stated with any exactness what constitutes success for the inductive method or any other. Moreover, the uniformity of nature is not an all-or-none affair. Nature appears to be uniform to some extent and also to be lacking in uniformity to some degree. As we have already seen, it is not easy to state a principle of uniformity that is strong enough to assure the success of inductive inference and weak enough to be plausible. The vagueness of the foregoing argument is not, however, its fundamental drawback. It can be made precise, and I shall do so below in connection with the discussion of the frequency interpretation of probability.[39] When it is made precise,...it suffers the serious defect of equally justifying too wide a variety of rules for ampliative inference.

I have presented Reichenbach's argument rather loosely in order to make intuitively clear its basic strategy. The sense in which it is a pragmatic justification should be clear. Unlike many authors who have sought a justification of induction, Reichenbach does not try to prove the truth of any synthetic proposition. He recognizes that the problem concerns the justification of a rule, and rules are neither true nor false. Hence, he tries to show that the adoption of a standard inductive rule is practically useful in the attempt to learn about and deal with the unobserved. He maintains that this can be shown even though we cannot prove the truth of the assertion that inductive methods will lead to predictive success. This pragmatic aspect is, it seems to me, the source of the fertility of Reichenbach's approach. Even though his argument does not constitute an adequate justification of induction, it seems to me to provide a valid core from which we may attempt to develop a more satisfactory justification.

III. Significance of the Problem

Hume's problem of induction evokes, understandably, a wide variety of reactions. It is not difficult to appreciate the response of the man engaged in active scientific research or practical affairs who says, in effect, "Don't bother me with these silly puzzles; I'm too busy doing science, building bridges, or managing affairs of state." No one, including Hume, seriously suggests any suspension of scientific investigation or practical decision pending a solution of the problem of induction. The problem concerns the *foundations* of science. As Hume eloquently remarks in *Enquiry Concerning Human Understanding:*

> Let the course of things be allowed hitherto ever so regular; that alone, without some new argument or inference, proves not that, for the future, it will continue so. In vain do you pretend to have learned the nature of bodies from your past experience. Their secret nature, and consequently all their effects and influence, may change, without any change in their sensible qualities. This happens sometimes, and with regard to some objects: Why may it not happen always, and with regard to all objects? What logic, what process of argument secures you against this supposition? My practice, you say, refutes my doubts. But you mistake the purport of my question. As an agent, I am quite satisfied in the point; but as a philosopher, who has some

> share of curiosity, I will not say scepticism, I want to learn the foundation of this inference.

We should know by now that the foundations of a subject are usually established long after the subject has been well developed, not before. To suppose otherwise would be a glaring example of "naive first-things-firstism."[40]

Nevertheless, there is something intellectually disquieting about a serious gap in the foundations of a discipline, and it is especially disquieting when the discipline in question is so broad as to include the whole of empirical science, all of its applications, and indeed, all of common sense. As human beings we pride ourselves on rationality--so much so that for centuries rationality was enshrined as the very essence of humanity and the characteristic that distinguishes man from the lower brutes. Questionable as such pride may be, our intellectual consciences should be troubled by a gaping lacuna in the structure of our knowledge and the foundations of scientific inference. I do not mean to suggest that the structure of empirical science is teetering because of foundational difficulties; the architectural metaphor is really quite inappropriate. I do suggest that intellectual integrity requires that foundational problems not be ignored.

Each of two opposing attitudes has its own immediate appeal. One of these claims that the scientific method is so obviously the correct method that there is no need to waste our time trying to show that this is so. There are two difficulties. First, we have enough painful experience to know that the appeal to obviousness is dangerously likely to be an appeal to prejudice and superstition. What is obvious to one age or culture may well turn out, on closer examination, to be just plain false. Second, if the method of science is so obviously superior to other methods we might adopt, then I should think we ought to be able to point to those characteristics of the method by which it gains its obvious superiority.

The second tempting attitude is one of pessimism. In the face of Hume's arguments and the failure of many attempts to solve the problem, it is easy to conclude that the problem is hopeless. Whether motivated by Hume's arguments or, as is probably more often the case, by simple impatience with foundational problems, this attitude seems quite widespread. It is often expressed by the formula that science is, at bottom, a matter of faith. While it is no part of my purpose to launch a wholesale attack on faith as such, this attitude toward the foundations of scientific inference is unsatisfactory. The crucial fact is that science makes a *cognitive claim,* and this cognitive claim is a fundamental part of the rationale for doing science at all. Hume has presented us with a serious challenge to that cognitive claim. If we cannot legitimize the cognitive claim, it is difficult to see what reason remains for doing science. Why not turn to voodoo, which would be simpler, cheaper, less time consuming, and more fun?

