II

Probability and Inductive Logic

II.1. INTRODUCTION. What is logic as a whole and how do inductive and deductive logic fit into the big picture? How does inductive logic use the concept of probability? What does logic have to do with arguments? In this chapter we give a preliminary discussion of these large questions to provide a perspective from which to approach the rest of the book.

II.2. ARGUMENTS. The word "argument" is used to mean several different things in the English language. We speak of two people having an argument, of one person advancing an argument, and of the value of a mathematical function as depending on the values of its arguments. One of these various senses of "argument" is selected and refined by the logician for the purposes at hand.

When we speak of a person advancing an argument, we have in mind his giving certain reasons designed to persuade us that a certain claim he is making is correct. Let us call that claim which is being argued for the *conclusion* of the argument, and the reasons advanced in support of it the *premises*. If we now abstract from the concrete situation in which one person is trying to convince others and consider the bare bones of this conception of an argument, we arrive at the following definition: An *argument* is a list of *sentences*, one of which is designated as the conclusion, and the rest of which are designated as premises.

But if we consider the matter closely, we see that this definition will not do. Questions, commands, and curses can all be expressed by sentences, but they do not make factual claims nor can they stand as reasons supporting such claims. Suppose someone said, "The Dirty Sox star pitcher has just broken both his arms and legs, their catcher has glaucoma, their entire outfield has come down with bubonic plague, and their shortstop has been deported. Therefore, they cannot possibly win the pennant." He would clearly be advancing an argument, to the effect that the Dirty Sox cannot win the pennant. But if someone said, "How's your sister? Stand up on the table. May you perish in unspeakable slime!" he would, whatever else he was doing, *not* be advancing an argument. That is, he would not be advancing evidence in support of a *factual claim*.

Let us call a sentence that makes a definite factual claim a *statement*. "Hannibal crossed the Alps," "Socrates was a corruptor of youth," "Every body attracts every other body with a force directly proportional to the sum of their

masses and inversely proportional to the square of their distance," and "The moon is made of avocado paste" are all statements, some true, some false. We may now formulate the logician's definition of an argument:

> **Definition 1:** An argument is a list of *statements*, one of which is designated as the conclusion and the rest of which are designated as premises.

In ordinary speech we seldom come right out and say, "A, B, C are my premises and D is my conclusion." However, there are several "indicator words" that are commonly used in English to point out which statement is the conclusion and which are the premises. The word "therefore" signals that the premises have been run through, and that the conclusion is about to be presented (as in the Dirty Sox example). The words "thus," "consequently," "hence," "so," and the phrase "it follows that" function in exactly the same way.

In ordinary discourse the conclusion is sometimes stated first, followed by the premises advanced in support of it. In these cases, different indicator words are used. Consider the following argument: "Socrates is basically selfish, since after all Socrates is a man, and all men are basically selfish." Here the conclusion is stated first and the word "since" signals that reasons in support of that conclusion follow. The words "because" and "for" and the phrase "by virtue of the fact that" are often used in the same way. There is a variation on this mode of presenting an argument, where the word "since" or "because" is followed by a list of premises and then the conclusion; for example, "Since all men are basically selfish and Socrates is a man, Socrates is basically selfish."

These are the most common ways of stating arguments in English, but there are other ways, too numerous to catalog. However, you should have no trouble identifying the premises and conclusion of a given argument if you remember that:

> The conclusion states the point being argued for and the premises state the reasons being advanced in support of the conclusion.

Since in logic we are interested in clarity rather than in literary style, one simple, clear method of stating an argument (and indicating which statements are the premises and which the conclusion) is preferred to the rich variety of forms available in English. We will put an argument into standard logical form simply by listing the premises, drawing a line under them, and writing the conclusion under the line. For example, the argument "Diodorus was not an Eagle Scout, since Diodorus did not know how to tie a square knot and all Eagle Scouts know how to tie square knots" would be put into standard logical form as follows:

Diodorus did not know how to tie a square knot.
All Eagle Scouts know how to tie square knots.

Diodorus was not an Eagle Scout.

Exercises

1. Which of the following sentences are statements?
   a. Friends, Romans, countrymen, lend me your ears.
   b. The sum of the squares of the sides of a right triangle equals the square of the hypotenuse.
   c. Hast thou considered my servant Job, a perfect and an upright man?
   d. My name is Faust; in all things thy equal.
   e. $E = mc^2$.
   f. May he be boiled in his own pudding and buried with a stick of holly through his heart.
   g. Ptolemy maintained that the sun revolved around the Earth.
   h. Ouch!
   i. Did Sophocles write *Electra*?
   j. The sun never sets on the British Empire.
2. Which of the following selections advance arguments? Put all arguments in standard logical form.
   a. All professors are absent-minded, and since Dr. Wise is a professor he must be absent-minded.
   b. Since three o'clock this afternoon I have felt ill, and now I feel worse.
   c. Candidate X is certain to win the election because his backers have more money than do Candidate Y's, and furthermore Candidate X is more popular in the urban areas.
   d. Iron will not float when put in water because the specific gravity of water is less than that of iron.
   e. In the past, every instance of smoke has been accompanied by fire, so the next instance of smoke will also be accompanied by fire.

II.3. LOGIC. When we *evaluate* an argument, we are interested in two things:

i. Are the premises true?

ii. Supposing that the premises are true, what sort of support do they give to the conclusion?

The first consideration is obviously of great importance. The argument "All college students are highly intelligent, since all college students are insane, and all people who are insane are highly intelligent" is not very compelling, simply because it is a matter of common knowledge that the premises are false. But important as consideration (i) may be, it is not the business of a logician to judge whether the premises of an argument are true or false.[1] After all, any statements whatsoever can stand as premises to some argument, and the logician has no special access to all human knowledge. If the premises of an argument make a claim about the internal structure of the nucleus of the carbon atom, one is likely to get more reliable judgments as to their truth from a physicist than from a logician. If the premises claim that a certain mechanism is responsible for the chameleon's color changes, one would ask a biologist, not a logician, whether they are true.

Consideration (ii), however, is the logician's stock in trade. Supposing that the premises are true, does it follow that the conclusion must be true? Do the premises provide strong but not conclusive evidence for the conclusion? Do they provide any evidence at all for it? These are questions which it is the business of logic to answer.

> **Definition 2:** *Logic* is the study of the strength of the evidential link between the premises and conclusions of arguments.

In some arguments the link between the premises and the conclusion is the strongest possible in that the truth of the premises *guarantees* the truth of the conclusion. Consider the following argument: "No Athenian ever drank to excess, and Alcibiades was an Athenian. Therefore, Alcibiades never drank to excess." Now if we suppose that the premises "No Athenian ever drank to excess" and "Alcibiades was an Athenian" are true, then we must also suppose that the conclusion "Alcibiades never drank to excess" is also true, for there is no way in which the conclusion could be false while the premises were true. Thus, for this argument we say that the truth of the premises guarantees the truth of the conclusion, and the evidential link between premises and conclusion is as strong as possible. This is in no way altered by the fact that the first premise and the conclusion are false. What is important for evaluating the strength of the evidential link is that, *if* the premises were true, the conclusion would also have to be true.

In other arguments the link between the premises and the conclusion is not so strong, but the premises nevertheless provide some evidential support for the conclusion. Sometimes the premises provide strong evidence for the conclusion, sometimes weak evidence. In the following argument the truth of

[1] Except in certain very special cases which need not concern us here.

the premises does not guarantee the truth of the conclusion, but the evidential link between the premises and the conclusion is still quite strong:

> Smith has confessed to killing Jones. Dr. Zed has signed a statement to the effect that he saw Smith shoot Jones. A large number of witnesses heard Jones gasp with his dying breath, "Smith did it." Therefore Smith killed Jones.

Although the premises are good evidence for the conclusion, we know that the truth of the premises does not *guarantee* the truth of the conclusion, for *we can imagine circumstances in which the premises would be true and the conclusion false.*

> Suppose, for instance, that Smith was insane and that he confessed to every murder he ever heard of, but that this fact was generally unknown because he had just moved into the neighborhood. This peculiarity was, however, known to Dr. Zed, who was Jones's psychiatrist. For his own malevolent reasons, Dr. Zed decided to eliminate Jones and frame Smith. He convinced Jones under hypnosis that Smith was a homocidal maniac bent on killing Jones. Then one day Dr. Zed shot Jones from behind a potted plant and fled.

Let it be granted that these circumstances are highly improbable. If they were not, the premises could not provide strong evidential support for the conclusion. Nevertheless, the circumstances are not impossible and thus the truth of the premises does not guarantee the truth of the conclusion.

The following is an argument in which the premises provide some evidence for the conclusion, but in which the evidential link between the premises and the conclusion is much weaker than in the foregoing example:

> Student 777 arrived at the health center to obtain a medical excuse from his final examination. He complained of nausea and a headache. The nurse reported a temperature of 100 degrees. Therefore, student 777 was really ill.

Given that the premises of this argument are true, it is not as improbable that the conclusion is false as it was in the preceding argument. Hence, the argument is a weaker one, though not entirely without merit.

