CARL G. HEMPEL

TURNS IN THE EVOLUTION OF THE PROBLEM OF INDUCTION*

1. THE STANDARD CONCEPTION: INDUCTIVE "INFERENCE"

Since the days of Hume's skeptical doubt, philosophical conceptions of the problem of induction and of ways in which it might be properly solved or dissolved have undergone a series of striking metamorphoses.

In my paper, I propose to examine some of those turnings, which seem to me to raise particularly important questions about the nature of empirical knowledge and especially scientific knowledge.

Many, but by no means all, of the statements asserted by empirical science at a given time are accepted on the basis of previously established evidence sentences. Hume's skeptical doubt reflects the realization that most of those indirectly, or inferentially, accepted assertions rest on evidence that gives them no complete, no logically conclusive, support. This is, of course, the point of Hume's observation that even if we have examined many occurrences of *A* and have found them all to be accompanied by *B*, it is quite conceivable, or logically possible, that some future occurrence of *A* might not be accompanied by *B*. Nor, we might add, does our evidence guarantee that past or present occurrences of *A* that we have not observed were – or are – accompanied by *B*, let alone that all occurrences ever of *A* are, without exception, accompanied by *B*.

Yet, in our everyday pursuits as well as in scientific research we constantly rely on what I will call the method of inductive acceptance, or MIA for short: we adopt beliefs, or expectations, about empirical matters on logically incomplete evidence, and we even base our actions on such beliefs – to the point of staking our lives on some of them.

The *problem of induction* is usually understood as the question of what can be said in justification of this procedure.

Any attempt to answer that question requires, first of all, a clear characterization of the *method* of inductive acceptance, presumably in terms of rules that specify under what conditions a given hypothesis may be inductively inferred from, or inductively accepted on the basis of, a given body of evidence. Only if MIA has been characterized by reference to such rules can the question of justification be significantly raised.

Such rules have indeed been proposed in the literature. Consider first a familiar and simple type, which seems to me to reflect induction as seen by Hume and, of course, by many later thinkers:

To argue from 'All examined instances of A have been B' to 'All A are B'.

This formulation seems quite plausible; but it is fundamentally defective: it does not make clear at all just what claim is being made for this rule, or what it means to say that this is a valid rule of inductive reasoning. For a rule of *deductive* reasoning, such as *modus tollens*, the claim is clear: if the premisses to which the rule is applied are true, then invariably so will be the conclusion.

But this claim of deductive validity is too strong, of course, for a rule of induction. Suppose, then, we were to read the rule as saying that under the specified conditions, it is rational to *accept* the conclusion, or perhaps: to act *as if* the conclusion were known to be true.

Now, the notion of accepting a hypothesis is surely in need of clarification: this is an issue to which I will address myself later. But even without entering into the details, it can be shown that on this construal our rule is untenable because it would oblige us to accept logically incompatible hypotheses on one and the same body of evidence.

Suppose, for example, that we have measured the length of a given silver bar at different temperatures and now plot the associated values as data points in a rectangular coordinate system. We then can draw different curves through the data points, representing l as different functions, $f_1, f_2, \ldots,$ of T.

Now, our rule entitles us to argue from:

All examined associated values of T and l satisfy f_1
All examined associated values of T and l satisfy f_2

to: It is reasonable to accept the hypothesis that *all* associated values of T, l satisfy f_1

and It is reasonable to accept the hypothesis that *all* associated values of T, l satisfy f_2

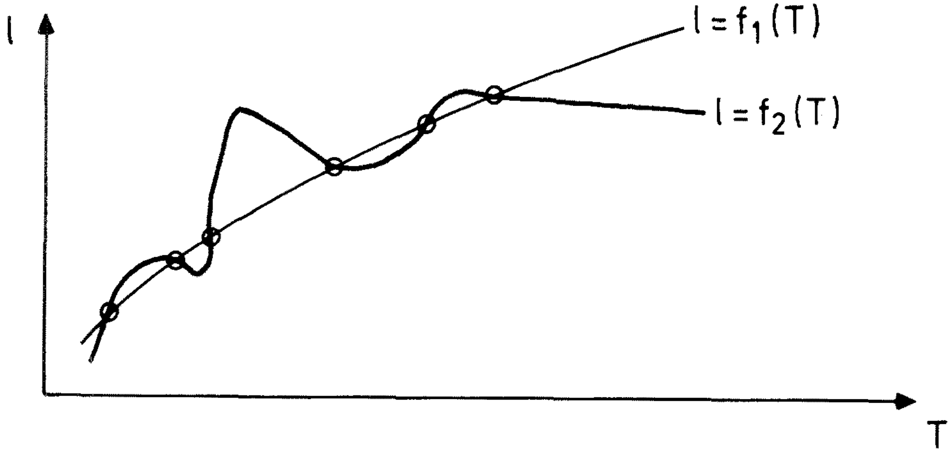

But the two general hypotheses thus inductively inferred are logically incompatible.

