

fest to non-manifest cases is obviously not very different from the problem of going from known to unknown or from past to future cases. The problem of dispositions looks suspiciously like one of the philosopher's oldest friends and enemies: the problem of induction. Indeed, the two are but different aspects of the general problem of proceeding from a given set of cases to a wider set. The critical questions throughout are the same: when, how, why is such a transition or expansion legitimate? In the next lecture, then, we must see how matters stand at present with the familiar problem of induction.

Thus passes the possible. It passes, indeed, only into another and exceedingly difficult problem. But that problem has been troubling our sleep for a long time on its own account. There is perhaps some solace in the thought that at least the ghost of the possible will no longer be thumping in the attic.

III

THE NEW RIDDLE OF INDUCTION

1. *The Old Problem of Induction*

At the close of the preceding lecture, I said that today I should examine how matters stand with respect to the problem of induction. In a word, I think they stand ill. But the real difficulties that confront us today are not the traditional ones. What is commonly thought of as the Problem of Induction has been solved, or dissolved; and we face new problems that are not as yet very widely understood. To approach them, I shall have to run as quickly as possible over some very familiar ground.

The problem of the validity of judgments about future or unknown cases arises, as Hume pointed out, because such judgments are neither reports of experience nor logical consequences of it. Predictions, of course, pertain to what has not yet been observed. And they cannot be logically inferred from what has been observed; for what *has* happened imposes no logical restrictions on what *will* happen. Although Hume's dictum that there are no necessary connections of matters of fact has been challenged at times, it has withstood all attacks. Indeed, I should be inclined not merely to agree that there are no necessary connections of matters of fact, but to ask whether there

are any necessary connections at all¹—but that is another story.

Hume's answer to the question how predictions are related to past experience is refreshingly non-cosmic. When an event of one kind frequently follows upon an event of another kind in experience, a habit is formed that leads the mind, when confronted with a new event of the first kind, to pass to the idea of an event of the second kind. The idea of necessary connection arises from the felt impulse of the mind in making this transition.

Now if we strip this account of all extraneous features, the central point is that to the question "Why one prediction rather than another?", Hume answers that the elect prediction is one that accords with a past regularity, because this regularity has established a habit. Thus among alternative statements about a future moment, one statement is distinguished by its consonance with habit and thus with regularities observed in the past. Prediction according to any other alternative is errant.

How satisfactory is this answer? The heaviest criticism has taken the righteous position that Hume's account at best pertains only to the source of predictions, not their legitimacy; that he sets forth the circumstances under which we make given predictions—and in this sense explains why we make them—but leaves untouched the

¹ Although this remark is merely an aside, perhaps I should explain for the sake of some unusually sheltered reader that the notion of a necessary connection of ideas, or of an absolutely analytic statement, is no longer sacrosanct. Some, like Quine and White, have forthrightly attacked the notion; others, like myself, have simply discarded it; and still others have begun to feel acutely uncomfortable about it.

question of our license for making them. To trace origins, runs the old complaint, is not to establish validity: the real question is not why a prediction is in fact made but how it can be justified. Since this seems to point to the awkward conclusion that the greatest of modern philosophers completely missed the point of his own problem, the idea has developed that he did not really take his solution very seriously, but regarded the main problem as unsolved and perhaps as insoluble. Thus we come to speak of 'Hume's problem' as though he propounded it as a question without answer.

All this seems to me quite wrong. I think Hume grasped the central question and considered his answer to be passably effective. And I think his answer is reasonable and relevant, even if it is not entirely satisfactory. I shall explain presently. At the moment, I merely want to record a protest against the prevalent notion that the problem of justifying induction, when it is so sharply dissociated from the problem of describing how induction takes place, can fairly be called Hume's problem.

I suppose that the problem of justifying induction has called forth as much fruitless discussion as has any half-way respectable problem of modern philosophy. The typical writer begins by insisting that some way of justifying predictions must be found; proceeds to argue that for this purpose we need some resounding universal law of the Uniformity of Nature, and then inquires how this universal principle itself can be justified. At this point, if he is tired, he concludes that the principle must be accepted as an indispensable assumption; or if he is energetic and ingenious, he goes on to devise some subtle justification for it. Such an invention, however, seldom satisfies

anyone else; and the easier course of accepting an unsubstantiated and even dubious assumption much more sweeping than any actual predictions we make seems an odd and expensive way of justifying them.

