

"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; "grue," 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 qualitiveness 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 the clinically pure cases that, though seldom encountered in practice, nevertheless dis- play 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 lawlike or confirmable hypotheses from accidental or nonconfirmable 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.

Other Approaches

Leonard J. Savage IMPLICATIONS OF PERSONAL PROBABILITY FOR INDUCTION*

INTRODUCTION

Statistical inference and philosophy evidently bear on each other. Exploration of their connections and common ground is accelerating but is not easy. Philosophers find the statistical literature dilute, discordant, philosophically unrigorous, and technical, and therefore hard to winnow. As a statistician driven toward philosophy by interest in the foundations of statistics, I find myself impeded by corresponding difficulties, unable to cover even the most pertinent chapters of philosophy or to determine which are pertinent.

Notwithstanding, the reference to implication in my title, I shall attempt no demonstrations here. Rather, I shall grope to share with you some possible insights into induction inspired by study of personal probability and statistics. Genuine demonstrations in philosophy seem rare or nonexistent, though philosophical discussion is often couched in pithy little logical-sound-

ing arguments and sometimes in even more treacherous long ones. How often is such seeming logic advanced with a conviction of rigor and how often as a sort of figure of speech hinting at something vague and insecure?

Three-line arguments utterly demolishing the concept of personal probability are widespread. Each has an even shorter and more devastating refutation, and so on. Such repartee can bear fruit, but only by slow growth on the soil of humility.

Some of the most untrustworthy of philosophical demonstrations have been among the most valuable. Warriors can overtake more to motion than meets the eye. And the importance of Hume's argument against induction—the keynote of this symposium—is undoubted; though perhaps most philosophers view it, like the arguments of Zeno, only as a challenge to search out manifest fallacy. Some of us, however, find Hume's conclusion not paradoxical but close to the mark.

The next section introduces personal probability, necessarily briefly. Though thorough discussion of personal probability cannot be an objective of this paper, a critical section will intensify its introduction, forestall unnecessary misunderstanding, and provide a natural setting for some remarks on induction. We can then look

*To be presented in an APA Symposium on Subjective Probability, December 27, 1967. Commentators will be J. Saver Minas (University of Waterloo, Ontario) and Ernest Nagel (Columbia University).

The research for this paper was supported by the Army, Navy, Air Force, and NASA under a contract administered by the Office of Naval Research.

Reprinted by permission of the estate of the author and *The Journal of Philosophy*, Vol. LXIV, no. 19 (October 5, 1967), pp. 593-607.

directly at the riddle of induction through the eyes of personal probability. Finally, several questions commonly associated with induction will be touched upon in a section on universal propositions.

PERSONAL PROBABILITY

The concept of personal probability was discovered several times between 1921 and 1940 and has older roots. Since about 1950, it has been known to statisticians and is having an increasing influence on them. For history, readings, and bibliography, see the anthology of Kyburg and Smokler.¹ Not all theories of personal probability are quite the same, and in presenting the concept to you, I shall attempt to portray scarcely any view but my own.

Personal probability can be regarded as part of a certain theory of coherent preference in the face of uncertainty. This preference theory is normative; its goal is to help us make better decisions by exposing to us possible incoherencies in our attitudes toward real and hypothetical alternatives.

If as a daily beverage I prefer water to wine and vinegar to water, you may disagree with me and even pity me, but my bizarre tastes are no ground for taxing me with incoherency. If, however, I go on to express a preference for wine as opposed to vinegar, this preference (however normal in itself) is absurd in the presence of my other preferences. If it is called to my attention, I would do well to review my expressed preferences and alter at least one of them.

Various systems of postulates, such as the postulate of transitivity of preference, though qualitative in approach, lead to an arithmetization of the value judgments and opinions of an ideally coherent person. In technical terms, such a person acts in the face of uncertainty so as to maximize the expected *utility* of his experiences with respect to his *personal probability measure* of the events that might affect those experiences.

