Chapter 6

A Philosophical Interlude:

Problems and Puzzles of Induction

The issues discussed in Part One of this book are, by and large, issues involving basic topics in the history and philosophy of science, and that moreover provide background material for topics explored in Parts Two and Three. In this chapter, we take a sort of philosophical interlude. The problems and puzzles we explore here are primarily philosophical, in the sense that they arose from and are mainly discussed by philosophers, rather than being issues having a practical effect on the everyday workings of science. These topics also provide something of an interlude, in the sense that, unlike the other topics of this part of the book, these issues are not necessarily needed as background for material in later chapters. Nonetheless, the problems and puzzles we discuss are of general interest, in that they illustrate some of the deeply puzzling aspects of some of the reasoning that is most basic to science.

I should mention that these issues do not tend to strike people, on first hearing of them, as deep or puzzling or profound. I recall my first reaction to these problems, when first learning

about them years ago, was along the lines of thinking they seemed to consist mainly of philosophical nonsense. And they did not initially strike me as deep or difficult at all–my initial sense was that they could all be resolved without a great deal of hard thought.

But after a while, one realizes that these problems do not admit of easy answers, and that they raise issues that are deeply puzzling. My main goal in this chapter is to introduce you to a few of these philosophical problems, all of which concern inductive reasoning. I encourage you to let these problems simmer for a while, and in so doing (one hopes) come to appreciate how puzzling these issues are. We will, in particular, look at Hume's problem of induction, Hempel's raven paradox, and Goodman's new riddle of induction. We will begin with Hume's problem.

Hume's Problem of Induction

David Hume (1711 - 1776) was apparently the first to notice a puzzling aspect of inductive reasoning, and his observation now is typically referred to as Hume's problem of induction. Understanding Hume's point requires reaching one of those "Aha..." moments. If you really grasp Hume's point, you will see that it is an extraordinarily puzzling point about our most common types of everyday reasoning, in particular, about inferences concerning the future. Let's begin with a quick point about reasoning in general.

When we reason, for example, when we present or consider arguments, our arguments almost always contain implied premises. Implied premises, as the name suggests, are premises that are necessary in order for the reasoning to be plausible, but that are implied rather than explicitly stated. For example, suppose we agree to meet for lunch downtown this Sunday, but your car is in the shop and you are not sure how you can get to the restaurant. Suppose I tell you that there is a local bus that runs from your house to the restaurant, and so you can take the bus

to our lunch meeting. Implied in this informal bit of reasoning, but not explicitly stated, is the premise that the buses run on Sundays. If we use brackets to indicate implied premises, the reasoning can be summarized as follows:

There is a bus that runs from your house to the restaurant.
[The buses run on Sundays.]
so You can take the bus to our lunch meeting this Sunday.

Again, almost all reasoning contains implied premises, and there is nothing particularly surprising or unusual about this fact.

As noted, Hume's problem of induction concerns inferences involving the future, so let's now consider a typical inference about the future. Consider, for example, the following perfectly ordinary piece of inductive reasoning:

In our past experience, the sun has always risen in the east.
so In the future, the sun will probably continue to rise in the east.

Note the logical form of this piece of reasoning, which is:

In our past experience, has always (or at least regularly) occurred.
so In the future, will probably continue to occur.

So far, there is nothing particularly unusual about this reasoning. We simply have a typical inductive inference with a quite common logical form, of the sort we use all the time. But

Hume was apparently the first to notice something interesting about reasoning of this sort. In particular, Hume noticed that this sort of reasoning contains an implied, but crucial, premise. In particular, this sort of reasoning requires the following crucial but unstated premise:

The future will continue to be like the past.

Given this, and again using brackets to indicate implied premises, the reasoning above is more accurately captured as:

In our past experience, the sun has always risen in the east.
[The future will continue to be like the past.]
so In the future, the sun will probably continue to rise in the east.

