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.

And to assess this evidence, we must appeal to the second level of scientific inductive logic.

Fortunately there are grounds for hope. Those who have tried to construct a system of scientific inductive logic have made some solid advances. Although the intellectual jigsaw puzzle has not been put together, we at least know what some of the pieces look like. Later we shall examine some of these "building blocks" of inductive logic, but first we shall try to put the problem of constructing a system of scientific induction in perspective by examining one of the main obstacles to this goal.

IV.2. REGULARITIES AND PROJECTION. At this point you may be puzzled as to why the construction of a system of scientific inductive logic is so difficult. After all, we know that scientific induction assumes that nature is uniform and that the future will be like the past, so if, for example, all observed emeralds have been green, the premise embodying this information confers high probability on the conclusion that the next emerald to be observed will be green. We say that scientific inductive logic *projects an observed regularity* into the future because it assigns high inductive probability to the argument:

All observed emeralds have been green.
___
The next emerald to be observed will be green.

In contrast, counterinduction would assume that the observed regular connection between being an emerald and being green would not hold in the future, and thus would assign high inductive probability to the argument:

All observed emeralds have been green.
___
The next emerald to be observed will not be green.

So it seems that scientific induction, in a quite straightforward manner, takes observed patterns or regularities in nature and assumes that they will hold in the future. Along these same lines, the premise that 99 percent of the observed emeralds have been green would confer a slightly lower probability on the conclusion that the next emerald to be observed would be green. Why can we not simply say, then, that arguments of the form

All observed *X*'s have been *Y*'s.
___
The next observed *X* will be a *Y*.

have an inductive probability of 1, and that all arguments of the form

Ninety-nine percent of the observed *X*'s have been *Y*'s.
___
The next observed *X* will be a *Y*.

have an inductive probability of 99/100?

That is, why can we not simply construct a system of scientific induction by giving the following rule on each level?

Rule S: An argument of the form

N percent of the observed *X*'s have been *Y*'s.
___
The next observed *X* will be a *Y*.

is to be assigned the inductive probability *N*/100.

Rule S does project observed regularities into the future, but there are several reasons why it cannot constitute a system of scientific inductive logic.

The most obvious inadequacy of Rule S is that it only applies to arguments of a specific form, and we are interested in assessing the inductive strength of arguments of different forms. Consider arguments which, in addition to a premise stating the percentage of observed *X*'s that have been *Y*'s, have another premise stating how many *X*'s have been observed. Here the rule does not apply, for the arguments are not of the required form. For example, Rule S does not tell us how to assign inductive probabilities to the following arguments:

I	II
Ten emeralds have been observed.	One million emeralds have been observed.
Ninety percent of the observed emeralds have been green.	Ninety percent of the observed emeralds have been green.
The next emerald to be observed will be green.	The next emerald to be observed will be green.

Obviously scientific inductive logic should tell us how to assign inductive probabilities to these arguments, and in assigning these probabilities it should take into account that the premises of Argument II bring a much greater amount of evidence to bear than the premises of Argument I.

Another type of argument that Rule S does not tell us how to evaluate is one that includes a premise stating in what variety of circumstances the regularity has been found to hold. That is, Rule S does not tell us how to assign inductive probabilities to the following arguments:

III	IV
Every person who has taken drug *X* has exhibited no adverse side reactions.	Every person who has taken drug *X* has exhibited no adverse side reactions.
Drug *X* has only been administered to persons between 20 and 25 years of age who are in good health.	Drug *X* has been administered to persons of all ages and varying degrees of health.
The next person to take drug *X* will have no adverse side reactions.	The next person to take drug *X* will have no adverse side reactions.

Again, scientific inductive logic should tell us how to assign inductive probabilities to these arguments, and in doing so it should take into account the fact that the premises of Argument IV tell us that the regularity has been found to hold in a great variety of circumstances, whereas the premises of Argument III inform us that the regularity has been found to hold in only a limited area.

There are many other types of argument that Rule S does not tell us how to evaluate, including most of the arguments advanced as examples in Chapter I. We can now appreciate why an adequate system of rules for scientific inductive logic must be a fairly complex structure. But there is another shortcoming of Rule S which has to do with arguments to which it does apply, that is, arguments of the form:

N percent of the observed *X*'s have been *Y*'s.

The next observed *X* will be a *Y*.