The utility function of a person is a certain behaviorally defined expression of the value for him of the experiences to which it applies. The theory makes no attempt to brand some utility functions as more appropriate than others; *de gustibus non disputandum est*. This is not to say that drinking vinegar in preference to wine is normal but simply that there cannot be an objective right and wrong about such matters, as there is about an expression of intransitivity of preference.

As utilities express values or tastes, so a person's system of personal probabilities expresses his opinion in an arithmetic way. Nothing in the theory of personal probability precludes his believing that Elizabeth I wrote *Hamlet*. Though bizarre to you and me, in the light of what we and the person all know, this opinion need not be incoherent. Yet, being subject to a personal probability does impose much objective discipline on a person's opinions. He can believe that Elizabeth I probably wrote *Hamlet*; also that it will probably snow in Rio tomorrow. But then coherency will require him to consider that it will snow in Rio tomorrow more probable than that *Hamlet* was written by a commoner.

Though utility is no less fundamental to the preference theory than is personal probability, the latter is much more important for this paper and must therefore be more fully described.

The ideally coherent person, frequently called by the apt technical term 'you', is said to regard the event *A* as more probable than *B* under this condition: If you could receive a particular prize if and only if *A* obtains or else if and only if *B* obtains, then you would prefer the alternative that associates the prize with *A*.

Coherency seems to demand that, if *C* is incompatible with *A* and *B*, then the union of *A* and *C* will be more probable for you than the union of *B* and *C* if and only if *A* is more probable for you than *B*. If also there exist for you partitions of the universe into arbitrarily many equally probable events,

then there is necessarily a unique (finitely additive) probability measure so defined on all events that *A* is more probable for you than *B* if and only if the numerical probability of *A* exceeds the numerical probability of *B*. The partition assumption, which can be somewhat weakened² is not really an assumption of coherency but rather an assumption of a sufficient richness of contemplated events.

All currently active versions of the preference theory, explicitly or implicitly, exclude dependence of your personal probabilities on what the prize is. The personal probability measure *P* does vary with the person and with his initial body of knowledge, or data, but we need not here complicate the notation with an explicit indication of this dependence. It is, however, important to describe how opinion changes under the impact of new bits of knowledge such as that the event *D* obtains; so *conditional probability*, or *probability given D* is introduced.

You are said to hold *A* to be more probable than *B* given *D* under this condition: You would rather have a prize contingent on the intersection of *A* and *D* than the same prize contingent on the intersection of *B* and *D*. The situation is almost verbatim the same as before and, therefore, leads generally to a new probability measure on events *A*, dependent on the conditioning event *D*. If the initial probability *P(D)* is not 0, then the conditional probability of *A* given *D* is $P(A|D) = P(A \cap D)/P(D)$. This is not a mere convention, but an immediate deduction from the qualitative definition of conditional probability. Nor is the qualitative definition itself unmotivated, as will now be explained.

Suppose that you are to be allowed to associate the prize with *A* or *B* but may defer your choice until learning which element *D_i* of a partition (that is, a disjoint and exhaustive finite sequence of events) actually obtains. This amounts to making several simultaneous decisions, one for each *i*. According to one of the criteria of coherency, you must definitely prefer to

associate the prize with *A* in case *D_i* does obtain if and only if *A* is more probable than *B* given *D_i* in the qualitative sense. This not only clarifies the definition but illustrates how there is within the preference theory a natural interpretation (of at least one important sense) of the phrase 'learning by experience'.

If you are to choose one or another act in the light of which of the possible outcomes *D_i* occur in some experiment, then, planning now, you would agree to have your behavior after the experiment governed by your conditional probabilities, given whichever *D_i* actually obtains. Therefore, since you are coherent, these conditional probabilities will indeed be your effective probabilities when you have seen the outcome of the experiment. (This, incidentally, shows why, at any given moment, you must use the probabilities conditional on all that you have thus far learned, a point which has sometimes seemed puzzling.³) By elaborating, any sort of contingency planning can be represented in the preference theory. This is noteworthy: for the theory itself is atemporal and makes no scientific or philosophical commitments about time.