And more generally speaking, the form of reasoning noted above is better captured as:

In our past experience, has always (or at least regularly) occurred.
[The future will continue to be like the past.]
So In the future, will probably continue to occur.

The first important point to note is why this implied premise is necessary. The reason this implied premise is necessary in inferences about the future is simply that, if the future does not continue to be like the past, then there is no reason to think that past experience will be any guide to what future experience will be like. In other words, if the statement noted above, that the future will continue to be like the past, is not correct, then past experience is no guide to the

future. And so inferences about the future would not be reliable.

This is a crucial point to understand, so we should pause for a moment to get clear on this issue. To help illustrate, consider Robert Heinlein's novel *Job*. In this novel, the two key characters continually wake up to find themselves in worlds that are slightly different from the world they inhabited the day before. For example, one day they might wake up in a world in which the monetary system is slightly different from the world of the day before (and hence any money they have on them from the previous day is no longer worth anything). One day they might inhabit a world in which everyone obeys traffic laws, and the next day they wake up in a world in which violating traffic laws is the norm. And in general, each day their world is different from the day before. Since the world they inhabit is constantly changing, they don't know what to expect from day to day. For them, the future is *not* like the past. As a result, they are unable to make the sorts of inductive inferences about the future that we take for granted. (Actually, about the only inductive inference they can make about the future is that the future will *not* continue to be like the past. And this, of course, is not a particularly helpful inference.)

In short, the statement noted above, that the future will continue to be like the past, is a necessary, though generally unrecognized, implied premise in every piece of reasoning we do about the future. Recognizing that this implied premise is a necessary premise in reasoning about the future is the first key point in understanding Hume's problem of induction. So this is a good place to pause and to make sure this point is clear.

Now, if the statement "the future will continue to be like the past" is a necessary implied premise in any reasoning about the future, then it is clear that the degree of confidence we have in inferences about the future depends crucially on the degree of confidence we have in this statement. The obvious next question, then, is what reason do we have to think the future will continue to be like the past?

Our main (perhaps only) reason for believing the future will continue to be like the past seems to boil down to the fact that today was pretty much like yesterday (heavy things still fall down today, the sun again rose in the east, day was followed by night, and so on). Yesterday was pretty much like the day before it, and that day was pretty much like the day before it, and so on. In short, in our past experience, each day has tended to be more or less like the day before. And this seems to be the basis for our believing that, in the future, things will be more or less like they always have been. In short, if we ask the question "why believe the future will continue to be like the past?," the best reason we can give is summed up in the following inference:

In our past experience, the future was like the past.

so The future will probably continue to be like the past.

But note that this inference is an inference about the future. And again, any inference about the future, including the inference immediately above, depends on the implied premise that the future will continue to be like the past. When this implied premise is made clear, the inference above is better represented as:

In our past experience, the future was like the past.

[The future will continue to be like the past.]

so The future will probably continue to be like the past.

But this reasoning is blatantly circular, that is, it assumes, as one of its premises, the very conclusion it is trying to establish. In other words, the inference summarized above depends on the assumption that its own conclusion is true. And that is blatantly circular and hence provides

no justification for accepting the conclusion.

In summary, Hume's point is that every instance of inductive reasoning depends upon the implied premise that the future will continue to be like the past. But the main (and seemingly only) way to justify this implied premise is circular, and hence it appears that this crucial implied premise cannot be justified. So inferences about the future depend on an assumption that cannot be justified, and so inferences about the future cannot themselves be logically justified.

Before closing this section, a few final comments are in order. First, note how general is Hume's point. Hume's point covers *all* inferences about the future, whether they are garden-variety inferences (such as the sun rising in the east), or inferences about scientific laws continuing to hold in the future, or beliefs about mathematics being the same in the future as it has been in the past, and so on.

Second–and this is an important point in understanding Hume–Hume was *not* trying to convince us that we ought not to make inferences about the future. Hume thought that making inferences about the future is part of our natures–we can no more stop making inferences about the future than we can voluntarily stop breathing. His question was whether or not we can *logically justify* our inferences about the future, and his answer was that we can not.