The following two arguments are of that form, so we can apply Rule S to evaluate them:

V	VI
One hundred percent of the observed samples of pure water have had a freezing point of +32 degrees Fahrenheit.	One hundred percent of the recorded economic depressions have occurred at the same time as large sunspots.
The next observed sample of pure water will have a freezing point of + 32 degrees Fahrenheit.	The next economic depression will occur at the same time as a large sunspot.

If we apply Rule S we find that it assigns an inductive probability of 1 to each of these arguments. But surely Argument V has a much higher degree of inductive strength than Argument VI! We feel perfectly justified in projecting into the future the observed regular connection between a certain type of chemical compound and its freezing point. But we feel that the observed regular connection between economic cycles and sunspots is a coincidence, an accidental regularity or spurious correlation, which should not be projected into the future. We shall say that the observed regularity reported in the premise of Argument V is *projectible*, while the regularity reported in the premise of Argument VI is not. We must now sophisticate our conception of scientific inductive logic still further. Scientific inductive logic does project observed regularities into the future, but only projectible regularities. It does assume that nature is uniform and that the future will resemble the past, but only in certain respects. It does assume that observed patterns in nature will be repeated, but only certain types of patterns. Thus, Rule S is not adequate for scientific inductive logic because it is incapable of taking into account differences in projectibility of regularities.

Exercises

1. Construct five inductively strong arguments to which Rule S does not apply.
2. Give two new examples of projectible regularities and two new examples of unprojectible regularities.
3. For each of the following arguments, state whether Rule S is applicable. If it is applicable, what inductive probability does it assign to the argument?
   a. One hundred percent of the crows observed have been black.
      The next crow to be observed will be black.
   b. One hundred percent of the crows observed have been black.
      All crows are black.
   c. Every time I have looked at a calendar, the date has been before January 1, 2010.
      The next time I look at a calendar the date will be before January 1, 2010.
   d. Every time fire has been observed, it has continued to burn according to the laws of nature until extinguished.
      All unobserved fires continue to burn according to the laws of nature until extinguished.

e. Eighty-five percent of the time when I have dropped a piece of silverware, company has subsequently arrived.

The next time I drop a piece of silverware company will subsequently arrive.

IV.3. THE GOODMAN PARADOX. If one tries to construct various examples of projectible and unprojectible regularities, he will soon come to the conclusion that projectibility is not simply a "yes or no" affair, but rather a matter of degree. Some regularities are highly projectible, some have a middling degree of projectibility, and some are quite unprojectible. Just how unprojectible a regularity can be has been demonstrated by Nelson Goodman in his famous "grue-bleen" paradox.

Goodman invites us to consider a new color word, "grue." It is to have the general logical features of our old color words such as "green," "blue," and "red." That is, we can speak of things being a certain color at a certain time—for example, "John's face is red now"—and we can speak of things either remaining the same color or changing colors. The new color word "grue" is defined in terms of the familiar color words "green" and "blue" as follows:

> **Definition 6:** A certain thing, *X*, is said to be *grue* at a certain time *t* if and only if:
>
> *X* is green at *t and t* is before the year 2100
>
> *or*
>
> *X* is blue at *t and t* is during or after the year 2100.

Let us see how this definition works. If you see a green grasshopper today, you can correctly maintain that you have seen a grue grasshopper today. Today is before the year 2100, and before the year 2100 something is grue just when it is green. But if you or one of your descendants sees a green grasshopper during or after the year 2100, it would then be incorrect to maintain that a grue grasshopper had been seen. During and after the year 2100, something is grue just when it is blue. Thus, after the year 2100, a blue sky would also be a grue sky.

Suppose now that a chameleon were kept on a green cloth until the beginning of the year 2100 and then transferred to a blue cloth. In terms of green and blue we would say that the chameleon changed color from green to blue. But in terms of the new color word "grue" we would say that it remained the same color: "grue." The other side of the coin is that when something remains the same color in terms of the old color words, it will change color in terms of the new one. Suppose we have a piece of glass that is green now and that will remain green during and after the year 2100. Then we would have to say that it was grue before the year 2100 but was not grue during and after the year

2100. At the beginning of the year 2100 it changed color from grue to some other color. To name the color that it changed to we introduce the new color word "bleen." "Bleen" is defined in terms of "green" and "blue" as follows:

> **Definition 7:** A certain thing, *X*, is said to be *bleen* at a certain time *t* if and only if:
>
> *X* is blue at *t and t* is before the year 2100
>
> *or*
>
> *X* is green at *t and t* is during or after the year 2100.