If science is basically a matter of faith, then the scientific faith exists on a par with other faiths. Although we may be culturally conditioned to accept this faith, others are not. Science has no ground on which to maintain its *cognitive* superiority to any form of irrationalism, however repugnant. This situation is, it seems to me, intellectually and socially undesirable. We have had enough experience with various forms of irrationalism to recognize the importance of being able to distinguish them logically from genuine science. I find it intolerable to suppose that a theory of biological evolution, supported as it is by extensive scientific evidence, has no more rational foundation than has its rejection by ignorant fundamentalists. I, too, have faith that the scientific method is especially well suited for establishing knowledge of the unobserved, but I believe this faith should be justified. It seems to me extremely important that some people should earnestly seek a solution to this problem concerning the foundations of scientific inference.

One cannot say in advance what consequences will follow from a solution to a foundational problem. It would seem to depend largely upon the nature of the solution. But a discipline with well-laid foundations is surely far more satisfactory than one whose foundations are in doubt. We have only to compare the foundationally insecure calculus of the seventeenth and eighteenth centuries with the calculus of the late nineteenth century to appreciate the gains in elegance, simplicity, and rigor. Furthermore, the foundations of calculus provided a basis for a number of other developments, interesting in their own right and *greatly extending the power and fertility of the original theory.* Whether similar extensions will occur as a result of a satisfactory resolution of Hume's problem is a point on which it would be rash to hazard any prediction, but we know from experience that important consequences result from the most unexpected sources. The subsequent discussion of the foundations of probability will indicate directions in which some significant consequences may be found, but for the moment it will suffice to note that a serious concern for the solution of Hume's problem cannot fail to deepen our understanding of the nature of scientific inference. This, after all, is the ultimate goal of the whole enterprise.

Notes

This book *[The Foundations of Scientific Inference]* is based upon five lectures in the Philosophy of Science Series at the University of Pittsburgh. The first two lectures, *Foundations of Scientific Inference:* I. *The Problem of Induction,* II.*Probability and Induction,* were presented in March 1963. The next two lectures, *Inductive Inference in Science:* I. *Hypothetico-Deductive Arguments,* II. *Plausibility Arguments,* were delivered in October 1964. The final lecture, *A Priori Knowledge,* was given in October 1965. The author wishes to express his gratitude to the National Science Foundation and the Minnesota Center for Philosophy of Science for support of research on inductive logic and probability.

1. David Hume, *Enquiry Concerning Human Under- standing,* see IV, I.

2. Ibid.

3. For a more detailed account of the relation between deductive validity and factual content, see p.24.

4. The problem of the synthetic a priori is discussed earlier, in sec. II, 4, pp. 27-40 [of Salmon's *Foundations of Scientific Inference,* from which this selection is excerpted].

5. Hume, *Human Understanding.*

6. Ibid.

7. Max Black, *Problems of Analysis* (Ithaca: Cornell University Press, 1954), Chap. II.

8. Ibid., pp. 196-97.

9. Lewis Carroll, "What the Tortoise Said to Achilles," in *The Complete Works of Lewis Carroll* (New York: Random House, n.d.).

10. I presented the following self-supporting argument for the counterinductive method in "Should We Attempt to Justify Induction?" *Philosophical Studies*, 8 (April 1957), pp. 45-47. Max Black in "Self-supporting Inductive Arguments,"*Models and Metaphors* (Ithaca: Cornell University Press, 1962), Chap. 12, replies to my criticism, but he does not succeed in shaking the basic point: The counterinductive rule is related to its self-supporting argument in precisely the same way as the standard inductive rule is related to its self-supporting argument. This is the "cash value" of claiming that the self-supporting argument is circular. Peter Achinstein, "The Circularity of a Self- supporting Inductive Argument," *Analysis,* 22 (June 1962), considers neither my formulation nor Black's answer sufficient, so he makes a further attempt to show circularity. Black's reply is found in "Self-Support and Circularity: A Reply to Mr. Achinstein," *Analysis.* 23 (December 1962). Achinstein's rejoinder is "Circularity and Induction," *Analysis*, 23 (June 1963).

11. Max Black, "The Justification of Induction," *Language and Philosophy* (Ithaca: Cornell Univer- sity Press, 1949), Chap. 3. The view he expresses in this essay, I believe, is closely related to the "probabilistic approach" I discuss in sec. II, 7, pp. 280-82.