2. *Dissolution of the Old Problem*

Understandably, then, more critical thinkers have suspected that there might be something awry with the problem we are trying to solve. Come to think of it, what precisely would constitute the justification we seek? If the problem is to explain how we know that certain predictions will turn out to be correct, the sufficient answer is that we don't know any such thing. If the problem is to *find* some way of distinguishing antecedently between true and false predictions, we are asking for prevision rather than for philosophical explanation. Nor does it help matters much to say that we are merely trying to show that or why certain predictions are *probable*. Often it is said that while we cannot tell in advance whether a prediction concerning a given throw of a die is true, we can decide whether the prediction is a probable one. But if this means determining how the prediction is related to actual frequency distributions of future throws of the die, surely there is no way of knowing or proving this in advance. On the other hand, if the judgment that the prediction is probable has nothing to do with subsequent occurrences, then the question remains in what sense a probable prediction is any better justified than an improbable one.

Now obviously the genuine problem cannot be one of attaining unattainable knowledge or of accounting for knowledge that we do not in fact have. A better under-

standing of our problem can be gained by looking for a moment at what is involved in justifying non-inductive inferences. How do we justify a *deduction*? Plainly, by showing that it conforms to the general rules of deductive inference. An argument that so conforms is justified or valid, even if its conclusion happens to be false. An argument that violates a rule is fallacious even if its conclusion happens to be true. To justify a deductive conclusion therefore requires no knowledge of the facts it pertains to. Moreover, when a deductive argument has been shown to conform to the rules of logical inference, we usually consider it justified without going on to ask what justifies the rules. Analogously, the basic task in justifying an inductive inference is to show that it conforms to the general rules of *induction*. Once we have recognized this, we have gone a long way towards clarifying our problem.

Yet, of course, the rules themselves must eventually be justified. The validity of a deduction depends not upon conformity to any purely arbitrary rules we may contrive, but upon conformity to valid rules. When we speak of *the* rules of inference we mean the valid rules—or better, *some* valid rules, since there may be alternative sets of equally valid rules. But how is the validity of rules to be determined? Here again we encounter philosophers who insist that these rules follow from some self-evident axiom, and others who try to show that the rules are grounded in the very nature of the human mind. I think the answer lies much nearer the surface. Principles of deductive inference are justified by their conformity with accepted deductive practice. Their validity depends upon accordance with the particular deductive inferences we actually make and sanction. If a rule yields unacceptable inferences,

we drop it as invalid. Justification of general rules thus derives from judgments rejecting or accepting particular deductive inferences.

This looks flagrantly circular. I have said that deductive inferences are justified by their conformity to valid general rules, and that general rules are justified by their conformity to valid inferences. But this circle is a virtuous one. The point is that rules and particular inferences alike are justified by being brought into agreement with each other. *A rule is amended if it yields an inference we are unwilling to accept; an inference is rejected if it violates a rule we are unwilling to amend.* The process of justification is the delicate one of making mutual adjustments between rules and accepted inferences; and in the agreement achieved lies the only justification needed for either.

All this applies equally well to induction. An inductive inference, too, is justified by conformity to general rules, and a general rule by conformity to accepted inductive inferences. Predictions are justified if they conform to valid canons of induction; and the canons are valid if they accurately codify accepted inductive practice.

A result of such analysis is that we can stop plaguing ourselves with certain spurious questions about induction. We no longer demand an explanation for guarantees that we do not have, or seek keys to knowledge that we cannot obtain. It dawns upon us that the traditional smug insistence upon a hard-and-fast line between justifying induction and describing ordinary inductive practice distorts the problem. And we owe belated apologies to Hume. For in dealing with the question how normally

accepted inductive judgments are made, he was in fact dealing with the question of inductive validity.² The validity of a prediction consisted for him in its arising from habit, and thus in its exemplifying some past regularity. His answer was incomplete and perhaps not entirely correct; but it was not beside the point. The problem of induction is not a problem of demonstration but a problem of defining the difference between valid and invalid predictions.