CRITICAL DISCUSSION OF PERSONAL PROBABILITY

My central claim for personal probability is that the preference theory, of which personal probability is an aspect, is a valuable framework for disciplining our behavior and attitudes in the face of uncertainty. How well the claim can be defended is of course open to debate and experience, and only after some discussion of that can we turn to the natural, but secondary, question of what personal probability has to do with probability.

Save through the criterion of coherency, the preference theory makes no distinction between right and wrong opinion. It does not censure the neighbor whom we find superstitious or paranoid nor recognize any

notion of the correct inference from data beyond what is implied by the definition and analysis of conditional probability. Some find in this open-mindedness a deadly objection against the theory. Two lines of reply suggest themselves. First, a theory that does some things well is not to be discarded merely for not doing everything. Second, a coherent person strongly but not absolutely rigidly convinced, for example, that 13 is a lucky number for him at roulette would reach a different opinion if he failed to win with exceptionally high frequency in a trial of many bets on that number. The theory does thus require holders of extremely diverse systems of opinion to agree closely with one another when presented with suitable common evidence. This alone seems to be an adequate model of the ostensible objectivity of scientific knowledge.

Sometimes theories of statistics based on a frequency concept of probability are called "objective" and the theory of statistics based on personal probability is called "subjective." This is natural, because the probabilities of a successful frequency theory would be objective, and personal probabilities are clearly subjective. But the employment of frequentist theories of statistics also involves subjective judgments, as is usually recognized by their proponents. In such theories, the subjective judgments are not fully under that orderly discipline, coherence, which is demanded by the preference theory. Thus arises a paradox: Some frequency enthusiasts disparage personalistic statistics for dealing in opinions rather than facts, though their own theory of statistics actually proves to be more subjective than the personalistic one and in fact virtually becomes the personalistic theory when certain criteria of coherence are recognized.¹

Holders of what I have called *necessary* views of probability hope, in effect, to improve upon the concept of personal probability by finding such strong rules governing the probability of one event (or proposition) in the light of another that there will be no room for personal dif-

ferences, given common knowledge. Should this program be possible, it could not but be welcome, but all of its proponents admit to being very short of their goal. No purported steps toward it seem valid to me, and it might even be "demonstrated" that none are possible. Attempts to construct necessary probability seem generally to be affected, explicitly or implicitly, by the dubious notion, promulgated in Wittgenstein's *Tractatus*, of atomic propositions as the natural irreducible propositions of which all others are disjunctions. Necessary theories are, apparently inevitably, based on notions of symmetry, such as that knowledge of each of two or more things is exactly the same in every relevant respect. However, the judgment that those attributes which distinguish the similar objects or events are irrelevant is really a subjective one, for which there has not been and, in my judgment, cannot be any valid objective prescription. Successful construction of necessary probability would, it seems to me, negate just what is most convincing in Hume's skepticism. Modern necessary theories, descended as they are from naive old notions of equally likely cases, are designed to escape certain well-known disasters, but succeed only in postponing them.

According to a frequent criticism, a person's probability for an event will be high if he desires that event and low if he does not, or just the opposite, depending on his temperament; so the theory of personal probability is thought to encourage the errors of optimism and pessimism. These are indeed errors, and it is an important psychological truth that we cannot protect ourselves against them merely by logical care, as in principle we can against outright fallacies. However, the preference theory by its very structure exhorts us to appraise the probabilities of events, apart from any actual consequences they may have for us, by considering only certain hypothetical consequences. The counsel of dispassionate comparison is built in, though of course men of flesh and blood will not always be able to follow it.

Revolving as it does around pleasure and pain, profit and loss, the preference theory is sometimes thought to be too mundane to guide pure science or idle curiosity. Should there indeed be a world of action and a separate world of the intellect and should the preference theory be a valid guide for the one, yet utterly inferior to some other guide for the other, then even its limited range of applicability would be vast in interest and importance; but this dualistic possibility is for me implausible on the face of it and not supported by the theories advanced in its name.

Of course, the goods and ills to which the theory refers need not be mundane, but may reflect the most heroic aspirations—or the most vile. The philosophical puzzle is how the theory can bear on situations in which any notion of motive seems inapplicable. For my part, though perhaps without justification, I can hardly imagine betting at even odds against *A* rather than for *A* (should such a choice be forced upon me) if from the point of view of pure science or idle curiosity I felt quite sure of *A*.

