

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 homicidal 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.

- 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.
- 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.
- My nose itches; therefore I am going to have a fight.
- Brutus said that Caesar was ambitious, and Brutus was an honorable man. Therefore Caesar must have been ambitious.
- 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;

²Although 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:



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:

- The meeting place is either the gym or the cafeteria.
The meeting place is not the gym.
 The meeting place is the cafeteria.

- A good meal has always made me feel better.
A good meal today will make me feel better.
- Many great leaders have been crazy.
Everyone who isn't a leader is sane.
- 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.
- 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 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 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.