Thus, before the year 2100 something is grue just when it is green and bleen just when it is blue. In or after the year 2100 something is grue just when it is blue and bleen just when it is green. In terms of the old color words the piece of glass remains the same color (green), but in terms of the new color words the piece of glass changes color (from grue to bleen).

Imagine a tribe of people speaking a language that had "grue" and "bleen" as basic color words rather than the more familiar ones that we use. Suppose we describe a situation in our language—for example, the piece of glass being green before the year 2100 and remaining green afterward—in which we would say that there is no change in color. But if they correctly describe the same situation in their language, then, *in their terms,* there is a change. This leads to the important and rather startling conclusion that whether a certain situation involves change or not may depend on the descriptive machinery of the language used to discuss that situation.

One might object that "grue" and "bleen" are not acceptable color words because they have reference to a specific date in their definitions. It is quite true that *in our language,* in which blue and green are the basic color words, grue and bleen must be defined not only in terms of blue and green but also in terms of the date "2100 A.D." But a speaker of the grue-bleen language could maintain that definitions of our color words in his language must also have reference to a specific date. In the grue-bleen language, "grue" and "bleen" are basic, and "blue" and "green" are defined as follows:

> **Definition 8:** A certain thing, *X*, is said to be *green* at a certain time *t* if and only if:
>
> *X* is grue at *t and t* is before the year 2100
>
> *or*
>
> *X* is bleen at *t and t* is during or after the year 2100.
>
> **Definition 9:** A certain thing *X* is said to be *blue* at a certain time *t* if and only if:

X is bleen at *t and t* is before the year 2100

or

X is grue at *t and t* is during or after the year 2100.

Defining the old color words in terms of the new requires reference to a specific date as much as defining the new words in terms of the old. So the formal structure of their definitions gives no reason to believe that "grue" and "bleen" are not legitimate, although unfamiliar, color words.

Let us see what can be learned about regularities and projectibility from these new color words. We have already shown that whether there is change in a given situation may depend on what linguistic machinery is used to describe that situation. We shall now show that what regularities we find in a given situation also may depend on our descriptive machinery. Suppose that at one minute to midnight on December 31, 2099, a gem expert is asked to predict what the color of a certain emerald will be after midnight. He knows that all observed emeralds have been green. He projects this regularity into the future and predicts that the emerald will remain green. Notice that this is in accordance with Rule S, which assigns an inductive probability of 1 to the argument:

> One hundred percent of the times that emeralds have been observed they have been green.
>
> ---
>
> The next time that an emerald is observed it will be green.

But if the gem expert were a speaker of the grue-bleen language, he would find a different regularity in the color of observed emeralds. He would notice that every time an emerald had been observed it had been grue. (Remember that before the year 2100 everything that is green is also grue.) Now if he followed Rule S he would project *this* regularity into the future, for Rule S also assigns an inductive probability of 1 to the argument:

> One hundred percent of the times emeralds have been observed they have been "grue."
>
> ---
>
> The next time an emerald is observed it will be "grue."

And if he projected the regularity that all observed emeralds have been grue into the future, he would predict that the emerald will remain grue. But during the year 2100 a thing is "grue" only if it is blue. So by projecting this regularity he is in effect predicting that the emerald will change from green to blue.

Now, we will all agree that this is a ridiculous prediction to make on the basis of the evidence. And no one is really claiming that it should be made. But it cannot be denied that this prediction results from the projection into the

future of an observed regularity in accordance with Rule S. The point is that the regularity of every observed emerald having been grue is a totally unprojectible regularity. And the prediction of our hypothetical grue-bleen-speaking gem expert is an extreme case of the trouble we get into when we try to project, via some rule such as Rule S, regularities that are in fact unprojectible.