To illustrate with a sufficiently idle question: Was Caesar wearing a new toga when he was assassinated? I guess not; perhaps you do too, and reasons are easy to adduce. Correspondingly, an enforceable document entitling the bearer to ten dollars in case the toga was new would be worth about one dollar to me. The strong doubt and the preference to sell are for me inextricable. The notion of this hypothetical preference to sell does involve contractual propositions, which are philosophically puzzling. Such contractions seem essential to the whole theory, not only to its applications to idle curiosity. Whether that is bad and what to do about it are questions beyond my present depth.

Some personalists and necessarians are altogether immune to the criticism of worldliness because, for them, that *A* is more probable than *B* for Mr. Smith is an intuitive and unanalyzable notion. They have been justifiably suspected, however, of not knowing what they are talking about. What use

can their unanalyzable concept be for either the world of action or that of the intellect? This divergence between two kinds of personalists, though sharp when focused upon, has not been intensely disruptive or even prevented both kinds from residing in one head.

Interest without material interest is not without meaning, but, no matter how pure an investigator's science may be, he must come down to earth if he reflects how he should next spend his time, not to mention his or the government's money.

Does personal probability have any claim on the name "probability"? Does it alone have such a claim? These questions are, as I have said, relatively secondary, but we are bound to ask them sooner or later, and thinking about them increases familiarity with personal probability.

Personal probability does have many of the attributes suggested by "probability"; it is a probability measure (in the mathematical sense) that helps guide action. A natural and, in the presence of controversy, a generous supposition is that there are several kinds of probability, but I am unable to share that view.

In a trivial sense, there are indeed many kinds of probability because the mathematical properties of probability apply to many things having no connection, and only a formal parallelism, with any extra-mathematical notion of probability. For example, the distribution of the total mass of the furniture in this room would be a mathematical probability. But, speaking seriously of more than one valid interpretation of probability, we mean interpretations concerning uncertainty or indeterminable behavior. Carnap, for example, proposes to recognize both a necessary probability and a frequency probability. Others, while respecting the notion of personal probability, feel that a frequency probability is also meaningful and important. Though those who put forward pluralistic views normally seem to take for granted that the different kinds of probabilities have something to do with one another, they seem not to delve into the

relationships. Otherwise, they would, I suspect, find that one kind of probability subsumes the others as special cases. At any rate, that is close to the conclusion of radical personalists like me.

Each attempt to define probability is of course based on something genuine which any complete analysis must take into account. For example, probability calculations often are based on the judgment that certain events are equally probable; these situations evoke symmetry and, more recently, necessary definitions of probability. Certainly too, observation of a long part of an even longer sequence of events judged to be similar (in a sense to be made clear soon) leads us to evaluate the probability of each as yet untied event in the sequence as practically equal to the frequency of success, the success ratio, in the events that have been tried. Such situations evoke frequentist definitions of probability. The personalist understands both of these phenomena in terms of personal probability.

If, for example, a person judges in a certain situation that every card in the deck has the same probability of being drawn, then this probability must be $1/52$, the probability of a red card must be $1/2$, and so on. But an objective final criterion for such judgments of symmetry seems to be a will-o'-the-wisp. Each person must stand on his own two feet, judging for himself in the light of his other opinions and experiences.

Turning to frequentist views, they seem to us to involve serious circularities or lacunae. An excellent analysis of the situation that evokes ideas of frequency has been given by de Finetti (*op. cit.*), and though it is too long and mathematical to be fully repeated here, anyone seriously contesting that personal probability has within itself the tenable part of the notion of frequency probability ought to reply to de Finetti's analysis. Even certain fragments of it, which are brief and easy to state, are enlightening. A person accepts a sequence of events as similar in the spirit of the frequentist notion of probability, or exchangeable,

exactly when the probability for him that all of any given k of the events obtain does not depend on which k they are. In the presence of this symmetry judgment and certain other judgments, which can loosely be described as a moderately open mind about what the success ratio in any subset of events will be, the probability of any event of the sequence, given the outcomes of many others, must be nearly equal to the success ratio for the observed events. Once again, familiar calculations are based on a judgment of the symmetry, which is a special instance of the judgments called personal probabilities.