12. Max Black, *Problems of Analysis,* p. 191.

13. Ibid., p. 206.

14. Compare Richard Bevan Braithwaite, *Scientific Explanation* (New York: Hmper & Row, 1960), Chap. 8. I think the same general view is to be found in A. J. Ayer, *The Problem of Knowledge* (Baltimore: Penguin Books, 1956), p. 75. I have discussed Ayer's view in "The Concept of Inductive Evidence," *American Philosophical Quarterly,* 2 (October 1965).

15. See Braithwaite for a systematic exposition of this conception.

16. Sec. VII, pp. 108-31 [of Salmon's *Foundations of Scientific Inference]*. "The Confirmation of Scientific Hypotheses" is devoted to a detailed analysis of this type of inference.

17. See John Patrick Day, *Inductive Probability* (New York: Humanities Press, 1961), p. 6. The nineteenth-century notion that induction is a process of discovery and the problem of whether there can be a logic of discovery are discussed earlier in sec. VII, pp. 109-14 [of Salmon's *Foundations of Scientific Inference.]*

18. See e.g., Karl R. Popper, *The Logic of Scientific Discovery* (New York: Basic Books, 1959), sec. 30, and Thomas S. Kuhn, *The Structure of Scientific Revolutions* (Chicago: University of Chicago Press, 1962). A fuller discussion of the relations among such concepts as deductive validity and content is given earlier in sec. II, 4, especially p. 33 [of Salmon's *Foundations of Scientific Inference]*.

19. The most comprehensive statement of Popper's position is to be found in *The Logic of Scientific Discovery.* This is the English translation, with additions, of Karl R. Popper, *Logik der Forschung* (Vienna, 1934).

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

21. I believe Popper openly acknowledges the non-ampliative character of deduction. See "Why Are the Calculi of Logic and Arithmetic Applicable to Reality," in Karl R. Popper, *Conjectures* and *Refutations* (New York: Basic Books, 1962), Chap. 9.

22. See *The Logic of Scientific Discovery,* Chap. 10. 23. Ibid., p. 270.

24. I return to Popper's methodological views in the discussion of confirmation in sec. VII [of *Foundations of Scientific Inference.]* In that context I shall exhibit what I take to be the considerable valid content of Popper's account of the logic of science. See pp. 114-21.

25. David Hume, *Human Understanding,* sec. IV.

26. Ibid.

27. Ibid.

28. Ibid.

29. Ibid.

30. Ibid.

31. Ibid.

32. Wesley C. Salmon, "The Uniformity of Nature," *Philosophy* and *Phenomenological Research,* 14 (September 1953).

33. Max Black, "The Justification of Induction," in *Language* and *Philosophy.*

34. Among the authors who subscribe to approaches similar to this are A. *J.* Ayer, *Language. Truth and Logic* (New York: Dover Publications, 1952); Paul Edwards, "Russell's Doubts about Induction," *Mind,* 58 (1949), pp. 141-63; Asher Moore, "The Principle of Induction," *Journal of Philosophy,* 49 (1952), pp. 741-58; Arthur Pap, *Elements of Analytic Philosophy* (New York: Macmillan, 1949), and *An Introduction to the Philosophy of Science;* and P. F. Strawson, *Introduction to Logical Theory* (London: Methuen, 1952).

35. I have criticized this type of argument at some length in "Should We Attempt To Justify Induction?" *Philosophical Studies,* 8 (April 1957), and in "The Concept of Inductive Evidence," *American Philosophical Quarterly*, 2 (October 1965). This latter article is part of a "Symposium on Inductive Evidence" in which Stephen Barker and Henry E. Kyburg, Jr., defend against the attack. See their comments and my rejoinder.

36. This point has enormous import for any attempt to construct an inductive justification of induction. To decide whether the fact that induction has been successful in the past is positive evidence, negative evidence, or no evidence at all begs the very question at issue.

37. As I attempted to show in "Should We Attempt to Justify Induction?" this equivocation seems to arise out of a failure to distinguish *validation* and *vindication.* This crucial distinction is explicated by Herbert Feigl, "De Principiis non Disputandum . . . ?" in *Philosophical Analysis*, ed. Max Black (Ithaca: Cornell U. Press, 1950).

38. Hans Reichenbach, *Experience and Prediction* (Chicago: University of Chicago Press, 1938), Chap. 5, and *The Theory of Probability* (Berkeley: University of California Press, 1949), Chap. II.

39. Sec. V, 5, pp. 83-96 [of Salmon's Foundations of Scientific Inference].

40. Leonard J. Savage, *The Foundations of Statistics* (New York: Wiley, 1954), p. 1.