As this example illustrates, one and the same body of evidence can be described in different ways which, via our rule of induction, yield logically incompatible generalizations and also logically incompatible hypotheses concerning particular past, present, or future cases.

This point is also illustrated by Goodman's "New riddle of induction", exemplified by his "grue-bleen" paradox. That paradox raises yet another issue, however, which I will have to pass by here.

2. PROBABILISTIC CONSTRUAL OF INDUCTIVE REASONING

A way out of the difficulty illustrated by my example is suggested by a different construal of the basic form of inductive reasoning – a construal that has, indeed, come into wide acceptance. According to it, rules of inductive inference should not be conceived of as assigning certain specific "inductive conclusions" to a given body of evidence, but rather as principles which require that both the evidence and a specific contemplated hypothesis be *given*, and which then assign a certain probability or "rational credibility" to the given hypothesis relative to the given evidence.

A rule of inductive inference, thus construed, does not entitle us, as did our earlier rule, to argue from a given evidence sentence e to a hypothesis h specified by the rule. Rather, the rule now contemplated presupposes that both e and h are given and then specifies the probability of h on e:

$$p(h, e) = r$$

On this understanding, one body of evidence – e.g. a set of meteorological data available today – can assign certain probabilities to each of two incompatible hypotheses, e.g., 0.75 to 'Rain tomorrow', 0.25 to 'No rain tomorrow'. But no contradiction arises since our rule does not entitle us to assert either of the hypotheses, let alone both.

Indeed, in contrast to the case of deductive reasoning, there is no rule of detachment for inductive arguments which would entitle us, given that e is true, to detach, and assert without reference to e, an inductive conclusion such as 'probably h', or 'h holds with probability r'.

As these observations show, the new, probabilistic, construal of inductive inference avoids the earlier contradictions at a price: the new construal provides us with no principles of inductive acceptance or belief concerning empirical hypotheses.

But in all our practical and theoretical pursuits, we must eventually decide which, if any, of a set of alternative hypotheses under consideration we should accept as a basis for our expectations and actions.

3. A NEW TURN: TWO TYPES OF RULES OF INDUCTION

At this point, the problem of formulating rules of induction seems to split into two distinct issues:

(i) the problem of formulating *rules determining* the *probabilities* to be assigned to hypotheses on the basis of given evidence;

(ii) the problem of formulating *rules of acceptance* which would determine what hypotheses to accept, on given evidence, as a basis for our expectations and actions.

The first of these represents what Carnap considers as the central task of inductive logic. And indeed, Carnap's *Logical Foundations of Probability* offers an impressive formal theory of the probability of hypotheses on given evidence, for sentences expressible in a formalized language with the structure of first-order logic. Carnap and others subsequently changed and extended the approach taken in that work, and certain other conceptions of inductive probability have been developed as well; but it is certainly an open problem whether an adequate general explication of inductive probability can be given at all.

But let us set this question aside and assume for the moment that a suitable solution of the first problem is available. Then we are still left with the second task, that of formulating *rules of acceptance* for hypotheses on given evidence.

At first glance, it may seem plausible to think, and it has in fact often been suggested, that the acceptability of a hypothesis might be defined in terms of its probability, by a rule of this type: a hypothesis is to be accepted, on the relevant evidence available, just in case its probability on that evidence is greater than 1/2, or greater than some other fixed value, say 0.99.

But this rule runs afoul of the so-called lottery paradox, which shows that on certain types of evidence, several logically incompatible hypotheses would have to be adopted – just as in the case of the simple induction rule which we considered at the beginning.

This difficulty can be overcome in light of a very fruitful idea which was first developed precisely in the context of mathematical decision theory.

4. INDUCTION AND VALUATION

Briefly, the idea is that adequate criteria for the rational acceptability of a hypothesis must take account not only of the relevant evidence available and the support it lends to the hypothesis, but also of the values attached to avoiding the mistakes of accepting the hypothesis when it is, in fact, false; or of rejecting it when it is true. This idea gives a new turn to the conception of induction by attributing to it both a cognitive and a valuational component.