Thus we see that arguments may have various *degrees of strength.* When the premises present absolutely conclusive evidence for the conclusion—that is, when the truth of the premises guarantees the truth of the conclusion—then we have the strongest possible type of argument. There are cases ranging from this maximum possible strength down to arguments where the premises contradict the conclusion.

Exercises:

Arrange the following arguments in the order of the strength of the link between premises and conclusion.

1. No one who is not a member of the club will be admitted to the meeting.
   I am not a member of the club.
   ____
   I will not be admitted to the meeting.
2. The last three cars I have owned have all been sports cars. They have all performed beautifully and given me little trouble. Therefore, I am sure that the next sports car I own will also perform beautifully and give me little trouble.
3. My nose itches; therefore I am going to have a fight.
4. Brutus said that Caesar was ambitious, and Brutus was an honorable man. Therefore Caesar must have been ambitious.
5. The weatherman has said that a low-pressure front is moving into the area. The sky is gray and very overcast. On the street I can see several people carrying umbrellas. The weatherman is usually accurate. Therefore, it will rain.

II.4. INDUCTIVE VERSUS DEDUCTIVE LOGIC. When an argument is such that the truth of the premises guarantees the truth of the conclusion, we shall say that it is deductively valid. When an argument is not deductively valid but nevertheless the premises provide good evidence for the conclusion, the argument is said to be inductively strong. How strong it is depends on how much evidential support the premises give to the conclusion. In line with the discussion in the last section, we can define these two concepts more precisely as follows:

> **Definition 3:** An argument is *deductively valid* if and only if it is *impossible* that its conclusion is false while its premises are true.
>
> **Definition 4:** An argument is *inductively strong* if and only if it is *improbable* that its conclusion is false while its premises are true, and it is not deductively valid. The *degree* of inductive strength depends on how improbable it is that the conclusion is false while the premises are true.[2]

The sense of "impossible" intended in Definition 3 requires clarification. In a sense, it is impossible for me to fly around the room by flapping my arms;

[2]Although the "while" in Definition 3 may be read as "and" with the definition remaining correct, the "while" in Definition 4 should be read as "given that" and not "and." The reasons for this can be made precise only after some probability theory has been studied. However, the sense of Definition 4 will be explained later in this section.

this sense of impossibility is called *physical impossibility*. But it is not physical impossibility that we have in mind in Definition 3. Consider the following argument:

George is a man.
George is 100 years old.
George has arthritis.

George will not run a four-minute mile tomorrow.

Although it is physically impossible for the conclusion of the argument to be false (that is, that he will indeed run a four-minute mile) while the premises are true, the argument, although a pretty good one, is *not* deductively valid.

To uncover the sense of impossibility in the definition of deductive validity, let us look at an example of a deductively valid argument:

No gourmets enjoy banana–tuna fish soufflés with chocolate sauce.
Antoine enjoys banana–tuna fish soufflés with chocolate sauce.

Antoine is not a gourmet.

In this example it is impossible in a stronger sense—we shall say *logically impossible*—for the conclusion to be false while the premises are true. What sort of impossibility is this? For the conclusion to be false Antoine would have to be a gourmet. For the second premise to be true he would also have to enjoy banana–tuna fish soufflés with chocolate sauce. But for the first premise to be true there must be no such person. Thus, to suppose the conclusion is false is to contradict the factual claim made by the premises. To put the matter a different way, the factual claim made by the conclusion is already implicit in the premises. This is a feature of all deductively valid arguments.

> If an argument is deductively valid, its conclusion makes no factual claim that is not, at least implicitly, made by its premises.

Thus, it is logically impossible for the conclusion of a deductively valid argument to be false while its premises are true, because to suppose that the conclusion is false is to contradict some of the factual claims made by the premises.

We can now see why the following argument is not deductively valid:

George is a man.
George is 100 years old.
George has arthritis.

George will not run a four-minute mile tomorrow.

The factual claim made by the conclusion is *not* implicit in the premises, for there is no premise stating that no 100-year-old man with arthritis can run a four-minute mile. Of course, we all believe this to be a fact, but there is nothing in the premises that claims this to be a fact; if we *added* a premise to this effect, *then* we would have a deductively valid argument.

The conclusion of an *inductively strong argument*, on the other hand, ventures beyond the factual claims made by the premises. The conclusion asserts more than the premises, since we can describe situations in which the premises would be true and the conclusion false.

> If an argument is inductively strong, its conclusion makes factual claims that go beyond the factual information given in the premises.

Thus, an inductively strong argument risks more than a deductively valid one; it risks the possibility of leading from true premises to a false conclusion. But this risk is the price that must be paid for the advantage which inductively strong arguments have over deductively valid ones: the possibility of discovery and prediction of new facts on the basis of old ones.

Definition 4 stated that an argument is inductively strong if and only if it meets two conditions:

i. It is improbable that its conclusion is false, while its premises are true.

ii. It is not deductively valid.

Condition (ii) is required because all deductively valid arguments meet condition (i). It is *impossible* for the conclusion of a deductively valid argument to be false while its premises are true, so the probability that the conclusion is false while the premises are true is zero.

Condition (i), however, requires clarification. The "while" in this condition should be read as "given that," not as "and," so that the condition can be rephrased as:

i. It is improbable that its conclusion is false, *given that* its premises are true.

But just what do we mean by "given that?" And why is "It is improbable that its conclusion is false *and* its premises true" an incorrect formulation of condition (i)? What is the difference, in this context, between "and" and "given that"? At this stage these questions are best answered by examining several examples of arguments. The following is an inductively strong argument:

There is intelligent life on Mercury.
There is intelligent life on Venus.

There is intelligent life on Earth.
There is intelligent life on Jupiter.
There is intelligent life on Saturn.
There is intelligent life on Uranus.
There is intelligent life on Neptune.
There is intelligent life on Pluto.

There is intelligent life on Mars.

Note that the conclusion is not by itself probable. It is, in fact, probable that the conclusion is false. But it is improbable that the conclusion is false *given that* the premises are true. That is, if the premises were true, then on the basis of that information it would be probable that the conclusion would be true (and thus improbable that it would be false). This is not affected in the least by the fact that some of the premises themselves are quite improbable. Thus, although the conclusion taken by itself is improbable, and some of the premises taken by themselves are also improbable, the conclusion is probable *given the premises*. This example illustrates an important principle:

> The type of probability that grades the inductive strength of arguments—we shall call it *inductive probability*—does not depend on the premises alone or on the conclusion alone, but on the *evidential relation* between the premises and the conclusion.

Hopefully we have now gained a certain intuitive understanding of the phrase "given that." Let us now see why it is incorrect to replace it with "and" and thus incorrect to say that an argument is inductively strong if and only if it is improbable that its conclusion is false *and* its premises are true (and it is not deductively valid). Consider the following argument, which is not inductively strong:

There is a 2000-year-old man in Cleveland.

There is a 2000-year-old man in Cleveland who has three heads.

Now it is quite probable that the conclusion is false *given that* the premise is true. Given that there is a 2000-year-old man in Cleveland, it is quite likely that he has only one head. Thus, the argument is *not* inductively strong. But it is improbable that the conclusion is false *and* the premise is true. For the conclusion to be false and the premise true, there would have to be a non-three-headed 2000-year-old man in Cleveland, and it is quite improbable that there is *any* 2000-year-old man in Cleveland. Thus, it is improbable that the conclusion is false *and* the premise is true, simply because it is improbable that the premise is true.

We now see that the inductive strength of arguments cannot depend on the premises alone. Thus, although it is improbable that the conclusion is false *and* the premises true, it is probable that the conclusion is false *given that* the premises are true and the argument is *not* inductively strong.

An argument might be such that it is improbable that the premises are true *and* the conclusion false, simply because it is improbable that the conclusion is false; that is, it is probable that the conclusion is true. It is important to note that such conditions do not guarantee that the argument is inductively strong. Consider the following example of an argument that has a probable conclusion and yet is *not* inductively strong:

There is a man in Cleveland who is 1999 years and 11-months-old and in good health.

No man will live to be 2000 years old.

Now the conclusion itself is highly probable. Thus, it is improbable that the conclusion is false and consequently improbable that the conclusion is false *and* the premise true. But if the premise were true it would be likely that the conclusion would be false. By itself the conclusion is probable, but given the premise it is not.

The main points of this discussion of inductive strength can be summed up as follows:

1. The inductive probability of an argument is the probability that its conclusion is true given that its premises are true.
2. The inductive probability of an argument is determined by the evidential relation between its premises and its conclusion, not by the likelihood of the truth of its premises alone or the likelihood of the truth of its conclusion alone.
3. An argument is inductively strong if and only if:
   a. Its inductive probability is high.
   b. It is not deductively valid.

We defined logic as the study of the strength of the evidential link between the premises and conclusions of arguments. We have seen that there are two different standards against which to evaluate the strength of this link: deductive validity and inductive strength. Corresponding to these two standards are two branches of logic: deductive logic and inductive logic. *Deductive logic* is concerned with tests for deductive validity—that is, rules for deciding whether or not a given argument is deductively valid—and rules for constructing deductively valid arguments. *Inductive logic* is concerned with tests for measuring the inductive probability, and hence the inductive strength, of arguments and with rules for constructing inductively strong arguments.