This analysis of frequentist notions can be carried further. After a large number of events in the exchangeable sequence has been observed, the coherent person who began with a moderately open mind will necessarily be rather sure that the success ratio already observed is close to all the success ratios to be observed in that sequence in the indefinite future.

To summarize my opinion, the foundational difficulties in the definition of personal probability are less than those of other attempts to define probability, and the truth behind other attempts to define probability is correctly expressible through the theory of personal probability.

IS INDUCTION RATIONAL?

In a learned, thorough, and plainspoken article, which has been invaluable to me in preparing this paper, Wesley Salmon⁷ has stated the riddle of induction, given its history, and explored the strengths and weaknesses of attempts to answer or escape it. He concludes that the riddle has not been answered, but refuses to despair and makes a stirring comparison with the riddle of the infinitesimal calculus, which was solved only after more than a century of resistance, with enormous benefit to mathematics and to the human mind.

The riddle of induction can be put thus: What rational basis is there for any of our beliefs about the unobserved?

The theory of personal probability touches on the domain of the riddle and can even be construed as giving a partial answer. The theory prescribes, presumably compellingly, exactly how a set of beliefs should change in the light of what is observed. It can help you say, "My opinions today are the rational consequence of what they were yesterday and of what I have seen since yesterday." In principle, yesterday's opinions can be traced to the day before, but even given a coherent demigod able to trace his present opinions back to those with which he was born and to what he has experienced since, the theory of personal probability does not pretend to say with what system of opinions he ought to have been born. It leaves him, just as Hume would say, without rational foundation for his beliefs of today.

Can there be any such foundation? The theory as such is silent, but I am led by study of it to doubt that there is a rational basis for what we believe about the unobserved. In fact, Hume's arguments, and modern variants of them such as Goodman's discussion of 'bilen' and 'grue', appeal to me as correct and realistic. That all my beliefs are but my personal opinions, no matter how well some of them may coincide with opinions of others, seems to me not a paradox but a truism. The grandiose image of a demigod tracing his beliefs back to the cradle only to find an impasse there seems a valid metaphor. If there is rational basis for beliefs going beyond mere coherence, then there are some specific opinions that a rational baby demigod must have. Put that way, the notion of any such basis seems to me quite counterintuitive.

We may understand better if we explore why the skeptical answer repels so many. One philosopher will ask, "Can you be sincere in saying that you do not know that you have two hands?" Indeed, I believe firmly enough in my two hands to stake my whole life on them against a trifling gain, as when I climb a ladder for a look around, and almost too firmly to imagine what evidence would convince me that they are not there. Does not so perfect a degree of belief deserve to

be called knowledge? No, not in the spirit of the riddle; for the question is not how firm or widespread the belief is but whether it is rational in such a sense that I would be irrational to believe otherwise.

Turn now to a different example. If the first twenty balls drawn from a box are black, is it not rational to believe that most of the balls in the box are black? Not at all, as housewives and other statisticians know. Often the only sound cherries in a box are on top; why should it not be so with the black balls? Perhaps that is unfair, because the original statement was elliptic and should be amplified. How? Perhaps by some reference to shuffling. A good counsel, but with too much practical and physical experience behind it to find any place in a fundamental definition of the rational. Perhaps we were supposed to understand that in each drawing every ball still in the box had an equal chance of being drawn. That does go part way toward justifying the conclusion (via exchangeable events), but the only interpretation I know for 'equal chance' is symmetry of opinion.

Even if each ball does have an equal chance, the conclusion still does not follow. For if, before the twenty drawings, you were strongly of the opinion that white balls were slightly in the majority, you might still quite plausibly think so; for example, even had the first twenty babies born in Shanghai last year been girls, I would remain rather firmly of the opinion that most babies born in Shanghai last year were boys. But suppose, as a counteskeptical might insist, that there is no information bearing on the contents of the box. What can it mean to be devoid of relevant information? The correct opinion of the completely uninformed mind brings us to a less grandiose but no more tractable version of the demigod in his crib and the dubious program of defining necessary probability.