This point was emphasized forcefully in a provocative article by Richard Rudner, published in 1953 in *Philosophy of Science* under the title 'The scientist *qua* scientist makes value judgments', which attracted a great deal of attention and generated a fruitful controversy. Briefly, Rudner argues as follows: Scientists, in the course of their research, do again and again accept or reject hypotheses. But the evidence their decisions are based on is generally incomplete; therefore, in accepting a hypothesis the scientist must make the decision that the evidence is *sufficiently* strong or that the probability is *sufficiently* high to warrant the acceptance of the hypothesis. Obviously our decision ... respecting how strong is 'strong enough', is going to be a function of the *importance*, in the typically ethical sense, of making a mistake in accepting or rejecting the hypothesis.

Rudner illustrates the point by an example, which I will modify slightly for convenience in later reference. Suppose that a pharmaceutical firm has produced a large quantity of tablets, intended for the treatment of a certain disease. In too large amounts, the active ingredient is toxic; in too small amounts, ineffective. The hypothesis to be considered before releasing the whole batch for sale is to the effect that the amount of the active ingredient in each tablet falls within an interval that makes the drug both safe and effective. The evidence consists of findings obtained by chemical analysis of a random sample drawn from the entire production. For acceptance of the hypothesis that the pills are safe and effective, we would then require a very high degree of evidential support because the consequences of a mistake could be extremely grave by our moral standards.

If, by contrast, the hypothesis were to the effect that a given large quantity of machine-stamped belt buckles are non-defective – i.e. fall within a specified range in regard to certain required characteristics – a mistake would not be nearly so serious, and a considerably lower level of evidential support would suffice for the acceptability of the hypothesis.

In sum, then, Rudner argues that the inductive acceptability of a hypothesis depends on two factors: the relevant evidence available and value judgments about the importance of avoiding mistakes.

Rudner's argument suffers from a serious ambiguity, however. This was pointed out by Richard Jeffrey in a reply, 'Valuation and acceptance of scientific hypotheses' (*Philosophy of Science*, 1956). Jeffrey notes that, as Rudner's examples show, Rudner is concerned, not with the acceptance of a given hypothesis *tout court*, but rather with the adoption of a practical course of action based on the hypothesis, such as releasing the tablets for the treatment of human patients, or releasing the belt buckles for sale to belt makers. But, Jeffrey notes, the scientist *qua* scientist is not concerned with giving advice, or making decisions, on contemplated courses of practical action.

Furthermore, Jeffrey argues, on Rudner's view of the matter there can be no *one* level of evidential support that is high enough for the acceptance of a given hypothesis, since one and the same hypothesis might be put to use in different courses of action, for which the moral seriousness of a mistake may differ considerably. For example, the tablets might be used for treating a sick child or a sick household pet; and it seems clear that accepting, and acting on, the hypothesis of safety and effectiveness would, by our moral standards, require a higher level of evidential support in the former case than in the latter. There is no *one* degree of moral importance that attaches to the avoidance of mistakes in all possible practical applications of a given hypothesis.

In the spirit of these considerations, Jeffrey holds – along with a number of other leading thinkers in the field – that the notion of acceptance is not properly applicable to empirical hypotheses at all, but only to courses of action – specific practical applications of a given hypothesis, as we might say.

Jeffrey concludes: "it is not the business of the scientist as such... to accept or reject hypotheses"; rather, "the scientist's proper role is to provide the rational agents in the society which he represents with probabilities for the hypotheses which, on [Rudner's] account he simply accepts or rejects."

This is a startling conception indeed. In the context of the scientific pursuit of pure, unapplied, knowledge of the world, it entirely eliminates the classical problem of induction: if the pure scientist never accepts or rejects any hypotheses, then the problem of formulating rules of inductive acceptance, and the further problem of justifying those rules, simply does not arise.

There arises, however, a proxy problem, namely the question of rules for the assignment of probabilities to proposed hypotheses, and the question of how to justify those rules. But for the moment, let us leave this issue on one side.

Jeffrey himself is quick to point out that his conception faces serious difficulties of its own; among them, the following:

(i) it presupposes that an adequate theory can be provided for the assignment of probabilities to hypotheses on given evidence; this is part of the proxy problem just mentioned, and Jeffrey seems to regard it as dubious.

(ii) the role here assigned to the scientist "bears no resemblance to our ordinary conception of science. Books on electrodynamics, for example, simply list Maxwell's equations as laws; they do not add a degree of confirmation."

(iii) Rudner had anticipatively stressed that on a view of the type later proposed by Jeffrey, the pure scientist *still* has to accept certain hypotheses, namely those of the type 'the probability, or the evidential support, of hypothesis h on the available evidence in p'. Jeffrey suggests that it is not the task of the scientist to accept even hypotheses of this type, but he acknowledges that Rudner's objection is one of the weightiest difficulties for a probabilistic view of science.