Some books appear to suggest that there are two different types of arguments, deductive and inductive, and that deductive logic is concerned with deductive arguments and inductive logic with inductive arguments. That is, they suggest the following classification, together with the assumption that every argument falls in one and only one category:

	Deductive arguments	Inductive arguments
Good	Valid	Strong
Bad	Invalid	Weak

Nothing, however, is further from the truth, for, as we have seen, all inductively strong arguments are deductively invalid.

It is more correct to picture arguments as being arranged on a scale of descending strength, as follows:

Arguments

Deductively valid

Degrees of inductive strength

Worthless

Deductive and inductive logic are not distinguished by the different types of arguments with which they deal, but by the different standards against which they evaluate arguments.

Exercises:

Decide whether each of the following arguments is deductively valid, inductively strong, or neither:

1. The meeting place is either the gym or the cafeteria.
   The meeting place is not the gym.
   ______
   The meeting place is the cafeteria.

2. A good meal has always made me feel better.
   ______
   A good meal today will make me feel better.
3. Many great leaders have been crazy.
   ______
   Everyone who isn't a leader is sane.
4. On all the birthdays I have ever had I have been less than 30 years old.
   ______
   On my next birthday I will be less than 30 years old.
5. No pigs can fly.
   Some horses can fly.
   ______
   Some horses aren't pigs.

II.5. EPISTEMIC PROBABILITY. We have seen that the concept of inductive probability applies to arguments. The inductive probability of an argument is the probability that its conclusion is true given that its premises are true. Thus, the inductive probability of an argument is a measure of the strength of the evidence that the premises provide for the conclusion. It is correct to speak of the inductive probability of an argument, but incorrect to speak of the inductive probability of statements. Since the premises and conclusion of any argument are statements, it is incorrect to speak of the inductive probability of a premise or of a conclusion.

There is, however, some sense of probability in which it is intuitively acceptable to speak of the probability of a premise or conclusion. When we said that it is improbable that there is a 2000-year-old man in Cleveland, we were relying on some such intuitive sense of probability. There must then be a type of probability, other than inductive probability, that applies to statements rather than arguments.

Let us call this type of probability *epistemic probability* because the Greek stem *episteme* means knowledge, and the epistemic probability of a statement depends on just what our stock of relevant knowledge is. Thus, *the epistemic probability of a statement can vary from person to person and from time to time*, since different people have different stocks of knowledge at the same time and the same person has different stocks of knowledge at different times. *For me*, the epistemic probability that there is a 2000-year-old man now living in Cleveland is quite low, since I have certain background knowledge about the current normal life span of human beings. I feel safe in using this statement as an example of a statement whose epistemic probability is low because I feel safe in assuming that your stock of background knowledge is similar in the relevant respects and that *for you* its epistemic probability is also low.

It is easy to imagine a situation in which the background knowledge of two people would differ in such a way as to generate a difference in the epistemic probability of a given statement. For example, the epistemic probability that Pegasus will show in the third race may be different for a fan in the grandstand than for Pegasus' jockey, owing to the difference in their knowledge of the relevant factors involved.

It is also easy to see how the epistemic probability of a given statement can change over time for a particular person. The fund of knowledge that each of us possesses is constantly in a state of flux. We are all constantly learning new things directly through experience and indirectly through information which is communicated to us. We are also, unfortunately, continually forgetting things that we once knew. This holds true for societies and cultures as well as for individuals, and human knowledge is continually in a dynamic process of simultaneous growth and decay.

It is important to see how upon the addition of new knowledge to a previous body of knowledge the epistemic probability of a given statement could either increase or decrease. Suppose we are interested in the epistemic probability of the statement that Mr. X is an Armenian and the only relevant information we have is that Mr. X is an Oriental rug dealer in Allentown, Pa., that 90 percent of the Oriental rug dealers in the United States are Armenian, and that Allentown, Pa., is in the United States. On the basis of this stock of relevant knowledge, the epistemic probability of the statement is equal to the inductive probability of the following argument:

> Mr. X is an Oriental rug dealer in Allentown, Pa.
> Allentown, Pa., is in the United States.
> Ninety percent of the Oriental rug dealers in the United States are Armenian.
> ___
> Mr. X is an Armenian.

The inductive probability of this argument is quite high. If we are now given the new information that although 90 percent of the Oriental rug dealers in the United States are Armenian, only 2 percent of the Oriental rug dealers in Allentown, Pa., are Armenian, while 98 percent are Syrian, the epistemic probability that Mr. X is Armenian decreases drastically, for it is now equal to the inductive probability of the following argument:

> Mr. X is an Oriental rug dealer in Allentown, Pa.
> Allentown, Pa., is in the United States.
> Ninety percent of the Oriential rug dealers in the United States are Armenian.
> Ninety-eight percent of the Oriental rug dealers in Allentown, Pa., are Syrian.
> Two percent of the Oriental rug dealers in Allentown, Pa., are Armenian.
> ___
> Mr. X is an Armenian.

The inductive probability of this argument is quite low. Note that the decrease in the epistemic probability of the statement "Mr. X is an Armenian" results not from a change in the inductive probability of a given argument but from the fact that, upon the addition of new information, a *different* argument with more premises becomes relevant in assessing its epistemic probability.

Suppose now we are given still more information, to the effect that Mr. X is a member of the Armenian Club of Allentown and that 99 percent of the members of the Armenian Club are actually Armenians. Upon addition of this information the epistemic probability that Mr. X is an Armenian again becomes quite high, for it is now equal to the inductive probability of the following argument:

> Mr. X is an Oriental rug dealer in Allentown, Pa.
> Allentown, Pa., is in the United States.
> Ninety percent of the Oriental rug dealers in the United States are Armenian.
> Ninety-eight percent of the Oriental rug dealers in Allentown, Pa., are Syrian.
> Two percent of the Oriental rug dealers in Allentown, Pa., are Armenian.
> Mr. X is a member of the Armenian Club of Allentown, Pa.
> Ninety-nine percent of the members of the Armenian Club are Armenian.
> ___
> Mr. X is an Armenian.

Notice once more that the epistemic probability of the statement changes because, with the addition of new knowledge, it became equal to the inductive probability of a new argument with additional premises.

Epistemic probabilities are *important* to us. They are the probabilities upon which we base our decisions. From a stock of knowledge we will arrive at the associated epistemic probability of a statement by the application of inductive logic. Exactly how inductive logic gets us epistemic probabilities from a stock of knowledge depends on how we characterize a stock of knowledge. Just what knowledge is; how we get it; what it is like once we have it; these are deep questions. At this stage, we will work within a simplified model of knowing—the Certainty Model.

The Certainty Model: Suppose that our knowledge originates in observation; that observation makes particular sentences (observation reports) certain and that the probability of other sentences is attributable to the certainty of these. In such a situation we can identify our stock of knowledge with a *list* *of sentences*, those observation reports that have been rendered certain by observational experience. It is then natural to evaluate the probability of a statement by looking at an argument with all our stock of knowledge as premises and the statement in question as the conclusion. The inductive strength of that argument will determine the probability of the statement in question. In the certainty model, the relation between epistemic probability and inductive probability is quite simple:

> **Definition 5:** The *epistemic* probability of a statement is the *inductive* probability of that argument which has the statement in question as its conclusion and whose premises consist of all of the observation reports which comprise our stock of knowledge.

The Certainty Model lives up to its name by assigning epistemic probability of one to each observation report in our stock of knowledge.

The certainty of observation reports may be something of an idealization. But it is a useful idealization, and we will adopt it for the present. Later in the course we will discuss some other models of observation.

Exercise

1. Construct several new examples in which the epistemic probability of a statement is increased or decreased by the addition of new information to a previous stock of knowledge.

II.6. PROBABILITY AND THE PROBLEMS OF INDUCTIVE LOGIC. Deductive logic, at least in its basic branches, is well-developed. The definitions of its basic concepts are precise, its rules are rigorously formulated, and the interrelations between the two are well understood. Such is not the case, however, with inductive logic. There are no universally accepted rules for constructing inductively strong arguments; no general agreement on a way of measuring the inductive strength of arguments; no precise, uncontroversial definition of inductive probability. Thus, inductive logic cannot be learned in the sense in which one learns algebra or the basic branches of deductive logic. This is not to say that inductive logicians are wallowing in a sea of total ignorance; many things are known about inductive logic, but many problems still remain to be solved. We shall try to get an idea of just what the problems are, as well as what progress has been made toward their solution.

Some of the main problems of inductive logic can be framed in terms of the concept of inductive probability. I said that there is no precise, uncontroversial definition of inductive probability. I did give a definition of inductive probability. Was it controversial? I think not, but, if you will remember, it was imprecise. I said that the inductive probability of an argument is the probability that its conclusion is true, given that its premises are true. But at that point I could not give an exact definition of "the probability that an argument's conclusion is true, given that its premises are true." I was, instead, reduced to giving examples so that you could get an intuitive feeling for the meaning of this phrase. The logician, however, is not satisfied with an intuitive feeling for the meaning of key words and phrases. He wishes to analyze the concepts involved and arrive at precise, unambiguous definitions. Thus, one of the problems of inductive logic which remains outstanding is, what exactly is inductive probability?