The trouble we get into is indeed deep, for the prediction so arrived at will conflict with the prediction arrived at by projecting a projectible regularity. If we project the projectible regularity that every time an emerald has been observed it has been green, then we arrive at the prediction that the emerald will remain green. If we project the unprojectible regularity that every time an emerald has been observed, it has been grue, then we arrive at the prediction that the emerald will change from green to blue. These two predictions clearly are in conflict.[1]

Thus, the mistake of projecting an unprojectible regularity may not only lead to a ridiculous prediction. It may, furthermore, lead to a prediction that conflicts with a legitimate prediction which results from projecting a projectible regularity discovered in *the same set of data.* An acceptable system of scientific inductive logic must provide some means to escape this conflict. It must incorporate rules that tell us which regularities are projectible. From the discussion of accidental regularities and the sunspot theory of economic cycles, we already know that scientific inductive logic must have rules for determining projectibility. But the Goodman paradox gives this point new urgency by demonstrating how unprojectible a regularity can be and how serious are the consequences of projecting a totally unprojectible regularity.

Let us summarize what is to be learned from the discussion of "grue" and "bleen":

1. Whether we find change or not in a certain situation may depend on the linguistic machinery we use to describe that situation.
2. What regularities we find in a sequence of occurrences may depend on the linguistic machinery used to describe that sequence.
3. We may find two regularities in a sequence of occurrences, one projectible and one unprojectible, such that the predictions that arise from projecting them both are in conflict.

---

[1] Actually they are inconsistent only under the assumption that the emerald will not be destroyed before 2100 A.D., but presumably we will have independent inductive evidence for this assumption.

Exercise:

Define "grue" in terms of "blue," "green," and "bleen" without mentioning the year 2100. You can use "and," "or," and "not."

IV.4. THE GOODMAN PARADOX, REGULARITY, AND THE PRINCIPLE OF THE UNIFORMITY OF NATURE. We saw, in the last section, that projecting observed regularities into the future is not as simple as it first appears. The regularities found in a certain sequence of events may depend on the language used to describe that sequence of events. The Goodman paradox showed that if we try to project all regularities that can be found by using any language, our predictions may conflict with one another. This is a startling result, and it dramatizes the need for rules for determining projectibility in scientific induction. (This might be accomplished through the specification of the most fruitful language for scientific description of events.)

This need is further dramatized by the following, even more startling result: For any prediction whatsoever, we can find a regularity whose projection licenses that prediction. Of course, most of these regularities will be unprojectible. The point is that we need rules to eliminate those predictions based on unprojectible regularities. I shall illustrate this principle in three ways: (1) in an example that closely resembles Goodman's "grue-bleen" paradox, (2) with reference to the extrapolation of curves on graphs, (3) with reference to the problem, often encountered on intelligence tests, of continuing a sequence of numbers. The knowledge gained from this discussion will then be applied to a reexamination of the principle of the uniformity of nature.

Example 1

Suppose you are presented with four boxes, each labeled "Excelsior!" In the first box you discover a green insect; in the second, a yellow ball of wax; in the third, a purple feather. You are now told that the fourth box contains a mask and are asked to predict its color. You must look for a regularity in this sequence of discoveries, whose projection will license a prediction as to the color of the mask. Although on the face of it, this seems impossible, with a little ingenuity a regularity can be found. What is more, for any prediction you wish to make, there is a regularity whose projection will license that prediction. Suppose you want to predict that the mask will be red. The regularity is found in the following manner.

Let us define a new word, "snarf." A snarf is something presented to you in a box labeled "Excelsior!" and is either an insect, a ball of wax, a feather, or a mask. Now you have observed three snarfs and are about to observe a fourth. This is a step toward regularity, but there is still the problem that the three

observed snarfs have been different colors. One more definition is required in order to find regularity in apparent chaos. A thing X is said to be "murkle" just when:

X is an insect *and* X is green

or

X is a ball of wax *and* X is yellow

or

X is a feather *and* X is purple

or

X is some other type of thing *and* X is red.

Now we have found the regularity: all observed snarfs have been murkle. If we project this regularity into the future, assuming that the next snarf to be observed will be murkle, we obtain the required prediction.[2] The next snarf to be observed will be a mask, and for a mask to be murkle it must be red. Needless to say, this regularity is quite unprojectible. But it is important to see that we could discover an unprojectible regularity that, if it were projected, would lead to the prediction that the mask is red. And it is easy to see that, if we wanted to discover a regularity that would lead to a prediction that the mask will be a different color, a few alterations to the definition of "murkle" would accomplish this aim. This sort of thing can always be done and, as we shall see, in some areas we need not even resort to such exotic words as "snarf," "murkle," "grue," and "bleen."