Hempel's Raven Paradox

Carl Hempel (1905 - 1997) was an influential 20th century philosopher, working mainly in the philosophy of science. As you might guess, his raven paradox was originally presented using ravens as an example, though it might be easier to see the relevance of the paradox if we use a somewhat different example. As an example to illustrate Hempel's raven paradox, suppose you and I are astronomers, and our main project involves gathering information on quasars. As a

brief background, quasars are a relatively recent discovery, with the first being discovered about 40 years ago. Even after 40 years of research not a great deal is known about quasars (though there are some interesting and reasonably plausible theories about quasars that have been developed recently). At any rate, some of the basic facts about quasars are that they seem to emit an enormous amount of energy, and they all seem to be located a great distance from Earth.

Now, suppose we are working in the early years of research on quasars. Suppose we notice that the first few quasars detected were all located a great distance from Earth, and one question we become interested in is whether all quasars are located a great distance from Earth. As the years go on, we (and other astronomers) continue to observe more quasars, and continue to note that every one we observe is located far away. So far, so good. We seem to be dealing with a fairly common situation, in which our observations are providing inductive support for the statement "all quasars are located a great distance from Earth."

There is nothing particularly puzzling about the situation described so far. When we are considering a general statement such as the one about quasars, and we observe a good number of instances that are in agreement with the statement, and none that run contrary to it, we tend to take this as inductive support for the statement.

The puzzle arises, as Hempel noted, when we consider the logical structure of general statements such as "all quasars are located a great distance from Earth." A general statement such as this is logically equivalent to its contrapositive, in this case, to the statement "all objects not located a great distance from Earth are not quasars." In other words, the statement

(i) All quasars are located a great distance from Earth

and the statement

(ii) All objects not located a great distance from Earth are not quasars

are logically equivalent statements.

Above we noted that each time we observed a quasar that was located a great distance from Earth (and again assuming we observe no instances running contrary to the statement), each observation helps support the statement that all quasars are located a great distance from Earth. To be consistent, then, each time we observe an object not located a great distance from Earth that is not a quasar, we have to accept that this observation supports (ii), that is, that all objects not located a great distance from Earth are not quasars.

This, too, by itself is not necessarily a problem or a puzzle. But now recall what we noted above, namely, that (i) and (ii) are equivalent. If (i) and (ii) are equivalent, then any support for (i) would have to count equally as support for (ii), and likewise, any support for (ii) would have to count equally as support for (i). And this gets right to the heart of the puzzle: whenever we have an observation that supports (ii), it seems that the observation must equally support (i).

So, for example, the book in your hand is an object not located a great distance from the Earth and is also not a quasar, so observing this book supports (ii). And for the reasons just noted, this observation should equally support (i). But this seems just crazy–surely you cannot help confirm a substantive scientific claim about quasars by making a trivial observation about the book in your hand.

The two paragraphs above are the heart of Hempel's raven paradox. At bottom, the paradox revolves around the fact that statements such as (i) and (ii) are logically equivalent. So support for one should constitute support for the other. But this seems entirely wrong–the claim that all quasars are located a great distance from Earth is a non-trivial scientific claim, and surely it cannot be supported by, say, trivially observing objects in the room around us.

As with Hume's problem, do not misconstrue Hempel's point. He is certainly not claiming that a trivial observation about a book in front of you actually helps support a substantive scientific claim about quasars. But he is pointing out that there is something odd about what seems a very basic pattern of inductive reasoning. Also, as noted above, Hempel's raven paradox does not constitute a practical problem, in that it is not generally a problem that affects the actual conduct of science. But undoubtably, inductive reasoning supporting general statements, such as that all quasars are located a great distance from the Earth, is an important component of science. And Hempel's raven paradox suggests that there is something deeply puzzling about the nature of such reasoning.