This clears the air but leaves a lot to be done. As principles of deductive inference, we have the familiar and highly developed laws of logic; but there are available no such precisely stated and well-recognized principles of inductive inference. Mill's canons hardly rank with Aristotle's rules of the syllogism, let alone with *Principia*

² A hasty reader might suppose that my insistence here upon identifying the problem of justification with a problem of description is out of keeping with my parenthetical insistence in the preceding lecture that the goal of philosophy is something quite different from the mere description of ordinary or scientific procedure. Let me repeat that the point urged there was that the organization of the explanatory account need not reflect the manner or order in which predicates are adopted in practice. It surely must describe practice, however, in the sense that the extensions of predicates as explicated must conform in certain ways to the extensions of the same predicates as applied in practice. Hume's account is a description in just this sense. For it is an attempt to set forth the circumstances under which those inductive judgments are made that are normally accepted as valid; and to do that is to state necessary and sufficient conditions for, and thus to define, valid induction. What I am maintaining above is that the problem of justifying induction is not something over and above the problem of describing or defining valid induction.

Mathematica. Elaborate and valuable treatises on probability usually leave certain fundamental questions untouched. Only in very recent years has there been any explicit and systematic work upon what I call the constructive task of confirmation theory.

3. *The Constructive Task of Confirmation Theory*

The task of formulating rules that define the difference between valid and invalid inductive inferences is much like the task of defining any term with an established usage. If we set out to define the term "tree", we try to compose out of already understood words an expression that will apply to the familiar objects that standard usage calls trees, and that will not apply to objects that standard usage refuses to call trees. A proposal that plainly violates either condition is rejected; while a definition that meets these tests may be adopted and used to decide cases that are not already settled by actual usage. Thus the interplay we observed between rules of induction and particular inductive inferences is simply an instance of this characteristic dual adjustment between definition and usage, whereby the usage informs the definition, which in turn guides extension of the usage.

Of course this adjustment is a more complex matter than I have indicated. Sometimes, in the interest of convenience or theoretical utility, we deliberately permit a definition to run counter to clear mandates of common usage. We accept a definition of "fish" that excludes whales. Similarly we may decide to deny the term "valid induction" to some inductive inferences that are commonly considered valid, or apply the term to others not

usually so considered. A definition may modify as well as extend ordinary usage.³

Some pioneer work on the problem of defining confirmation or valid induction has been done by Professor Hempel.⁴ Let me remind you briefly of a few of his results. Just as deductive logic is concerned primarily with a relation between statements—namely the consequence relation—that is independent of their truth or falsity, so inductive logic as Hempel conceives it is concerned primarily with a comparable relation of confirmation between statements. Thus the problem is to define the relation that obtains between any statement S_1 and another S_2 if and only if S_1 may properly be said to confirm S_2 in any degree.

With the question so stated, the first step seems obvious. Does not induction proceed in just the opposite direction from deduction? Surely some of the evidence-statements that inductively support a general hypothesis are consequences of it. Since the consequence relation is already well defined by deductive logic, will we not be on firm ground in saying that confirmation embraces the converse relation? The laws of deduction in reverse will then be among the laws of induction.

Let's see where this leads us. We naturally assume fur-

³ For a fuller discussion of definition in general see Chapter I of *The Structure of Appearance*.

⁴ The basic article is 'A Purely Syntactical Definition of Confirmation', cited in Note I.10. A much less technical account is given in 'Studies in the Logic of Confirmation', *Mind*, n.s., vol. 54 (1945), pp. 1-26 and 97-121. Later work by Hempel and others on defining *degree* of confirmation does not concern us here.

ther that whatever confirms a given statement confirms also whatever follows from that statement.⁵ But if we combine this assumption with our proposed principle, we get the embarrassing result that every statement confirms every other. Surprising as it may be that such innocent beginnings lead to such an intolerable conclusion, the proof is very easy. Start with any statement S_1 . It is a consequence of, and so by our present criterion confirms, the conjunction of S_1 and any statement whatsoever—call it S_2 . But the confirmed conjunction, $S_1 \cdot S_2$, of course has S_2 as a consequence. Thus every statement confirms all statements.