Example 2

When basing predictions on statistical data we often make use of graphs, which help summarize the evidence and guide us in making our predictions. To illustrate, suppose a certain small country takes a census every 10 years, and has taken three so far. The population was 11 million at the time of the first census, 12 million at the second census, and 13 million at the third. This information is represented on a graph in Figure IV.1. Each dot represents the information as to population size gained from one census. For example, the middle dot represents the second census, taken in the year 10, and showing a population of 12 million. Thus, it is placed at the intersection of the vertical

[2] This projection is in accordance with Rule S, which assigns an inductive probability of 1 to the argument:

All observed snarfs have been murkle.

The next snarf to be observed will be murkle.

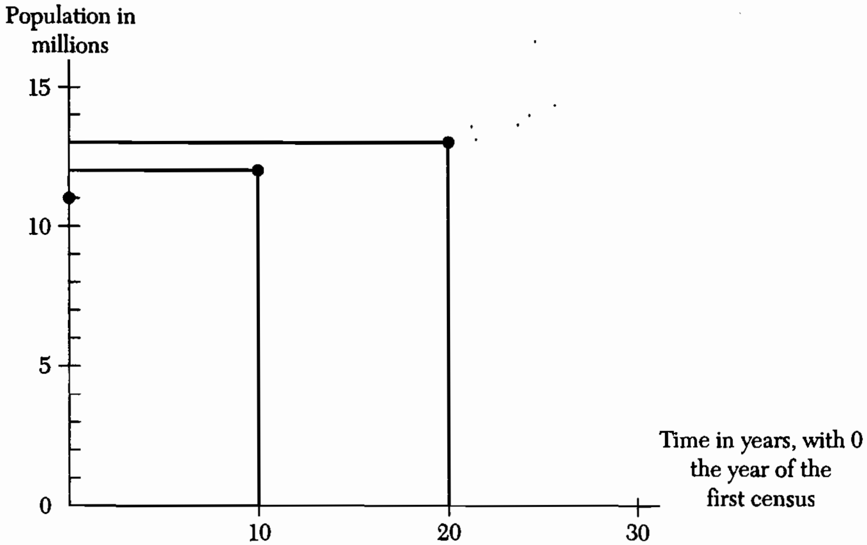

Figure IV. 1

line drawn from the year 10 and the horizontal line drawn from the population of 12 million.

Suppose now you are asked to predict the population of this country at the time of the fourth census, that is, in the year 30. You would have to look for a regularity that could be projected into the future. In the absence of any further information, you would probably proceed as follows: First you would notice that the points representing the first three census all fall on the straight line labeled A in Figure IV.2, and would then project this regularity into the future. This is in accordance with Rule S, which assigns an inductive probability of 1 to the following argument:

All points representing census so far taken have fallen on line A.

The point representing the next census to be taken will fall on line A.

This projection would lead you to the prediction that the population at the time of the fourth census will be 14 million, as shown by the dotted lines in Figure IV.2. The process by which you would arrive at your prediction is called *extrapolation*. If you had used similar reasoning to estimate the population during the year 15 at 12.5 million, the process would be called *interpolation*. Interpolation is estimating the position of a point that lies *between* the points representing the data. Extrapolation is estimating the position of a point that lies *outside* the points representing the data. So your prediction would be

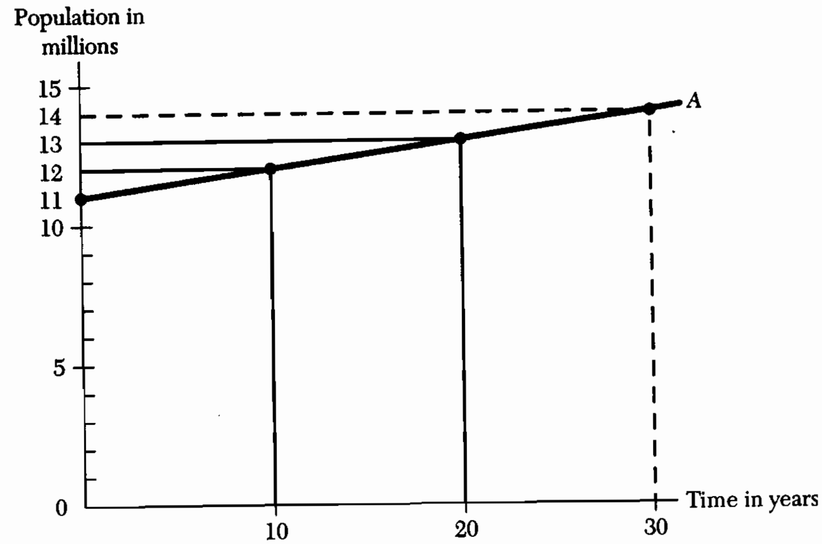

Figure IV.2

obtained by extrapolation, and your extrapolation would be a projection of the regularity that all the points plotted so far fell on line A.