This problem is intimately connected with two other problems: How is the inductive probability of an argument measured? And, what are the rules for constructing inductively strong arguments? Obviously we cannot develop an exact measure of inductive probability if we do not know precisely what it is. And before we can devise rules for constructing inductively strong arguments, we must have ways of telling which arguments measure up to the required degree of inductive strength. Thus, the solution to the problem of providing a precise definition of inductive probability determines what solutions are available for the problems of determining the inductive probabilities of arguments and constructing systematic rules for generating inductively strong arguments.

Let us call a precise definition of inductive probability, together with the associated method of determining the inductive probability of arguments and rules for constructing inductively strong arguments, an *inductive logic*. Thus, different definitions of inductive probability give rise to different inductive logics. Now we are not interested in finding just any system of inductive logic. We want a system that accords well with common sense and scientific practice. We want a system that gives the result that most of the cases that we would intuitively classify as inductively strong arguments do indeed have a high inductive probability. We want a system that accords with scientific practice and common sense, but that is more precise, more clearly formulated, and more rigorous than they are; a system that codifies, explains, and refines our intuitive judgments. We shall call such a system of inductive logic a *scientific inductive logic*. The problem that we have been discussing can now be reformulated as *the problem of constructing a scientific inductive logic*.

The second major problem of inductive logic, and the one that has been more widely discussed in the history of philosophy, is *the problem of rationally justifying the use of a system of scientific inductive logic* rather than some other system of inductive logic. After all, there are many different possible inductive logics. Some might give the result that arguments that we think are inductively strong are, in fact, inductively weak, and arguments that we think inductively weak are, in fact, inductively strong. That is, there are possible inductive logics which are diametrically opposed to scientific inductive logic, which are in total disagreement with scientific practice and common sense. Why should we not employ one of these systems rather than scientific induction?

Any adequate answer to this question must take into account the uses to which we put inductive logic (or, at present, the vague intuitions we use in place of a precise system of inductive logic). One of the most important uses of inductive logic is to frame our expectations of the future on the basis of our knowledge of the past and present. We must use our knowledge of the past and present as a guide to our expectations of the future; it is the only guide we have. But it is impossible to have a *deductively valid* argument whose premises contain factual information solely about the past and present and whose conclusion makes factual claims about the future. For the conclusion of a deductively valid argument makes no factual claim that is not already made by the premises. Thus, the gap separating the past and present from the future cannot be bridged in this way by deductively valid arguments, and if the arguments we use to bridge that gap are to have any strength whatsoever they must be inductively strong.

Let us look a little more closely, then, at the way in which inductive logic would be used to frame our expectations of the future. Suppose our plans depend critically on whether it will rain tomorrow. Then the reasonable thing to do, before we decide what course of action to take, is to ascertain the epistemic probability of the statement. "It will rain tomorrow." This we do by putting all the relevant information we now have into the premises of an argument whose conclusion is "It will rain tomorrow" and ascertaining the inductive probability of that argument. If the probability is high, we will have a strong expectation of rain and will make our plans on that basis. If the probability is near zero, we will be reasonably sure that it will not rain and act accordingly.

Now although it is doubtful that anyone carries out the formal process outlined above when he plans for the future, it is hard to deny that, if we were to make our reasoning explicit, it would fall into this pattern. Thus, the making of rational decisions is dependent, via the concept of epistemic probability, on our inductive logic. The second main problem of inductive logic, then, leads us to the following question: How can we rationally justify the use of scientific inductive logic, rather than some other inductive logic, as an instrument for shaping our expectations of the future?

The two main problems of inductive logic are:

1. The construction of a system of scientific inductive logic.
2. The rational justification of the use of that system rather than some other system of inductive logic.

It would seem that the first problem must be solved before the second, since we can hardly justify the use of a system of inductive logic before we know what it is. Nevertheless, I shall discuss the second problem first. It makes sense to do this because we can see why the second problem is *such* a problem without having to know all the details of scientific inductive logic. Furthermore, philosophers historically came to appreciate the difficulty of the second problem much earlier than they realized the full force of the first problem. This second problem, the traditional problem of induction, is discussed in the next chapter.

III

The Traditional Problem of Induction

III.1. INTRODUCTION. In Chapter II we saw that inductive logic is used to shape our expectations of that which is as yet unknown on the basis of those facts that are already known; for instance, to shape our expectations of the future on the basis of our knowledge of the past and present. Our problem is the rational justification of the use of a system of scientific inductive logic, rather than some other system of inductive logic, for this task.

The Scottish philosopher David Hume first raised this problem, which we shall call the *traditional problem of induction*, in full force. Hume gave the problem a cutting edge, for he advanced arguments designed to show that no such rational justification of inductive logic is possible, no matter what the details of a system of scientific inductive logic turn out to be. The history of philosophical discussion of inductive logic since Hume has been in large measure occupied with attempts to circumvent the difficulties he raised. This chapter examines these difficulties and the various attempts to overcome them.

II.2. HUME'S ARGUMENT. Before we can meaningfully discuss arguments which purport to show that it is impossible to rationally justify scientific induction, we must be clear on what would be required to rationally justify a system of inductive logic. Presumably we could rationally justify such a system if we could show that it is well-suited for the uses to which it is put. One of the most important uses of inductive logic is in setting up our predictions of the future. Inductive logic figures in these predictions by way of *epistemic probabilities*. If a claim about the future has high epistemic probability, we predict that it will prove true. And, more generally, we expect something more or less strongly as its epistemic probability is higher or lower. The epistemic probability of a statement is just the inductive probability of the argument which embodies all available information in its premises. Thus, the epistemic probability of a statement depends on two things: (i) the stock of knowledge, and (ii) the inductive logic used to grade the strength of the argument from that stock of knowledge to the conclusion.

Now obviously what we want is for our predictions to be correct. If we could get by with deductively valid arguments we could be assured of true predictions all the time. Deductively valid arguments lead from true premises always to

true conclusions and the statements comprising our stock of knowledge are known to be true. But deductively valid arguments are too conservative to leap from the past and present to the future. For this sort of daring behavior we will have to rely on inductively strong arguments—and we will have to give up the comfortable assurance that we will be right all the time.

How about most of the time? Let us call the sort of argument used to set up an epistemic probability an e-argument. That is, an *e-argument* is an argument which has, as its premises, some stock of knowledge. We might hope, then, that inductively strong e-arguments will give us *true conclusions most of the time*. Remember that there are *degrees* of inductive strength and that, on the basis of our present knowledge, we do not always simply predict or not-predict that an event will occur, but anticipate it with various *degrees of confidence*. We might hope further that inductively *stronger* e-arguments have true conclusions *more often* than inductively *weaker* ones. Finally, since we think that it is useful to gather evidence to enlarge our stock of knowledge, we might hope that inductively strong e-arguments give us true conclusions more often when the stock of knowledge embodied in the premises is great than when it is small.

The last consideration really has to do with justifying epistemic probabilities as tools for prediction. The epistemic probability is the inductive probability of an argument embodying *all* our stock of knowledge in its premises. The requirement that it embody *all* our knowledge, and not just some part of it, is known as the Total Evidence Condition. If we could show that basing our predictions on more knowledge gives us better success ratios, we would have justified the total evidence condition.

The other considerations have to do with justifying the other determinant of epistemic probability—the inductive logic which assigns inductive probabilities to arguments.

We are now ready to suggest what is required to rationally justify a system of inductive logic:

Rational Justification

> *Suggestion I:* A system of inductive logic is rationally justified if and only if it is shown that the arguments to which it assigns high inductive probability yield true conclusions from true premises most of the time, and the e-arguments to which it assigns higher inductive probability yield true conclusions from true premises more often than the arguments to which it assigns lower inductive probability.

It is this sense of rational justification, or something quite close to it, that Hume has in mind when he advances his arguments to prove that a rational justification of scientific induction is impossible.

If scientific induction is to be rationally justified in the sense of Suggestion I, we must establish that the arguments to which it assigns high inductive probability yield true conclusions from true premises most of the time. By what sort of reasoning, asks Hume, could we establish such a conclusion? If the argument that we must use is to have any force whatsoever, it must be either deductively valid or inductively strong. Hume proceeds to show that neither sort of argument could do the job.

Suppose we try to rationally justify scientific inductive logic by means of a deductively valid argument. The only premises we are entitled to use in this argument are those that state things we know. Since we do not know what the future will be like (if we did, we would have no need of an inductive logic on which to base our predictions), the premises can contain knowledge of only the past and present. But if the argument is deductively valid, then the conclusion can make no factual claims that are not already made by the premises. Thus, the conclusion of the argument can only refer to the past and present, not to the future, for the premises made no factual claims about the future. Such a conclusion cannot, however, be adequate to rationally justify scientific induction.

To rationally justify scientific induction we must show that e-arguments to which it assigns high inductive probability yield true conclusions from true premises most of the time. And "most of the time" does not mean most of the time in only the past and present; it means most of the time, *past*, *present*, and *future*. It is conceivable that a certain type of argument might have given us true conclusions from true premises in the past and might cease to do so in the future. Since our conclusion cannot tell us how successful arguments will be in the future, it cannot establish that the e-arguments to which scientific induction assigns high probability will give us true conclusions from true premises *most of the time*. Thus, we cannot use a deductively valid argument to rationally justify induction.