The fault lies in careless formulation of our first proposal. While some statements that confirm a general hypothesis are consequences of it, not all its consequences confirm it. This may not be immediately evident; for indeed we do in some sense furnish support for a statement when we establish one of its consequences. We settle one of the questions about it. Consider the heterogeneous conjunction:

⁵ I am not here asserting that this is an indispensable requirement upon a definition of confirmation. Since our commonsense assumptions taken in combination quickly lead us to absurd conclusions, some of these assumptions have to be dropped; and different theorists may make different decisions about which to drop and which to preserve. Hempel gives up the converse consequence condition, while Carnap (*Logical Foundations of Probability*, Chicago and London, 1950, pp. 474–6) drops both the consequence condition and the converse consequence condition. Such differences of detail between different treatments of confirmation do not affect the central points I am making in this lecture.

8497 is a prime number and the other side of the moon is flat and Elizabeth the First was crowned on a Tuesday.

To show that any one of the three component statements is true is to support the conjunction by reducing the net undetermined claim. But support⁶ of this kind is not confirmation; for establishment of one component endows the whole statement with no credibility that is transmitted to other component statements. Confirmation of a hypothesis occurs only when an instance imparts to the hypothesis some credibility that is conveyed to other instances. Appraisal of hypotheses, indeed, is incidental to prediction, to the judgment of new cases on the basis of old ones.

Our formula thus needs tightening. This is readily accomplished, as Hempel points out, if we observe that a hypothesis is genuinely confirmed only by a statement that is an instance of it in the special sense of entailing not the hypothesis itself but its relativization or restriction to the class of entities mentioned by that statement. The relativization of a general hypothesis to a class results from restricting the range of its universal and existential quantifiers to the members of that class. Less technically, what the hypothesis says of all things the evidence statement says of one thing (or

⁶ Any hypothesis is 'supported' by its own positive instances; but support—or better, direct factual support—is only one factor in confirmation. This factor has been separately studied by John G. Kemeny and Paul Oppenheim in 'Degree of Factual Support', *Philosophy of Science*, vol. 19 (1952), pp. 307–24. As will appear presently, my concern in these lectures is primarily with certain other important factors in confirmation, some of them quite generally neglected.

of one pair or other n -ad of things). This obviously covers the confirmation of the conductivity of all copper by the conductivity of a given piece; and it excludes confirmation of our heterogeneous conjunction by any of its components. And, when taken together with the principle that what confirms a statement confirms all its consequences, this criterion does not yield the untoward conclusion that every statement confirms every other.

New difficulties promptly appear from other directions, however. One is the infamous paradox of the ravens. The statement that a given object, say this piece of paper, is neither black nor a raven confirms the hypothesis that all non-black things are non-ravens. But this hypothesis is logically equivalent to the hypothesis that all ravens are black. Hence we arrive at the unexpected conclusion that the statement that a given object is neither black nor a raven confirms the hypothesis that all ravens are black. The prospect of being able to investigate ornithological theories without going out in the rain is so attractive that we know there must be a catch in it. The trouble this time, however, lies not in faulty definition, but in tacit and illicit reference to evidence not stated in our example. Taken by itself, the statement that the given object is neither black nor a raven confirms the hypothesis that everything that is not a raven is not black as well as the hypothesis that everything that is not black is not a raven. We tend to ignore the former hypothesis because we know it to be false from abundant other evidence—from all the familiar things that are not ravens but are black. But we are required to assume that no such evidence is available. Under this circumstance, even a much stronger hypothesis is also obviously confirmed: that nothing is

either black or a raven. In the light of this confirmation of the hypothesis that there are no ravens, it is no longer surprising that under the artificial restrictions of the example, the hypothesis that all ravens are black is also confirmed. And the prospects for indoor ornithology vanish when we notice that under these same conditions, the contrary hypothesis that no ravens are black is equally well confirmed.⁷