But it is obvious that there are quite a few other regularities to be found in the data which you did not choose to project. As shown in Figure IV.3 there is the regularity that all the points plotted so far fall on curve B, and the regularity that all the points plotted so far fall on curve C. The projection of one of these regularities will lead to a different prediction.

If you extrapolate along curve B, you can predict that the population in the year 30 will be back to 11 million. If you extrapolate along curve C, you can predict that the population will leap to 17 million. There are indeed an infinite number of curves that pass through all the points and thus an infinite number of regularities in the data. Whatever prediction you wish to make, a regularity can be found whose projection will license that prediction.

Example 3

Often intelligence and aptitude tests contain problems where one is given a sequence of numbers and asked to continue the sequence; for example:

i. 1, 2, 3, 4, 5, . . . ;
ii. 2, 4, 6, 8, 10, . . . ;
iii. 1, 3, 5, 7, 9, ...

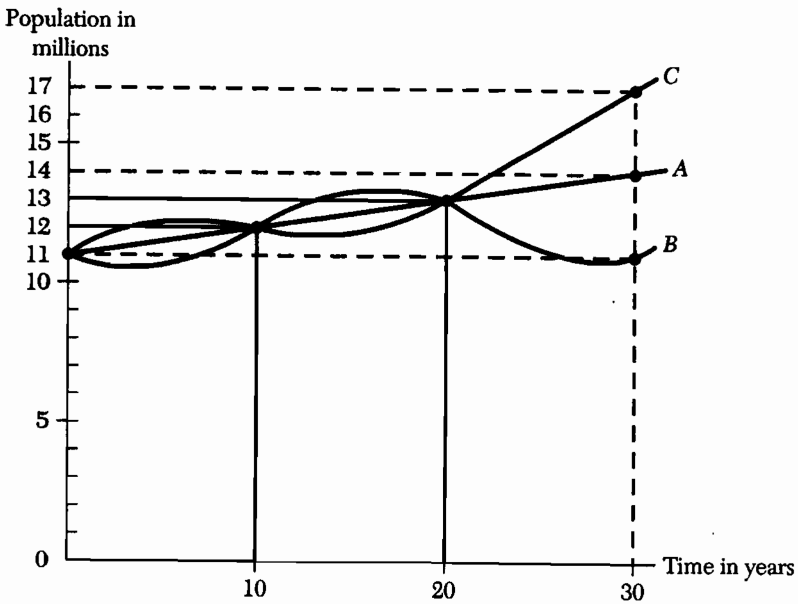

Figure IV.3

The natural way in which to continue sequence (i) is to add 6 to the end, for sequence (ii) to add 12, and for sequence (iii) to add 11. These problems are really problems of inductive logic on the intuitive level; one is asked to discover a regularity in the segment of the series given and to project that regularity in order to find the next number of the series.

Let us make this reasoning explicit for the three series given. In example (i) the first member of the series is 1, the second member is 2, the third member is 3, and, in general, for all the members given, the kth member is k. If we project this regularity to find the next member of the series, we will reason that the sixth member is 6, which is the answer intuitively arrived at before. In example (ii) the first member is twice 1, the second is twice 2, and, in general, for all the members given, the kth member is twice k. If we project this regularity, we will reason that the sixth member is twice 6, or 12, which is the answer intuitively arrived at before. In example (iii) the first member is twice 1 less 1, the second member is twice 2 less 1, and the third member is twice 3 less 1. In general, for all the members given, the kth member is twice k less 1. If we project this regularity, we will reason that the sixth member of the series is twice 6 less 1, or 11, which is the result intuitively arrived at. We say that k is a *generating function* for the first series, $2k$ a generating function for the second series, and $2k - 1$ a generating function for the third series. Although

"generating function" may sound like a very technical term, its meaning is quite simple. It is a formula with k in it, such that if 1 is substituted for k it gives the first member of the series, if 2 is substituted for k it gives the second member, and so on.