To proceed not only without initial information but without any initial opinion, personal or public, is the slogan of most frequentist statisticians. Among the efforts under this slogan, at least two important

directions are distinguishable: fiducial probability, as in Fisher's volume of 1956⁶ and earlier; and inductive behavior as opposed to inductive inference, as initiated by Neyman in 1938.⁷ Neither direction seems successful to me. The first does apparently purport to answer the riddle in some cases, but the claim to rationality as opposed to mere objectivity eludes me and perhaps would not actually be pressed by its proponents. The second direction seeks to escape the riddle by emphasizing behavior as opposed to opinion, but the riddle can be asked as provocatively about behavior as it can about opinion. Actually, it is the followers of this direction who are even more subjective than we personalists.

If there is no rational basis for induction, why does induction work? As I have been trying to bring out by examples, no sharply defined method that works has actually been put forward. Rather there is a vaguely defined method, a general and indefinitely ramified art, that we all learn more or less well. To be sure, logic cannot be fully written down either: whether there is any analogy here, I can only ask, but surely it is far from perfect and does not make induction more rational.

If all coherent opinion must be equally respected, why is astronomy better than astrology? Should your whole outlook and the facts (which latter we may pretend to be common to us both) lead you to believe astrology to be effective, then reason alone cannot assail your position. Actually, practically all of us who disbelieve strongly in astrology do so for diffuse and subtle reasons; few have directly tried astrology and found it wanting or even reflected carefully on it. St. Augustine gave astrology some thought, perhaps overworking the pertinent fact that he and the son of his mother's slave were born at the same time. Jung's too was open-minded about astrology and made an empirical test of what he regarded as one of its predictions. The test proved difficult to appraise, but, in principle, though Jung and I differed greatly in outlook, we had

enough in common so that a better test might well have brought us to a common opinion about the aspect of astrology under test—negative, I would wager.

UNIVERSAL PROPOSITIONS

How do we know universals? Some philosophers seem to regard this as the first and simplest question about induction.⁸ But it seems to be an advanced and complicated one. For without it we have already been discussing induction, and universals differ greatly from one another in analysis and status.

Since I see no objective grounds for any specific belief beyond immediate experience, I see none for believing a universal other than one that is tautological, given what has been observed, as it is when it is a purely mathematical conclusion or when every possible instance has been observed. Reflection on the meaning and role of universals, especially through examples, will bring out some interesting points more or less pertinent to induction, though they can hardly be new to some of you.

Ironically, the most traditional universal, "All men are mortal", is by no means the least complicated, and one complicating feature makes the proposition particularly compelling. The counterinstance, "Smith is not mortal", is itself a universal and (in its ordinary mundane sense) in conflict not merely with what we believe about people but with what we believe about the solar system and the galaxy.

Consider more modest generalizations about the fragility of man. Has any man ever lived more than n years? For $n = 100$, surely; for $n = 150$, possibly; for $n = 200$, probably not; for $n = 500$, surely not. These responses are crude expressions of my own personal probability. Yours need not be the same. The example illustrates that what we would ordinarily call knowledge of a universal is acceptance with a high probability of a universal with finite domain or of many

such, vaguely specified. My opinion that no man has ever lived 500 years is of course justified in that arguments adducible for it would convince many, notwithstanding the reputation of Methuselah. But, as in any other induction, justification in the sense that one who rejects it is guilty of an error comparable to a logical fallacy is not available.

Will some man ever live 1,000 years? The respondent can hardly imagine all that he is being asked to contemplate. Try substituting for 'ever' the phrase 'within the next n years'. For modest values of n , speculation does not seem meaningless. Conceivably, though for me not probably, even the next few decades will bring such scientific control of aging that lifespans of 1,000 years for some men then alive will become plausible. The probability of such a technological advance within the next 100, 1,000, or 10,000 years is of course somewhat more probable. But even with rather high tolerance for hypothetical questions, one loses track of meanings if asked to speculate about humanity hundreds of thousands, let alone hundreds of millions, of years into the future.