On the other hand, our definition does err in not forcing us to take into account all the *stated* evidence. The unhappy results are readily illustrated. If two compatible evidence statements confirm two hypotheses, then naturally the conjunction of the evidence statements should confirm the conjunction of the hypotheses.⁸ Suppose our evidence consists of the statements E_1 saying that a given thing b is black, and E_2 saying that a second thing c is not black. By our present definition, E_1 confirms the hypothesis that everything is black, and E_2 the hypothesis that everything is non-black. The conjunction of these perfectly compatible evidence statements will then confirm the self-contradictory hypothesis that everything is both black and non-black. Simple as this anomaly is, it requires drastic modification of our definition. What given evidence confirms

⁷ An able and thorough exposition of this paragraph is given by Israel Scheffler in his *Anatomy of Inquiry*, New York, 1963, pp. 286-91.

⁸ The status of the conjunction condition is much like that of the consequence condition—see Note III.5. Although Carnap drops the conjunction condition also (p. 394), he adopts for different reasons the requirement we find needed above: that the total available evidence must always be taken into account (pp. 211-13).

is not what we arrive at by generalizing from separate items of it, but—roughly speaking—what we arrive at by generalizing from the total stated evidence. The central idea for an improved definition is that, within certain limitations, what is asserted to be true for the narrow universe of the evidence statements is confirmed for the whole universe of discourse. Thus if our evidence is E_1 and E_2 , neither the hypothesis that all things are black nor the hypothesis that all things are non-black is confirmed; for neither is true for the evidence-universe consisting of b and c . Of course, much more careful formulation is needed, since some statements that are true of the evidence-universe—such as that there is only one black thing—are obviously not confirmed for the whole universe. These matters are taken care of by the studied formal definition that Hempel develops on this basis; but we cannot and need not go into further detail here.

No one supposes that the task of confirmation-theory has been completed. But the few steps I have reviewed—chosen partly for their bearing on what is to follow—show how things move along once the problem of definition displaces the problem of justification. Important and long-unnoticed questions are brought to light and answered; and we are encouraged to expect that the many remaining questions will in time yield to similar treatment.

But our satisfaction is shortlived. New and serious trouble begins to appear.

4. *The New Riddle of Induction*

Confirmation of a hypothesis by an instance depends rather heavily upon features of the hypothesis other than

its syntactical form. That a given piece of copper conducts electricity increases the credibility of statements asserting that other pieces of copper conduct electricity, and thus confirms the hypothesis that all copper conducts electricity. But the fact that a given man now in this room is a third son does not increase the credibility of statements asserting that other men now in this room are third sons, and so does not confirm the hypothesis that all men now in this room are third sons. Yet in both cases our hypothesis is a generalization of the evidence statement. The difference is that in the former case the hypothesis is a *lawlike* statement; while in the latter case, the hypothesis is a merely contingent or accidental generality. Only a statement that is *lawlike*—regardless of its truth or falsity or its scientific importance—is capable of receiving confirmation from an instance of it; accidental statements are not. Plainly, then, we must look for a way of distinguishing lawlike from accidental statements.

So long as what seems to be needed is merely a way of excluding a few odd and unwanted cases that are inadvertently admitted by our definition of confirmation, the problem may not seem very hard or very pressing. We fully expect that minor defects will be found in our definition and that the necessary refinements will have to be worked out patiently one after another. But some further examples will show that our present difficulty is of a much graver kind.

Suppose that all emeralds examined before a certain time t are green.⁹ At time t , then, our observations support the

⁹ Although the example used is different, the argument to follow is substantially the same as that set forth in my note 'A Query on Confirmation', cited in Note I.16.

hypothesis that all emeralds are green; and this is in accord with our definition of confirmation. Our evidence statements assert that emerald *a* is green, that emerald *b* is green, and so on; and each confirms the general hypothesis that all emeralds are green. So far, so good.