Suppose we try to rationally justify scientific induction by means of an inductively strong argument. We construct our argument, whatever it may be, and present it as an inductively strong argument. "Why do you think that this is an inductively strong argument?" Hume might ask. "Because it has a high inductive probability," we would reply. "And what system of inductive logic assigns it a high probability?" "Scientific induction, of course." What Hume has pointed out is that if we attempt to rationally justify scientific induction by use of an inductively strong argument, we are in the position of having to *assume* that scientific induction is reliable in order to prove that scientific induction is reliable; we are reduced to begging the question. Thus, we cannot use an inductively strong argument to rationally justify scientific induction.

A common argument is that scientific induction is justified because it has been quite successful in the past. On reflection, however, we see that this argument is really an attempt to justify induction by means of an inductively strong argument, and thus begs the question. More explicitly, the argument reads something like this:

> Arguments that are judged by scientific inductive logic to have high inductive probability have given us true conclusions from true premises most of the time in the past.
>
> ---
>
> Such arguments will give us true conclusions from true premises most of the time, past, present, and future.

It should be obvious that this argument is not deductively valid. At best it is assigned high inductive probability by a system of scientific inductive logic. But the point at issue is whether we should put our faith in such a system.

We can view the traditional problem of induction from a different perspective by discussing it in terms of the *principle of the uniformity of nature*. Although we do not have the details of a system of scientific induction in hand, we do know that it must accord well with common sense and scientific practice, and we are reasonably familiar with both. A few examples will illustrate a general principle which appears to underlie both scientific and common-sense judgments of inductive strength.

If you were to order steak in a restaurant, and a friend were to object that steak would corrode your stomach and lead to quick and violent death, it would seem quite sufficient to respond that you had often eaten steak without any of the dire consequences he predicted. That is, you would intuitively judge the following argument to be inductively strong:

> I have eaten steak many times and it has never corroded my stomach.
>
> ---
>
> Steak will not now corrode my stomach.

Suppose a scientist is asked whether a rocket would work in reaches of space beyond the range of our telescopes. She replies that it would, and to back up her answer appeals to certain principles of theoretical physics. When asked what evidence she has for these principles, she can refer to a great mass of observed phenomena that corroborate them. The scientist is then judging the following argument to be inductively strong:

> Principles A, B, and C correctly describe the behavior of material bodies in all of the many situations we have observed.
>
> ---
>
> Principles A, B, and C correctly describe the behavior of material bodies in those reaches of space that we have not as yet observed.

There appears to be a common assumption underlying the judgments that these arguments are inductively strong. As a steak eater you assume that the future will be like the past, that types of food that proved healthful in the past will continue to prove so in the future. The scientist assumes that the distant reaches of space are like the nearer ones, that material bodies obey the same general laws in all areas of space. Thus, it seems that underlying our judgments of inductive strength in both common sense and science is the presupposition that nature is uniform or, as it is sometimes put, that like causes produce like effects throughout all regions of space and time. Thus, we can say that a system of scientific induction will base its judgments of inductive strength on the presupposition that *nature is uniform* (and in particular that the future will resemble the past).

We ought to realize at this point that we have only a vague, intuitive understanding of the principle of the uniformity of nature, gleaned from examples rather than specified by precise definitions. This rough understanding is sufficient for the purposes at hand. But we should bear in mind that the task of giving an *exact* definition of the principle, a definition of the sort that would be presupposed by a system of scientific inductive logic, is as difficult as the construction of such a system itself. One of the problems is that nature is simply not uniform in all respects, the future does not resemble the past in all respects. Bertrand Russell once speculated that the chicken on slaughter-day might reason that whenever the humans came it had been fed, so when the humans would come today it would also be fed. The chicken thought that the future would resemble the past, but it was dead wrong.

The future may resemble the past, but it does not do so in all respects. And we do not know beforehand what those respects are nor to what degree the future resembles the past. Our ignorance of what these respects are is a deep reason behind the total evidence condition. Looking at more and more evidence helps us reject spurious patterns which we might otherwise project into the future. Trying to say exactly *what* about nature we believe is uniform thus turns out to be a surprisingly delicate task.

But suppose that a subtle and sophisticated version of the principle of the uniformity of nature can be formulated which adequately explains the judgments of inductive strength rendered by scientific inductive logic. Then if nature is indeed uniform in the required sense (past, present, and future), e-arguments judged strong by scientific induction will indeed give us true conclusions most of the time. Therefore, the problem of rationally justifying scientific induction could be reduced to the problem of establishing that nature is uniform.

But by what reasoning could we establish such a conclusion? If an argument is to have any force whatsoever it must be either deductively valid or inductively strong. A deductively valid argument could not be adequate, for if the information in the premises consists solely of our knowledge of the past and present, then the conclusion cannot tell us that nature will be uniform in the future. The conclusion of a deductively valid argument can make no factual claims that are not already made by the premises, and factual claims about the future are not factual claims about the past and present. But if we claim to have established the principle of the uniformity of nature by an argument that is rated inductively strong by scientific inductive logic, we are open to a challenge as to why we should place our faith in such arguments. But we cannot reply "Because nature is uniform," for that is precisely what we are trying to establish.

Let us summarize the traditional problem of induction. It appears that to rationally justify a system of scientific inductive logic we would have to establish that the e-arguments it judges to be inductively strong give us true conclusions most of the time. If we try to prove that this is the case by means of a deductively valid argument whose premises state things we already know, then the conclusion must fall short of the desired goal. But to try to rationally justify scientific induction by means of an argument that scientific induction judges to be inductively strong is to beg the question. The same difficulties arise if we attempt to justify scientific inductive logic by establishing that nature is uniform.

III.3. THE INDUCTIVE JUSTIFICATION OF INDUCTION.

Hume has presented us with a dilemma. If we try to justify scientific inductive logic by means of a deductively valid argument with premises known to be true, our conclusion will be too weak. If we try to use an inductively strong argument, we are reduced to begging the question. The proponent of the inductive justification of induction tackles the second horn of the dilemma. He maintains that we can justify scientific induction by an inductively strong argument without begging the question. Although his attempt is not altogether successful, there is a great deal to be learned from it.

The answer to the question "Why should we believe that scientific induction is a reliable guide for our expectations?" that immediately occurs to everyone is "Because it has worked well so far." Hume's objection to this answer was that it begs the question, that it assumes scientific induction is reliable in order to prove that scientific induction is reliable. The proponents of the inductive justification of induction, however, claim that the answer only *appears* to beg the question, because of a mistaken conception of scientific induction. They claim that if we properly distinguish *levels* of scientific induction, rather than lumping all arguments that scientific induction judges to be strong in one category, we will see that the inductive justification of induction does not beg the question.

Just what then are these levels of scientific induction? And what is their relevance to the inductive justification of induction? We can distinguish different *levels of argument*, in terms of the things they talk about. Arguments on level 1 will talk about individual things or events; for instance:

> Many jub-jub birds have been observed, and they have all been purple.
> ______
> The next jub-jub bird to be observed will be purple.

Level 1 of scientific inductive logic would consist of rules for assigning inductive probabilities to arguments of level 1. Presumably the rules of level 1 of scientific induction would assign high inductive probability to the preceding argument. Arguments on level 2 will talk about arguments on level 1; for instance:

> Some deductively valid arguments on level 1 have true premises.
> All deductively valid arguments on level 1 which have true premises have true conclusions.
> ______
> Some deductively valid arguments on level 1 have true conclusions.

This is a deductively valid argument on level 2 which talks about deductively valid arguments on level 1. The following is also an argument on level 2 which talks about arguments on level 1:

> Some arguments on level 1 which the rules of level 1 of scientific inductive logic say are inductively strong have true premises.
> The denial of a true statement is a false statement.
> ______
> Some arguments on level 1 which the rules of level 1 of scientific inductive logic say are inductively strong have premises whose denial is false.

This is a deductively valid argument on level 2 that talks about arguments on level 1, which the rules of level 1 of scientific inductive logic classify as inductively strong.

There are, of course, arguments on level 2 that are not deductively valid, and there is a corresponding second level of scientific inductive logic which consists of rules that assign degrees of inductive strength to *these* arguments. There are arguments on level 3 that talk about arguments on level 2, arguments on level 4 that talk about arguments on level 3, and so on. For each level of argument, scientific inductive logic has a corresponding level of rules.

This characterization of the levels of argument, and the corresponding levels of scientific induction, is summarized in Table III.1.

Table III.1

Levels of argument	Levels of scientific inductive logic
k: Arguments about arguments on level $k-1$.	k: Rules for assigning inductive probabilities to arguments on level k.
⋮	⋮
2: Arguments about arguments on level 1.	2: Rules for assigning inductive probabilities to arguments on level 2.
1: Arguments about individuals.	1: Rules for assigning inductive probabilities to arguments on level 1.