Are all emeralds green? If by 'emerald' we mean a certain kind of green gem, then the answer is "yes," by convention. But perhaps geologists mean such a thing by 'emerald' that, for them, there are already nongreen emeralds, or at least the possibility of nongreen ones is not closed by convention.

What has just been said about the greenness of emeralds might be paraphrased for the blackness of crows, which suggests still other points. That it would be cheating to paint a crow nonblack helps to bring out once more what a vast amount is vaguely implicit in the simplest-sounding propositions of ordinary language. When we say that all crows are black, do we mean to imply that there are no albino crows or only that the classic downfall of the generalization about white swans is not about to repeat itself in some new Australia where whole

species of crows are white? For me, there may well be an occasional albino crow, and there might already, not to mention the distant future, be some island where all the crows are albino. Or, conceivably, some present or future biologist could show convincingly that albinism is not possible for any crow. The observation, by me and my neighbors, of millions of crows, all of them black, does not really go far toward establishing that all crows are black in any very wide sense, even granting that we would know a brown crow with a red breast if we saw one. Yet we may well be convinced of the blackness of the next thousand crows to be encountered by us in our usual haunts. (Stop press: *The Encyclopedia Britannica* mentions a species of gray crow in England.)

You can know that all the golf balls in a cigar box are white, in the practical everyday sense of knowledge, by looking into the open box and examining the half dozen balls that are there. Even this is not knowledge with a capital K. You cannot altogether escape the dangers of hallucination, fatigue, and tricks of light—a point to which I must return. The presence of several balls rather than one only is not really an interesting feature here; the 'all' in the example is trivial.

About 51 per cent of the babies born alive in Boston are boys. Such frequency generalizations seem closely related to universals, and many of the phenomena illustrated in connection with universals apply again. Thus, interpreted narrowly, the proposition might be merely a statement of a fact already observed; interpreted too broadly, it could involve science-fiction-like predictions that none of us would really venture. An idealized situation in which the proposition could be personally justifiable is this. The statement is taken to refer to some such finite class of live births regarded as exchangeable as first deliveries of young women in Boston in 1967, all live births in Boston in 1967, all live births in Boston in the past, all live births in Boston from 1,000 A.D. to 2,000 A.D., etc. The larger the class, the cruder will be the approximate

exchangeability. If, before making observations, you are not strongly opinionated against the immediate neighborhood of 51 per cent and if you find 51 per cent males in a sample of ten thousand or more, then coherency will require you to be pretty sure that the success ratio in the whole set (the population) is within less than 1 per cent of what it is in the observed set (the sample). In practice, exchangeability may only crudely approximate your opinion, though well enough to justify the conclusion.

A frequency generalization, like the universal propositions and theories of science, is supported not by just one line of evidence but by a network of evidence. For example, sex ratios vary more than can be reasonably accounted for by chance from one community and one circumstance to another, but they have always been close to 51 per cent. Therefore, far from having an initial prejudice against this frequency's applying to Boston, I happen to be prejudiced in favor of its vicinity.

Evidently, my attitude toward universals tends to be reductionist. I would analyze them away as elliptical and, often more or less deliberately, ambiguous statements of a variety of finite conjunctions. But I must confess a serious difficulty in this reductionist program. Any ordinary proposition of ostensibly particular form, such as "This ball is red", is intended and understood to imply many things not yet observed and indeed many propositions that would ordinarily be regarded as universals. For example, "This object will look red to me and others whenever examined by daylight for some ill-specified time to come". It has about the same diameter in every direction. "It will not soon change shape if left undisturbed", etc. We can attempt more cautious particular propositions, such as "I see white in the upper-left quadrant", hoping thus to avoid being deceived by appearances inasmuch as we report only appearances. But universals lurk even in such reports of sense data. The notions of "I", "upper", "right", and "white" all seem to

take their meanings from orderly experience. Indeed, I cannot imagine communication in the absence of expectation of continued order in domains as yet unperceived. To be sure, each universal implicit in an ostensible particular can itself be subjected to reductionist analysis like other universals, but the ideal of eliminating universals altogether seems impossible to me. We have come once more, but along a different path, to the place where personalists disagree with necessarians in expecting no solution to the problem of the *tabula rasa*.