Now let me introduce another predicate less familiar than "green". It is the predicate "grue" and it applies to all things examined before *t* just in case they are green but to other things just in case they are blue. Then at time *t* we have, for each evidence statement asserting that a given emerald is green, a parallel evidence statement asserting that that emerald is grue. And the statements that emerald *a* is grue, that emerald *b* is grue, and so on, will each confirm the general hypothesis that all emeralds are grue. Thus according to our definition, the prediction that all emeralds subsequently examined will be green and the prediction that all will be grue are alike confirmed by evidence statements describing the same observations. But if an emerald subsequently examined is grue, it is blue and hence not green. Thus although we are well aware which of the two incompatible predictions is genuinely confirmed, they are equally well confirmed according to our present definition. Moreover, it is clear that if we simply choose an appropriate predicate, then on the basis of these same observations we shall have equal confirmation, by our definition, for any prediction whatever about other emeralds—or indeed about anything else.¹⁰ As in our earlier example, only the predictions subsumed under law-

¹⁰ For instance, we shall have equal confirmation, by our present definition, for the prediction that roses subsequently examined will be blue. Let "emerose" apply just to emeralds examined before time *t*, and to roses examined later. Then all emeroses so far examined are grue, and this confirms the hypothesis that all

like hypotheses are genuinely confirmed; but we have no criterion as yet for determining lawlikeness. And now we see that without some such criterion, our definition not merely includes a few unwanted cases, but is so completely ineffectual that it virtually excludes nothing. We are left once again with the intolerable result that anything confirms anything. This difficulty cannot be set aside as an annoying detail to be taken care of in due course. It has to be met before our definition will work at all.

Nevertheless, the difficulty is often slighted because on the surface there seem to be easy ways of dealing with it. Sometimes, for example, the problem is thought to be much like the paradox of the ravens. We are here again, it is pointed out, making tacit and illegitimate use of information outside the stated evidence: the information, for example, that different samples of one material are usually alike in conductivity, and the information that different men in a lecture audience are usually not alike in the number of their older brothers. But while it is true that such information is being smuggled in, this does not by itself settle the matter as it settles the matter of the ravens. There the point was that when the smuggled information is forthrightly declared, its effect upon the confirmation of the hypothesis in question is immediately and properly registered by the definition we are using. On the other hand, if to our initial evidence we add statements concerning the conductivity of pieces of other materials or concerning the number of older brothers of members of

emeroses are grue and hence the prediction that roses subsequently examined will be blue. The problem raised by such antecedents has been little noticed, but is no easier to meet than that raised by similarly perverse consequents. See further IV, 4 below.

other lecture audiences, this will not in the least affect the confirmation, according to our definition, of the hypothesis concerning copper or of that concerning this lecture audience. Since our definition is insensitive to the bearing upon hypotheses of evidence so related to them, even when the evidence is fully declared, the difficulty about accidental hypotheses cannot be explained away on the ground that such evidence is being surreptitiously taken into account.

A more promising suggestion is to explain the matter in terms of the effect of this other evidence not directly upon the hypothesis in question but *indirectly* through other hypotheses that *are* confirmed, according to our definition, by such evidence. Our information about other materials does by our definition confirm such hypotheses as that all pieces of iron conduct electricity, that no pieces of rubber do, and so on; and these hypotheses, the explanation runs, impart to the hypothesis that all pieces of copper conduct electricity (and also to the hypothesis that none do) the character of lawlikeness—that is, amenability to confirmation by direct positive instances when found. On the other hand, our information about other lecture audiences *disconfirms* many hypotheses to the effect that all the men in one audience are third sons, or that none are; and this strips any character of lawlikeness from the hypothesis that all (or the hypothesis that none) of the men in *this* audience are third sons. But clearly if this course is to be followed, the circumstances under which hypotheses are thus related to one another will have to be precisely articulated.

The problem, then, is to define the relevant way in which such hypotheses must be alike. Evidence for the

hypothesis that all iron conducts electricity enhances the lawlikeness of the hypothesis that all zirconium conducts electricity, but does not similarly affect the hypothesis that all the objects on my desk conduct electricity. Wherein lies the difference? The first two hypotheses fall under the broader hypothesis—call it “*H*”—that every class of things of the same material is uniform in conductivity; the first and third fall only under some such hypothesis as—call it “*K*”—that every class of things that are either all of the same material or all on a desk is uniform in conductivity. Clearly the important difference here is that evidence for a statement affirming that one of the classes covered by *H* has the property in question increases the credibility of any statement affirming that another such class has this property; while nothing of the sort holds true with respect to *K*. But this is only to say that *H* is lawlike and *K* is not. We are faced anew with the very problem we are trying to solve: the problem of distinguishing between lawlike and accidental hypotheses.