As the table shows, scientific inductive logic is seen not as a simple, homogeneous system but rather as a complex structure composed of an infinite number of strata of distinct sets of rules. The sets of rules on different levels are not, however, totally unrelated. The rules on each level presuppose, in some sense, that nature is uniform and that the future will resemble the past. If this were not the case, we would have no reason for calling the whole system of levels a system of *scientific* inductive logic.

We are now in a position to see how the system of levels of scientific induction is to be employed in the inductive justification of induction. In answer to the question, "Why should we place our faith in the rules of level 1 of scientific inductive logic?" the proponent of the inductive justification of induction will advance an argument on level 2:

> Among arguments used to make predictions in the past, e-arguments on level 1 (which according to level 1 of scientific inductive logic are inductively strong) have given true conclusions most of the time.
> ______
> With regard to the next prediction, an e-argument judged inductively strong by the rules of scientific inductive logic will yield a true conclusion.

The proponent will maintain that the premise of this argument is true, and if we ask why he thinks that this is an inductively strong argument, he will reply that *the rules of level 2 of scientific inductive logic* assign it a high inductive probability. If we now ask why we should put our faith in *these* rules, he will advance a similar argument on level 3, justify that argument by appeal to the rules of scientific inductive logic on level 3, justify those rules by an argument on level 4, and so on.

The inductive justification of induction is summarized in Table III.2. The arrows in the table show the order of justification. Thus, the rules of level 1 are justified by an argument on level 2, which is justified by the rules on level 2, which are justified by an argument on level 3, and so on.

Let us now see how it is that the proponent of the inductive justification of induction can plead not guilty to Hume's charge of begging the question; that is, of presupposing exactly what one is trying to prove. In justifying the rules of level 1, the proponent of the inductive justification of induction does not presuppose that *these* rules will work the next time; in fact, he advances an argument (on level 2) to show that they will work next time. Now it is true that the use of this argument presupposes that the rules of level 2 will work next time. But there is another argument waiting on level 3 to show that the rules of level 2 will work. The use of that argument does not presuppose what it is trying to establish; it presupposes that the rules on level 3 will work. Thus, none of the arguments used in the inductive justification of induction presuppose what they are trying to prove, and the inductive justification of induction does not technically beg the question.

Table III.2

Levels of argument		Levels of scientific inductive logic
⋮	↖	⋮
3: Rules of level 2 of scientific inductive logic have worked well in the past. / They will work well next time.	→ ↖	3: Rules for assigning inductive probabilities to arguments on level 3.
2: Rules of level 1 of scientific inductive logic have worked well in the past.* / They will work well next time.	→ ↖	2: Rules for assigning inductive probabilities to arguments on level 2.
1:		1: Rules for assigning inductive probabilities to arguments on level 1.

* The statement "rules of level 1 of scientific inductive logic have worked well in the past," is to be taken as shorthand for "arguments on level 1, which according to the rules of level 1 of scientific inductive logic are inductively strong and which have been used to make predictions in the past, have given us true conclusions, when the premises were true, most of the time." Thus, the argument on level 2 used to justify the rules of level 1 is exactly the same one as put forth in the second paragraph on page 33.

Perhaps how these levels work can be made clearer by looking at a simple example. Suppose our only observations of the world have been of 100 jub-jub birds and they have all been purple. After observing 99 jub-jub birds, we advanced argument jj-99:

> We have seen 99 jub-jub birds and they were all purple.
>
> ---
>
> The next jub-jub bird we see will be purple.

This argument was given high inductive probability by rules of level 1 of scientific inductive logic. We knew its premises to be true, and we took its conclusion as a prediction. The 100th jub-jub bird can thus be correctly described as purple—or as the color that makes the conclusion of argument jj-99 true—or as the color that results in a successful prediction by the rules of level 1 of scientific inductive logic. Let us also suppose that similar arguments had been advanced in the past: jj-98, jj-97, etc. Each of these arguments was an e-argument to which the rules of level 1 assigned high inductive probability. Thus, the observations of jub-jub birds 98 and 99, etc., are also observations of successful outcomes to predictions based on assignments of probabilities to e-arguments by rules of level 1. This gives rise to an argument on level 2:

> e-arguments on level 1, which are assigned high inductive probability by rules of level 1, have had their conclusions predicted 98 times and all those predictions were successful.
>
> ---
>
> Predicting the conclusion of the next e-argument on level 1 which is assigned high inductive probability will also lead to success.

This argument is assigned high inductive probability by rules of level 2. If the next jub-jub bird to be observed is purple, it makes this level 2 argument successful in addition to making the appropriate level 1 argument successful. A string of such successes gives rise to a similar argument on level 3 and so on, up the ladder, as indicated in Table III.2.

If someone were to object that what is wanted is a justification of scientific induction as a whole and that this has not been given, the proponent of the inductive justification of induction would reply that for every level of rules of scientific inductive logic he has a justification (on a higher level), and that certainly if every level of rules is justified, then the whole system is justified. He would maintain that it makes no sense to ask for a justification for the system *over and above* a justification for each of its parts. This position, it must be admitted, has a good deal of plausibility; a final evaluation of its merits, however, must await some further developments.

The position held by the proponent of the inductive justification of induction contrasts with the position held by Hume in that it sets different requirements for the rational justification of a system of inductive logic. The following is implicit in the inductive justification of induction:

Rational Justification

Suggestion II: A system of inductive logic is rationally justified if, for every level (k) of rules of that system, there is an e-argument on the next highest level ($k + 1$) which:

i. Is judged inductively strong by its own system's rules (these will be rules of level $k + 1$).

ii. Has as its conclusion the statement that the system's rules on the original level (k) will work well next time.

It is important to see that *whether a system of induction meets these conditions depends not only on the system of induction itself but also on the facts, on the way that the world is*. We can imagine a situation in which scientific induction would indeed not meet these conditions. Imagine a world which has been so chaotic that scientific induction on level 1 has not worked well; that is, suppose that the e-arguments on level 1, which according to the rules of level 1 of scientific inductive logic are inductively strong and which have been used to make predictions in the past, have given us *false* conclusions from true premises most of the time. In such a situation the inductive justification of induction could not be carried through. For although the argument on level 2 used to justify the rules of level 1 of scientific induction, that is:

> Rules of level 1 of scientific inductive logic have worked well in the past.
>
> ---
>
> They will work well next time.

would still be judged inductively strong by the rules of level 2 of scientific inductive logic, its premise would not be true. Indeed, in the situation under consideration the following argument on level 2 *would* have a premise that was known to be true and would also be judged inductively strong by the rules of level 2 of scientific inductive logic:

> Rules of level 1 of scientific inductive logic have not worked well in the past.
>
> ---
>
> They will not work well next time.

Thus, we can conceive of situations in which level 2 of scientific induction, instead of justifying level 1 of scientific induction, would tell us that level 1 is unreliable.

We are not, in fact, in such a situation. Level 1 of scientific induction has served us quite well, and it is upon this fact that the inductive justification of induction capitalizes. This is indeed an important fact, but it remains to be seen whether it is sufficient to rationally justify a system of scientific inductive logic.

The proponent of the inductive justification of scientific inductive logic has done us a service in distinguishing the various levels of induction. He has also made an important contribution by pointing out that there are possible situations in which the higher levels of scientific induction do not always support the lower levels and that we are, in fact, not in such a situation. But as a justification of the system of scientific induction his reasoning is not totally satisfactory. While he has not technically begged the question, he has come very close to it. Although he has an argument to justify every level of scientific induction, and although none of his arguments presuppose exactly what they are trying to prove, the justification of each level presupposes the correctness of the level above it. Lower levels are justified by higher levels, but always higher levels of scientific induction. No matter how far we go in the justifying process, we are always within the system of scientific induction. Now, isn't this loading the dice? Couldn't someone with a completely different system of inductive logic execute the same maneuver? Couldn't he justify each level of *his* logic by appeal to higher levels of *his* logic? Indeed he could. Given the same factual situation in which the inductive justification of scientific induction is carried out, an entirely different system of inductive logic could also meet the conditions laid down under Rational Justification, Suggestion II. Let us take a closer look at such a contrasting system of inductive logic.

We said that scientific induction assumes that, in some sense, nature is uniform and the future will be like the past. Some such assumption is to be found backing the rules on each level of scientific inductive logic. The assumptions are not exactly the same on each level; they must be different because we can imagine a situation in which scientific induction on level 2 would tell us that scientific induction on level 1 will not work well. Thus, different principles of the uniformity of nature are presupposed on different levels of scientific inductive logic. But although they are not exactly the same, they are similar; they are all principles of the uniformity of nature. Thus, each level of scientific inductive logic presupposes that, in some sense, nature is uniform and the future will be like the past. A system of inductive logic that would be *diametrically opposed* to scientific inductive logic would be one which presupposed on all levels that the future will not be like the past. We shall call this system a system of *counterinductive logic*.

Let us see how counterinductive logic would work on level 1. Scientific inductive logic, which assumes that the future will be like the past, would assign the following argument a high inductive probability:

Many jub-jub birds have been observed and they have all been purple.

The next jub-jub bird to be observed will be purple.

Counterinductive logic, which assumes that the future will *not* be like the past, would assign it a low inductive probability and would instead assign a high inductive probability to the following argument:

Many jub-jub birds have been observed and they have all been purple.