Thus, the regularity we found in each of these series is that a certain generating function yielded all the given members of the series. This regularity was projected by assuming that the same generating function would yield the next member of the series, and so we were able to fill in the ends of the series. For example, the prediction that the sixth member of series (iii) is 11 implicitly rests on the following argument:

> For every given member of series (iii) the kth member of that series was $2k - 1$.
>
> ---
>
> For the next member of series (iii) the kth member will be $2k - 1$.

But, as you may expect, there is a fly in the ointment. If we look more closely at these examples, we can find *other* regularities in the given members of the various series. And the projection of these other regularities conflicts with the projection of the regularities we have already noted. The generating function $(k-1)\ (k-2)\ (k-3)\ (k-4)\ (k-5) + k$ also yields the five given members of series (i). (This can be checked by substituting 1 for k, which gives 1; 2 for k, which gives 2; and so on, up through 5.) But if we project this regularity, the result is that the sixth member of the series is 126!

Indeed, whatever number we wish to predict for the sixth member of the series, there is a generating function that will fit the given members of the series and that will yield the prediction we want. It is a mathematical fact that in general this is true. For any finite string of numbers which begins a series, there are generating functions that fit that string of given numbers and yield whatever next member is desired. Whatever prediction we wish to make, we can find a regularity whose projection will license that prediction.

Thus, if the intelligence tests were simply looking for the projection of a regularity, any number at the end of the series would be correct. What they are looking for is not simply the projection of a regularity but the projection of an intuitively projectible regularity.

If we have perhaps belabored the point in Examples (1), (2), and (3) we have done so because the principle they illustrate is so hard to accept. Any prediction whatsoever can be obtained by projecting regularities. As Goodman puts it, "To say that valid predictions are those based on past regularities, without being able to say *which* regularities, is thus quite pointless. Regularities are where you find them, and you can find them anywhere." An acceptable scientific inductive logic must have rules for determining the projectibility of regularities.

It remains to be shown how this discussion of regularities and projectibility bears on the principle of the uniformity of nature. Just as we saw that the naïve characterization of scientific inductive logic as a system that projects observed regularities into the future was pointless unless we can say which regularities it projects, so we shall see that the statement that scientific inductive logic presupposes the uniformity of nature is equally pointless unless we are able to say *in what respects* nature is presupposed to be uniform. For it is self-contradictory to say that nature is uniform in all respects, and trivial to say it is uniform in some respects.

In the original statement of the Goodman paradox, the gem expert, who spoke our ordinary language, assumed nature to be uniform with respect to the blueness or greenness of emeralds. Since observed emeralds had always been green, and since he was assuming that nature is uniform and that the future would resemble the past in this respect, he predicted that the emerald would remain green. But the hypothetical gem expert who spoke the grue-bleen language assumed nature to be uniform *with respect to the grueness or bleenness of emeralds*. Since observed emeralds had always been grue and since he was assuming that nature is uniform and that the future would resemble the past in this respect, he predicted that the emerald would remain grue. But we saw that these two predictions were in conflict. The future cannot resemble the past in both these ways. As we have seen, such conflicts can be multiplied *ad infinitum*. The future cannot resemble the past in all respects. It is self-contradictory to say that nature is uniform in all respects.

We might try to retreat to the claim that scientific induction presupposes that nature is uniform in some respects. But this claim is so weak as to be no claim at all. To say that nature is uniform in some respects is to say that it exhibits some patterns, that there are some regularities in nature taken as a whole (in both the observed and unobserved parts of nature). But as we have seen in this section, in any sequence of observations, no matter how chaotic the data may seem, there are always regularities. This holds not only for sequences of observations but also for nature as a whole. No matter how chaotic nature might be, it would always exhibit some patterns; it would always be uniform in some respects. These uniformities might seem highly artificial, such as a uniformity in terms of "grue" and "bleen" or "snarf" and "murkle." They might be fiendishly complex. But no matter how nature might behave, there would always be some uniformity, "natural" or "artificial," simple or complex. It is therefore trivial to say that nature is uniform in some respects. Thus, if the statement that scientific induction presupposes that nature is uniform is to convey any information at all, it must specify in what respects scientific induction presupposes that nature is uniform.