The most popular way of attacking the problem takes its cue from the fact that accidental hypotheses seem typically to involve some spatial or temporal restriction, or reference to some particular individual. They seem to concern the people in some particular room, or the objects on some particular person’s desk; while lawlike hypotheses characteristically concern all ravens or all pieces of copper whatsoever. Complete generality is thus very often supposed to be a sufficient condition of lawlikeness; but to define this complete generality is by no means easy. Merely to require that the hypothesis contain no term naming, describing, or indicating a particular thing or location will obviously not be enough. The troublesome

hypothesis that all emeralds are grue contains no such term; and where such a term does occur, as in hypotheses about men in *this room*, it can be suppressed in favor of some predicate (short or long, new or old) that contains no such term but applies only to exactly the same things. One might think, then, of excluding not only hypotheses that actually contain terms for specific individuals but also all hypotheses that are equivalent to others that do contain such terms. But, as we have just seen, to exclude only hypotheses of which *all* equivalents contain such terms is to exclude nothing. On the other hand, to exclude all hypotheses that have *some* equivalent containing such a term is to exclude everything; for even the hypothesis

All grass is green

has as an equivalent

All grass in London or elsewhere is green.

The next step, therefore, has been to consider ruling out predicates of certain kinds. A syntactically universal hypothesis is lawlike, the proposal runs, if its predicates are 'purely qualitative' or 'non-positional'.¹¹ This will obviously accomplish nothing if a purely qualitative

¹¹ Carnap took this course in his paper 'On the Application of Inductive Logic', *Philosophy and Phenomenological Research*, vol. 8 (1947), pp. 133-47, which is in part a reply to my 'A Query on Confirmation', cited in Note I.16. The discussion was continued in my note 'On Infirmities of Confirmation Theory', *Philosophy and Phenomenological Research*, vol. 8 (1947), pp. 149-51; and in Carnap's 'Reply to Nelson Goodman', same journal, same volume, pp. 461-2.

predicate is then conceived either as one that is equivalent to some expression free of terms for specific individuals, or as one that is equivalent to no expression that contains such a term; for this only raises again the difficulties just pointed out. The claim appears to be rather that at least in the case of a simple enough predicate we can readily determine by direct inspection of its meaning whether or not it is purely qualitative. But even aside from obscurities in the notion of 'the meaning' of a predicate, this claim seems to me wrong. I simply do not know how to tell whether a predicate is qualitative or positional, except perhaps by completely begging the question at issue and asking whether the predicate is 'well-behaved'—that is, whether simple syntactically universal hypotheses applying it are lawlike.

This statement will not go unprotested. "Consider", it will be argued, "the predicates 'blue' and 'green' and the predicate 'grue' introduced earlier, and also the predicate 'bleen' that applies to emeralds examined before time *t* just in case they are blue and to other emeralds just in case they are green. Surely it is clear", the argument runs, "that the first two are purely qualitative and the second two are not; for the meaning of each of the latter two plainly involves reference to a specific temporal position." To this I reply that indeed I do recognize the first two as well-behaved predicates admissible in lawlike hypotheses, and the second two as ill-behaved predicates. But the argument that the former but not the latter are purely qualitative seems to me quite unsound. True enough, if we start with "blue" and "green", then "grue" and "bleen" will be explained in terms of "blue" and "green" and a temporal term. But equally truly, if we start with "grue"

and “bleen”, then “blue” and “green” will be explained in terms of “grue” and “bleen” and a temporal term; “green”, for example, applies to emeralds examined before time t just in case they are grue, and to other emeralds just in case they are bleen. Thus qualitativeness is an entirely relative matter and does not by itself establish any dichotomy of predicates. This relativity seems to be completely overlooked by those who contend that the qualitative character of a predicate is a criterion for its good behavior.