The next jub-jub bird to be observed will not be purple.

In general, counterinductive logic assigns low inductive probabilities to arguments that are assigned high inductive probabilities by scientific inductive logic, and high inductive probabilities to arguments that are assigned low inductive probabilities by scientific inductive logic.

Now suppose that a counterinductivist decided to give an inductive justification of counterinductive logic. The scientific inductivist would justify his rules of level 1 by the following level 2 argument:

Rules of level 1 of scientific induction have worked well in the past.

They will work well next time.

The counterinductivist, on the other hand, would justify his rules of level 1 by another kind of level 2 argument:

Rules of level 1 of counterinductive logic have *not* worked well in the past.

They will work well next time.

By the counterinductivist's rules, this is an inductively strong argument, for on level 2 he also assumes that the future will be unlike the past. Thus, the counterinductivist is not at all bothered by the fact that his level 1 rules have been failures; indeed he takes this as evidence that they will be successful in the future. Granted his argument appears absurd to us, for we are all at heart scientific inductivists. But if the scientific inductivist is allowed to use his own rules on level 2 to justify his rules on level 1, how can we deny the same right to the counterinductivist? If asked to justify his rules on level 2, the counterinductivist will advance a similar argument on level

3, and so on. If an inductive justification of scientific inductive logic can be carried through, then a parallel inductive justification of counterinductive logic can be carried through. Table III.3 summarizes how this would be done.

Table III.3

Level of argument	Justifying arguments of the scientific inductivist	Justifying arguments of the counter-conductivist
⋮	⋮	⋮
3:	Rules of level 2 of scientific inductive logic have worked well in the past. / They will work well next time.	Rules of level 2 of counter-inductive logic have not worked well in the past. / They will work well next time.
2:	Rules of level 1 of scientific inductive logic worked well in the past. / They will work well next time.	Rules of level 1 of counter-inductive logic have not worked well in the past. / They will work well next time.

The counterinductivist is, of course, a fictitious character. No one goes through life consistently adhering to the canons of counterinductive logic, although some of us do occasionally slip into counterinductive reasoning. The poor poker player who thinks that his luck is due to change because he has been losing so heavily is a prime example. But aside from a description of gamblers' rationalizations, counterinductive logic has little practical significance.

It does, however, have great theoretical significance. For what we have shown is that if scientific inductive logic meets the conditions laid down under Rational Justification, Suggestion II, so does counterinductive logic. This is sufficient to show that Suggestion II is inadequate as a definition for rational justification. A rational justification of a system of inductive logic must provide reasons for using that system rather than any other. Thus, if two inconsistent systems, scientific induction and counterinduction, can meet the conditions of Suggestion II, then Suggestion II cannot be an adequate definition of rational justification. The arguments examined in this section do show that scientific inductive logic meets the conditions of Suggestion II, but these arguments do not rationally justify scientific induction.

This is not to say that what has been pointed out is not both important and interesting. Let us say that any system of inductive logic that meets the conditions of Suggestion II is *inductively coherent with the facts*. It may be true that for a system of inductive logic to be rationally justified it must be inductively coherent with the facts; that is, that inductive coherence with the facts may be a necessary condition for rational justification. But the example of the counterinductivist shows conclusively that inductive coherence with the facts is not by itself sufficient to rationally justify a system of inductive logic. Consequently, the inductive justification of scientific inductive logic fails.

We may summarize our discussion of the inductive justification of induction as follows:

1. The proponent of the inductive justification of scientific induction points out that scientific inductive logic is inductively coherent with the facts.
2. He claims that this is sufficient to rationally justify scientific inductive logic.
3. But it is not sufficient since counterinductive logic is also inductively coherent with the facts.
4. Nevertheless it is important and informative since we can imagine circumstances in which scientific inductive logic would not be inductively coherent with the facts.
5. The proponent of the inductive justification of scientific induction has also succeeded in calling to our attention the fact that there are various levels of induction.

Suggested readings

John Stuart Mill, "The Ground of Induction," reprinted in *A Modern Introduction to Philosophy*, Paul Edwards and Arthur Pap, Eds. (New York: The Free Press, 1973), pp. 133–41.

F. L. Will, "Will the Future Be Like the Past?" reprinted in *A Modern Introduction to Philosophy* (rev. ed.), Paul Edwards and Arthur Pap, Eds. (New York: The Free Press, 1973), pp. 148–58.

Max Black, "Inductive Support of Inductive Rules," *Problems of Analysis* (Ithaca, New York: Cornell University Press, 1954), pp. 191–208.

All of these authors are arguing for some type of inductive justification of induction, although none of them holds the exact position outlined in this section, which is a synthesis of several viewpoints.

III.4. THE PRAGMATIC JUSTIFICATION OF INDUCTION. Remember that the traditional problem of induction can be formulated as a dilemma: If the reasoning we use to rationally justify scientific inductive logic

is to have any strength at all it must be either deductively valid or inductively strong. But if we try to justify scientific inductive logic by means of a deductively valid argument with premises that are known to be true, our conclusion will be too weak. And if we try to use an inductively strong argument, we are reduced to begging the question. Whereas the proponent of the *inductive* justification of scientific induction attempts to go over the second horn of the dilemma, the proponent of the *pragmatic* justification of induction attacks the first horn; he attempts to justify scientific inductive logic by means of a deductively valid argument.

The pragmatic justification of induction was proposed by Herbert Feigl and elaborated by Hans Reichenbach, both founders of the logical empiricist movement. Reichenbach's pragmatic justification of induction is quite complicated, for it depends on what he believes are the details (at least the basic details) of scientific inductive logic. Thus, no one can fully understand Reichenbach's arguments until he has studied Reichenbach's definition of probability and the method he prescribes for discovering probabilities. We shall return to these questions later; at this point we will discuss a simplified version of the pragmatic justification of induction. This version is correct as far as it goes. Only bear in mind that there is more to be learned.

Reichenbach wishes to justify scientific inductive logic by a deductively valid argument. Yet he agrees with Hume that no deductive valid argument with premises that are known to be true can give us the conclusion that scientific induction will give us true conclusions most of the time. He agrees with Hume that the conditions of Rational Justification, Suggestion I, cannot be met. Since he fully intends to rationally justify scientific inductive logic, the only path open to him is to argue that the conditions of Rational Justification, Suggestion I, need not be met in order to justify a system of inductive logic. He proceeds to advance his own suggestion as to what is required for rational justification and to attempt to justify scientific inductive logic in these terms.

If Hume's arguments are correct, there is no way of showing that scientific induction will give us true conclusions from true premises most of the time. But since Hume's arguments apply equally well to any system of inductive logic there is no way of showing that any competing system of inductive logic will give us true conclusions from true premises most of the time either. Thus, scientific inductive logic has the same status as all other systems of inductive logic in this matter. No other system of inductive logic can be demonstrated to be superior to scientific inductive logic in the sense of showing that it gives true conclusions from true premises more often than scientific inductive logic.

Reichenbach claims that although it is impossible to show that any inductive method will be successful, it can be shown that scientific induction will be successful, if any method of induction will be successful. In other words, it is possible that no inductive logic will guide us to e-arguments that give us true conclusions most of the time, but if any method will then scientific inductive logic will also. If this can be shown, then it would seem fair to say that scientific induction has been rationally justified. After all, we must make some sort of judgments, conscious or unconscious, as to the inductive strength of arguments if we are to live at all. We must base our decisions on our expectations of the future, and we base our expectations of the future on our knowledge of the past and present. We are all gamblers, with the stakes being the success or failure of our plans of action. Life is an exploration of the unknown, and every human action presumes a wager with nature.

But if our decisions are a gamble and if no method is guaranteed to be successful, then it would seem rational to bet on that method which will be successful, if any method will. Suppose that you were forcibly taken into a locked room and told that whether or not you will be allowed to live depends on whether you win or lose a wager. The object of the wager is a box with red, blue, yellow, and orange lights on it. You know nothing about the construction of the box but are told that either all of the lights, some of them, or none of them will come on. You are to bet on one of the colors. If the colored light you choose comes on, you live; if not, you die. But before you make your choice you are also told that neither the blue, nor the yellow, nor the orange light can come on without the red light also coming on. If this is the only information you have, then you will surely bet on red. For although you have no guarantee that your bet on red will be successful (after all, all the lights might remain dark) you know that if any bet will be successful, a bet on red will be successful. Reichenbach claims that scientific inductive logic is in the same privileged position vis-à-vis other systems of inductive logic as is the red light vis-à-vis the other lights.

This leads us to a new proposal as to what is required to rationally justify a system of inductive logic:

Rational Justification

Suggestion III: A system of inductive logic is rationally justified if we can show that the e-arguments that it judges inductively strong will give us true conclusions most of the time, if e-arguments judged inductively strong by any method will.

Reichenbach attempts to show that scientific inductive logic meets the conditions of Rational Justification, Suggestion III, by a deductively valid argument. The argument goes roughly like this:

Either nature is uniform or it is not.
If nature is uniform, scientific induction will be successful.

If nature is not uniform, then no method will be successful.