NOTES

1. Henry E. Kyburg, Jr., and Howard E. Smokler, *Studies in Subjective Probability* (New York: Wiley, 1964). Includes an English translation of Bruno de Finetti, "La prevision: Ses lois logiques, ses sources subjectives," *Annales de l'Institut Henri Poincaré*, vii (1937): 1-68; and also a reprinting of my "The Foundations of Statistics Reconsidered," pp. 575-586 in *Proceedings of the Fourth [1960] Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, Jerzy Neyman, ed. (Berkeley: University of California Press, 1961). A few other works pertinent to personal probability are listed in my "Difficulties in the Theory of Personal Probability," *Philosophy of Science*, xxxiv (1967), to be published.
2. Leonard J. Savage, *The Foundations of Statistics* (New York: Wiley, 1954), chap. 4.
3. Alfred J. Ayer, "The Conception of Probability as a Logical Relation," in *The Problem of Knowledge* (New York: Penguin, 1956), pp. 67-73.
4. Leonard J. Savage, "The Foundations of Statistics Reconsidered," pp. 575-586 in *Proceedings of the Fourth [1960] Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, ed. Jerzy Neyman (Berkeley: University of California Press, 1961).
5. Wesley C. Salmon, "The Foundations of Scientific Inference," in *Mind and Cosmos: Essays in Contemporary Science and Philosophy*, vol. 1, University of Pittsburgh Series in the Philosophy of Science (Pittsburgh, Pa.: University Press, 1966).

6. Sir Ronald A. Fisher, *Statistical Methods and Scientific Inference* (New York: Hafner, 1956).
7. Jerzy Neyman, "L'estimation statistique, traitée comme un problème classique de probabilité," pp. 25-37 of *Actualités Scientifiques et Industrielles*, 739 (Paris: Hermann et Cie., 1938).
8. Carl G. Jung, "Synchronizität als ein Prinzip akausaler Zusammenhänge," *Naturerklärung und Psyche* (Studien aus dem C. G. Jung-Institut, IV; Zurich: Rascher, 1952). English translation:

9. Jean Nicod, "The Logical Problem of Induction" (orig. Paris, 1923), in *Foundations of Geometry and Induction* (New York: Harcourt, Brace & World, 1930).

Gilbert H. Harman

THE INFERENCE TO THE BEST EXPLANATION¹

I wish to argue that enumerative induction should not be considered a warranted form of nondeductive inference in its own right.² I claim that, in cases where it appears that a warranted inference is an instance of enumerative induction, the inference should be described as a special case of another sort of inference, which I shall call "the inference to the best explanation."

The form of my argument in the first part of this paper is as follows: I argue that even if one accepts enumerative induction as one form of nondeductive inference, one will have to allow for the existence of "the inference to the best explanation." Then I argue that all warranted inferences which may be described as instances of enumerative induction must also be described as instances of the inference to the best explanation.

So, on my view, either (a) enumerative induction is not always warranted or (b) enumerative induction is always warranted

but is an uninteresting special case of the more general inference to the best explanation. Whether my view should be expressed as (a) or (b) will depend upon a particular interpretation of "enumerative induction."

In the second part of this paper, I attempt to show how taking the inference to the best explanation (rather than enumerative induction) to be the basic form of nondeductive inference enables one to account for an interesting feature of our use of the word "know." This provides an additional reason for describing our inferences as instances of the inference to the best explanation rather than as instances of enumerative induction.

I

"The inference to the best explanation" corresponds approximately to what others have called "abduction," "the method of hypothesis," "hypothetic inference," "the method of elimination," "eliminative induction," and "theoretical inference." I prefer my own terminology because I believe that it

Reprinted by permission from the author and from *Philosophical Review*, Vol. 74, no. 1, January 1965.