Of course, one may ask why we need worry about such unfamiliar predicates as “grue” or about accidental hypotheses in general, since we are unlikely to use them in making predictions. If our definition works for such hypotheses as are normally employed, isn't that all we need? In a sense, yes; but only in the sense that we need no definition, no theory of induction, and no philosophy of knowledge at all. We get along well enough without them in daily life and in scientific research. But if we seek a theory at all, we cannot excuse gross anomalies resulting from a proposed theory by pleading that we can avoid them in practice. The odd cases we have been considering are clinically pure cases that, though seldom encountered in practice, nevertheless display to best advantage the symptoms of a widespread and destructive malady.

We have so far neither any answer nor any promising clue to an answer to the question what distinguishes law-like or confirmable hypotheses from accidental or non-confirmable ones; and what may at first have seemed a minor technical difficulty has taken on the stature of a major obstacle to the development of a satisfactory theory

of confirmation. It is this problem that I call the new riddle of induction.

5. *The Pervasive Problem of Projection*

At the beginning of this lecture, I expressed the opinion that the problem of induction is still unsolved, but that the difficulties that face us today are not the old ones; and I have tried to outline the changes that have taken place. The problem of justifying induction has been displaced by the problem of defining confirmation, and our work upon this has left us with the residual problem of distinguishing between confirmable and non-confirmable hypotheses. One might say roughly that the first question was “Why does a positive instance of a hypothesis give any grounds for predicting further instances?”; that the newer question was “What is a positive instance of a hypothesis?”; and that the crucial remaining question is “What hypotheses are confirmed by their positive instances?”

The vast amount of effort expended on the problem of induction in modern times has thus altered our afflictions but hardly relieved them. The original difficulty about induction arose from the recognition that anything may follow upon anything. Then, in attempting to define confirmation in terms of the converse of the consequence relation, we found ourselves with the distressingly similar difficulty that our definition would make any statement confirm any other. And now, after modifying our definition drastically, we still get the old devastating result that any statement will confirm any statement. Until we find a way of exercising some control over the hypotheses to be

admitted, our definition makes no distinction whatsoever between valid and invalid inductive inferences.

The real inadequacy of Hume's account lay not in his descriptive approach but in the imprecision of his description. Regularities in experience, according to him, give rise to habits of expectation; and thus it is predictions conforming to past regularities that are normal or valid. But Hume overlooks the fact that some regularities do and some do not establish such habits; that predictions based on some regularities are valid while predictions based on other regularities are not. Every word you have heard me say has occurred prior to the final sentence of this lecture; but that does not, I hope, create any expectation that every word you will hear me say will be prior to that sentence. Again, consider our case of emeralds. All those examined before time t are green; and this leads us to expect, and confirms the prediction, that the next one will be green. But also, all those examined are grue; and this does not lead us to expect, and does not confirm the prediction, that the next one will be grue. Regularity in greenness confirms the prediction of further cases; regularity in grueness does not. 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. As we have seen, Hume's failure to recognize and deal with this problem has been shared even by his most recent successors.

As a result, what we have in current confirmation theory is a definition that is adequate for certain cases that so far can be described only as those for which it is adequate. The theory works where it works. A hypothesis is

confirmed by statements related to it in the prescribed way provided it is so confirmed. This is a good deal like having a theory that tells us that the area of a plane figure is one-half the base times the altitude, without telling us for what figures this holds. We must somehow find a way of distinguishing lawlike hypotheses, to which our definition of confirmation applies, from accidental hypotheses, to which it does not.

Today I have been speaking solely of the problem of induction, but what has been said applies equally to the more general problem of projection. As pointed out earlier, the problem of prediction from past to future cases is but a narrower version of the problem of projecting from any set of cases to others. We saw that a whole cluster of troublesome problems concerning dispositions and possibility can be reduced to this problem of projection. That is why the new riddle of induction, which is more broadly the problem of distinguishing between projectible and non-projectible hypotheses, is as important as it is exasperating.

Our failures teach us, I think, that lawlike or projectible hypotheses cannot be distinguished on any merely syntactical grounds or even on the ground that these hypotheses are somehow purely general in meaning. Our only hope lies in re-examining the problem once more and looking for some new approach. This will be my course in the final lecture.