If any method of induction will be successful, then scientific induction will be successful.

There is no question that this argument is deductively valid, and the first and second premises are surely known to be true. But how do we know that the third premise is true? Couldn't there be some strange inductive method that would be successful even if nature were not uniform? How do we know that for any method to be successful nature must be uniform?

Reichenbach has a response ready for this challenge. Suppose that in a completely chaotic universe, some method, call it method X, were successful. Then there is still at least one outstanding uniformity in nature: the uniformity of method X's success. And scientific induction would discover *that* uniformity. That is, if method X is successful on the whole, if it gives us true predictions most of the time, then sooner or later the statement "Method X has been reliable in the past" will be true, and the following argument would be judged inductively strong by scientific inductive logic:

Method X has been reliable in the past.

Method X will be reliable in the future.

Thus, if method X is successful, scientific induction will also be successful in that it will discover method X's reliability, and, so to speak, license method X as a subsidiary method of prediction. This completes the proof that scientific induction will be successful if any method will.

The job may appear to be done, but in fact there is a great deal more to be said. In order to analyze just what has been proved and what has not, we shall use the idea of levels of inductive logic, which was developed in the last section. When we talk about a method, we are really talking about a system of inductive logic, while glossing over the fact that a system of inductive logic is composed of distinct levels of rules. Let us now pay attention to this fact. Since a system of inductive logic is composed of distinct levels of rules, in order to justify that system we would have to justify each level of its rules. Thus, to justify scientific inductive logic we would have to justify level 1 rules of scientific inductive logic, level 2 rules of scientific inductive logic, level 3 rules of scientific inductive logic, and so on. If each of these levels of rules is to be justified in accordance with the principle "It is rational to rely on a method that is successful if any method is successful," then the pragmatic justification of induction must establish the following:

1: Level 1 rules of scientific induction will be successful if level 1 rules of any system of inductive logic will be successful.

2: Level 2 rules of scientific induction will be successful if level 2 rules of any system of inductive logic will be successful.

⋮ ⋮

k: Level k rules of scientific induction will be successful if level k rules of any system of inductive logic will be successful.

But if we look closely at the pragmatic justification of induction, we see that it does not establish this but rather something quite different.

Suppose that system X of inductive logic is successful on level 1. That is, the arguments that it judges to be inductively strong give us true conclusions from true premises most of the time. Then sooner or later an argument on level 2 which is judged inductively strong by scientific inductive logic, that is:

Rules of level 1 of system X have been reliable in the past.

Rules of level 1 of system X will be reliable in the future.

will come to have a premise that is known to be true. If the rules on level 1 of system X give true predictions most of the time, then sooner or later it will be true that they have given us true predictions most of the time *in the past*. And once we have this premise, scientific induction on level 2 leads us to the conclusion that they will be reliable in the future.

Thus, what has been shown is that if any system of inductive logic has successful rules on level 1, then scientific induction provides a justifying argument for these rules on level 2. Indeed, we can generalize this principle and say that if a system of inductive logic has successful rules on a given level, then scientific induction provides a justifying argument on the next highest level. More precisely, the pragmatist has demonstrated the following: If system X of inductive logic has rules on level k which pick out, as inductively strong arguments of level k, those which give true predictions most of the time, then there is an argument on level $k + 1$, which is judged inductively strong by the rules of level $k + 1$ of scientific inductive logic, which has as its conclusion the statement that the rules of system X on level k are reliable, and which has a premise that will sooner or later be known to be true.

Now this is quite different from showing that if any method works on any level then scientific induction will also work on *that* level, or even from showing that if any method works on level 1 then scientific induction will work on level 1. Instead what has been shown is that if any other method is generally successful on level 1 then scientific induction will have at least one notable success on level 2: it will eventually predict the continued success of that other method on level 1.

Although this is an interesting and important conclusion, it is not sufficient for the task at hand. Suppose we wish to choose a set of rules for level 1. In order to be in a position analogous to the wager about the box with the colored lights, we would have to know that scientific induction would be successful on level 1 if any method were successful on level 1. But we do not know this. For all we know, scientific induction might fail on level 1 and another method might be quite successful. If this were the case, scientific induction on level 2 would eventually tell us so, but this is quite a different matter.

In summary, the attempt at a pragmatic justification of induction has made us realize that a deductively valid justification of scientific induction would be acceptable if it could establish that: if any system of inductive logic has successful rules on a given level, then scientific inductive logic will have successful rules on that level. But the arguments advanced in the pragmatic justification fail to establish this conclusion. Instead, they show that if any system of inductive logic has successful rules on a given level, then scientific inductive logic will license a justifying argument for those rules on the next higher level.

Both the attempt at a pragmatic justification and the attempt at an inductive justification have failed to provide an absolute justification of scientific induction. Nevertheless, both of them have brought forth useful facts. For instance, the pragmatic justification of induction shows one clear advantage of scientific induction over counterinduction. The counterinductivist cannot prove that if any method is successful on level 1, counterinduction on level 2 will eventually predict its continued success. In fact some care is required to even give a logically consistent formulation of counterinduction as a general policy.

It seems, then, that there is still room for constructive thought on the problem, and that we can learn much from previous attempts to solve it.

Suggested reading

Hans Reichenbach, *Experience and Prediction: An Analysis of the Foundations and the Structure of Knowledge*. (Chicago: University of Chicago Press, 1938).

III.5. SUMMARY. We have developed the traditional problem of induction and discussed several answers to it. We found that each position we discussed had a different set of standards for rational justification of a system of inductive logic.

I. *Position*: The original presentation of the traditional problem of induction.
Standard for Rational Justification: A system of inductive logic is rationally justified if and only if it is shown that the e-arguments that it judges inductively strong yield true conclusions most of the time.

II. *Position*: The inductive justification of induction.
Standard for Rational Justification: A system of inductive logic is

rationally justified if for every level (k) of rules of that system there is an e-argument on the next highest level ($k + 1$) which:

i. Is judged inductively strong by its own system's rules.

ii. Has as its conclusion the statement that the system's rules on the original level (k) will work well next time.

III. *Position*: The pragmatic justification of induction.

Standard for Rational Justification: A system of inductive logic is rationally justified if it is shown that the e-arguments that it judges inductively strong yield true conclusions most of the time, if e-arguments judged inductively strong by any method will.

The attempt at an inductive justification of scientific inductive logic taught us to recognize different levels of arguments and corresponding levels of inductive rules. It also showed that scientific inductive logic meets the standards for Rational Justification, Suggestion II. However, we saw that Suggestion II is really not a sense of rational justification at all, for both scientific inductive logic and counterinductive logic can meet its conditions. Thus, it cannot justify the choice of one over the other.

The attempt at a pragmatic justification of scientific inductive logic showed us that Suggestion III, properly interpreted in terms of levels of induction, would be an acceptable sense of rational justification, although it would be a weaker sense than that proposed in Suggestion I. However, the pragmatic justification fails to demonstrate that scientific induction meets the conditions of Suggestion III.

It seems that we cannot make more progress in justifying inductive logic until we make some progress in saying exactly what scientific inductive logic is. The puzzles to be discussed in the next chapter show that we have to be careful in specifying the nature of scientific inductive logic.

IV

The Goodman Paradox and The New Riddle of Induction

IV.1. INTRODUCTION. In Chapter III we presented some general specifications for a system of scientific inductive logic. We said it should be a system of rules for assigning inductive probabilities to arguments, with different levels of rules corresponding to the different levels of arguments. This system must accord fairly well with common sense and scientific practice. It must on each level presuppose, in some sense, that nature is uniform and that the future will resemble the past. These general specifications were sufficient to give us a foundation for surveying the traditional problem of induction and the major attempts to solve or dissolve it.

However, to be able to apply scientific inductive logic, as a rigorous discipline, we must know precisely what its rules are. Unfortunately no one has yet produced an adequate formulation of the rules of scientific inductive logic. In fact, inductive logic is in much the same state as deductive logic was before Aristotle. This unhappy state of affairs is not due to a scarcity of brainpower in the field of inductive logic. Some of the great minds of history have attacked its problems. The distance by which they have fallen short of their goals is a measure of the difficulty of the subject. Formulating the rules of inductive logic, in fact, appears to be a more difficult enterprise than doing the same for deductive logic. Deductive logic is a "yes or no" affair; an argument is either deductively valid or it is not. But inductive strength is a matter of degree. Thus, while deductive logic must *classify* arguments as valid or not, inductive logic must *measure* the inductive strength of arguments.

Setting up such rules of measurement is not an easy task. It is in fact beset with so many problems that some philosophers have been convinced it is impossible. They maintain that a system of scientific induction cannot be constructed; that prediction of the future is an art, not a science; and that we must rely on the intuitions of experts, rather than on scientific inductive logic, to predict the future. We can only hope that this gloomy doctrine is as mistaken as the view of those early Greeks who believed deductive logic could never be reduced to a precise system of rules and must forever remain the domain of professional experts on reasoning.

If constructing a system of scientific inductive logic were totally impossible, we would be left with an intellectual vacuum, which could not be filled by appeal to "experts." For, to decide whether someone is an expert predictor or a charlatan, we must assess the evidence that his predictions will be correct.