Goodman's Gruesome Problem

Hume's problem of induction, discussed above, is now sometimes referred to as the "old" riddle of induction, in contrast to the "new" riddle of induction put forth by Nelson Goodman. Goodman (1906 - 1998) was a broad-ranging philosopher, with interests ranging from logic to epistemology to the arts. He was apparently the first to notice another odd feature of certain types of inductive reasoning, and here we focus just on this issue.

Consider a statement such as "all emeralds are green." This statement seems to be highly supported by experience, in particular, every emerald we have observed has been green, and moreover, we have never observed an emerald that was not green. With respect to emeralds, the predicate "green" would seem to be what Goodman called a "projectible" predicate, that is, a predicate for which we can, based on our past experience that all observed emeralds have been green, project that in the future all observed emeralds will be green.

Now, define a new predicate, which Goodman named "grue." There are a variety of ways

to define “grue,” but for our purposes (and this follows Goodman’s formulation fairly closely), say an object is grue if it is green and first observed before New Year’s Day 2020, or blue and first observed after that day. As noted above, all emeralds observed so far have been green, and none have failed to be green. And again, this seems to give us some reason to think that any emeralds we observe in the future will also be green.

But now note that every emerald observed so far was green and was first observed before New Year’s Day 2020. In other words, every single emerald observed so far has been grue, and moreover, none have failed to be grue. In other words, at least with respect to emeralds observed so far, the inductive support for the statement that, in the future, all observed emeralds will be green, is *exactly* the same as the inductive support for the statement that, in the future, all observed emeralds will be grue.

But of course, we would never make the inference that in the future, all observed emeralds will be grue. That is, although we feel justified in thinking that emeralds we observe in the future will continue to be green, we are sure that emeralds observed in the future (in particular, those first observed after New Year’s Day 2020) will not be grue.

But if it is so obvious that emeralds observed after New Year’s Day 2020 will be green, but not grue, there must be some difference between the predicate “green” and the predicate “grue.” The former, to use Goodman’s terminology, is a projectible predicate (that is, we are justified in projecting its application to emeralds into the future), but the latter is not a projectible predicate. But in general, what is the difference between projectible and non-projectible predicates?

This question, which at first glance seems as if it would be quite easy to address, has proved to be difficult. Of the responses that come immediately to mind–that predicates such as “grue” are constructed rather than“natural” predicates, that unlike ordinary predicates they

involve references to time, and so on–none of these has withstood scrutiny. So although there are numerous suggestions for distinguishing projectible from non-projectible predicates, none of the suggestions have reached anything that might be considered a consensus view.

As with Hume's problem of induction, and Hempel's raven paradox, it is important not to misunderstand Goodman's point. He is certainly not suggesting that we should believe that all emeralds observed in the future will continue to be grue. Obviously they will not. But given how obvious the difference between a predicate such as "green" and a predicate such as "grue" seems to be, one would think it would be easy to provide an account that plausibly captures the difference between projectible and non-projectible predicates. Goodman's main question was what that difference was. And as noted, even though the problem seems at first glance an easy one to solve, after decades of proposed solutions, there is no agreement that any of the proposals provide an adequate solution. So once again, although Goodman's new riddle of induction is certain not a practical problem, in that it does not affect the everyday workings of science, the problem raises puzzling questions about inductive reasoning, and in particular, about the differences between predicates whose use we are confident can be projected into the future, and those that cannot be.

Concluding Remarks

As noted at the outset of this chapter, the issues discussed above are clearly philosophical issues, rather than issues that affect working scientists. And as noted, these problems tend to seem, at first, to be problems that should be easy to solve. Yet the fact that these problems resist solution suggests that there is something deeply puzzling about some of our most basic types of inductive reasoning.

Also as noted at the beginning of this chapter, it generally takes some time to fully appreciate these problems. With this in mind, I would encourage you to keep these problems in the back of your mind, and let them simmer a while. In the meantime, we will move on to discuss issues that will arise repeatedly in historical examples from the history of science, these being issues surrounding the notion of falsifiability.