The points about regularities and projectibility and the uniformity of nature are really two sides of the same coin. There are so many regularities in any sequence of observations and so many ways for nature to be uniform that the statements "Scientific induction projects observed regularities into the future" and "Scientific induction presupposes the uniformity of nature" lose all meaning. They can, however, be reinvested with meaning if we can formulate *rules of projectibility* for scientific inductive logic. Then we could say that scientific inductive logic projects regularities that meet these standards. And that would be saying something informative. We could reformulate the principle of the uniformity of nature to mean: Nature is such that projecting regularities that meet these standards will lead to correct predictions most of the time. Thus, the whole concept of scientific inductive logic rests on the idea of projectibility. The problem of formulating precise rules for determining projectibility is the new riddle of induction.

Exercise:

In the example of the four boxes labeled "Excelsior!" find a regularity in the observations whose projection would lead to the prediction that the mask will be blue.

IV.5. SUMMARY. This chapter described the scope of the problem of constructing a system of scientific inductive logic. We began with the supposition that scientific inductive logic could be simply characterized as the projection of observed regularities into the future in accordance with some rule, such as Rule S. We saw that this characterization of scientific inductive logic is inadequate for several reasons, the most important being that too many regularities are to be found in any given set of data. In one set of data we can find regularities whose projection leads to conflicting predictions. In fact, for any prediction we choose, there will be a regularity whose projection licenses that prediction.

Scientific inductive logic must select from the multitude of regularities present in any sequence of observations, for indiscriminate projection leads to paradox. Thus, in order to characterize scientific inductive logic we must specify the rules used to determine which regularities it considers to be projectible. The problem of formulating these rules is called the new riddle of induction.

Essentially the same problem reappears if we try to characterize scientific inductive logic as a system that presupposes that nature is uniform. To say that nature is uniform in *some* respects is trivial. To say that nature is uniform in *all* respects is not only false but self-contradictory. Thus, if we are to characterize scientific inductive logic in terms of some principle of the uniformity of nature which it presupposes, we must say in what respects nature is presupposed to be uniform, which in turn determines what regularities scientific inductive logic takes to be projectible. So the problem about the uniformity of nature is just a different facet of the new riddle of induction.

The problem of constructing a system of scientific inductive logic will not be solved until the new riddle of induction and other problems have been solved. Although these solutions have not yet been found, there have been developments in the history of inductive logic which constitute progress towards a system.

In the next chapter we shall pursue an analysis of causality which casts some light on well-known features of the experimental method. Then we will discuss the major achievement of the field, the probability calculus.

Suggested reading

Nelson Goodman, *Fact, Fiction and Forecast* (Cambridge, MA: Harvard University Press, 1983).

V

Mill's Methods of Experimental Inquiry and the Nature of Causality

V.1. INTRODUCTION. One of the purposes of scientific inductive logic is to assess the evidential warrant for statements of cause and effect. But what exactly do statements claiming causal connection *mean*, and what is their relation to statements describing *de facto* regularities? These are old and deep questions and we can give only partial answers here.

In his *System of Logic*, published in 1843, John Stuart Mill discussed five "methods of experimental inquiry" that he found used in the work of contemporary scientists. When we make some simple distinctions between different senses of "cause," we will find that we can use the basics of logic introduced in Chapter I to give a logical analysis of Mill's methods.

V.2. CAUSALITY AND NECESSARY AND SUFFICIENT CONDITIONS. Many of the inquiries of both scientific research and practical affairs may be characterized as the search for the causes of certain effects. The practical application of knowledge of causes consists either in producing the cause in order to produce the effect or in removing the cause in order to prevent the effect. Knowledge of causes is the key to control of effects. Thus, physicians search for the cause of certain diseases so that they may remove the cause and prevent the effect. On the other hand, advertising men engage in motivational research into the causes of consumer demand so that they can produce the cause and thus produce the effect of consumer demand for their products.

However, the word "cause" is used in English to mean several different things. For this reason, it is more useful to talk about *necessary conditions* and *sufficient conditions* rather than about causes.

> **Definition 10:** A property F is a *sufficient condition* for a property G if and only if *whenever F is present, G is present.*
>
> **Definition 11:** A property H is a *necessary condition* for a property I if and only if *whenever I is present, H is present.*

Being run over by a steamroller is a sufficient condition for death, but it is not a necessary condition. Whenever someone has been run over by a