(iv) Finally, I would add one further difficulty: Even if the scientist limits himself to determining probabilities for hypotheses, he must perform tests to obtain the evidence on the basis of which to calculate those probabilities. He must, therefore, it seems, accept certain empirical statements after all, namely the evidence sentences by which he judges the probability of contemplated hypotheses.

5. INDUCTION AND EPISTEMIC UTILITY

In sum, then, neither Rudner's nor Jeffrey's account of the matter can be quite right. But is seems to me that there is indeed something right about Rudner's claim concerning the relevance of valuation to scientific inquiry; therefore, I would like to consider the possibility of modifying Rudner's idea a bit in an effort to overcome some of the difficulties it faces.

Let me outline briefly one decision-theoretical approach to problems of the kind considered by Rudner concerning the acceptance or the rejection of a hypothesis. Take the pharmaceutical case. Suppose a random sample of the total production of tablets has been tested and that, on the evidence E thus obtained, our theory of evidential support assigns to the hypothesis H that all tablets in the whole batch are safe and effective the probability $p(H, E) = 0.75$.

A rational decision as to whether to accept or reject H (or rather, the corresponding course of action) will then depend on the importance, or the value, attached to achieving or to avoiding certain possible outcomes of our action.

The desirability or undesirability attached to different possible outcomes are often assumed to be expressible numerically as positive or negative "utilities". In our example, the relevant utilities might be as follows:

	H is true	H is false
Accept H	+200	−1000
Reject H	−60	+300

Then, given the probabilities $p(H, E) = 0.75$, $p(-H, E) = 0.25$, each of the possible outcomes can be assigned an expectable utility:

$$\bar{U}(\text{Acc } H, E) = (0.75 \times 200) - (0.25 \times 1000) = 150 - 250 = -150$$

$$\bar{U}(\text{Rej } H, E) = (0.75 \times -60) + (0.25 \times 300) = -45 + 75 = +30$$

And one often invoked rule for rational decision directs us to reject H, as the course of action which offers the greater expectable utility.

In case no probabilities are available, and in particular if the entire idea of assigning probabilities to hypotheses should be ruled out, decision theory still offers rules for what is called "decision-making under uncertainty". One such rule directs us to choose the course of action for which the minimum payoff (utility) is maximal; in our example, this maximin rule would again call for rejection of H, since $-60 > -1000$.

Now, it seems to me that despite Jeffrey's strictures, there is a tenable, and indeed important, point in Rudner's view of the matter. I think that the acceptance of a hypothesis in pure or basic science may be construed as an action, too, though not as an action of the practical kind considered in our previous example.

The action consists in including the given hypothesis into the corpus, K, of previously accepted hypotheses; and its purpose and potential value lies, not in solving any practical or technological problems, but in the increase of scientific knowledge. Since science aims at establishing true hypotheses, the addition of a hypothesis to the corpus of accepted statements might plausibly be assigned a positive utility in case the hypothesis is true; otherwise a negative utility of the same numerical magnitude. That magnitude itself could plausibly be taken to be given by the proportion of informational content that the hypothesis would add to the content of K.

I sketched this idea as a somewhat incidental point in an article published in 1962 ('Deductive-nomological *vs.* statistical explanation'; *Minnesota Studies in the Philosophy of Science*, vol. III; section 12), in which I proposed the term 'epistemic utility' to refer to the utility which the acceptance or rejection of a new hypothesis possesses for the purposes of pure, basic research.

By means of this concept and a suitable theory of probability for hypotheses, one could then formulate *a rule of inductive acceptance for pure or basic research*: in deciding whether to accept or to reject a given hypothesis or to leave it in suspense, choose the course of action which has the highest expectable epistemic utility attached to it. The idea is strictly analogous to that for practical actions, but with epistemic utilities taking the place of the utilities of practical consequences. Thus, on this construal, *the scientist qua scientist does indeed make value judgments*, but they are not of a *moral* kind; rather, they reflect the value to basic research of constructing sound and information-rich accounts of the world; and these are what I would call *epistemic* values.

Since it is often said that science *presupposes* value judgment, let me stress that epistemic judgments of value do not enter into the *content* of scientific hypotheses or theories; Kepler's laws, for example, do not presuppose or imply any value judgments at all – either epistemic or of other kinds. But epistemic valuation does enter into the *acceptance* of hypotheses or theories in this sense: the assertion that a given hypothesis H is *acceptable* in a given knowledge situation implies that the acceptance of H possesses a greater expectable epistemic value for science than does the acceptance of any rival hypothesis that may be under consideration.

That valuational considerations should play an essential role here is hardly surprising: how could a procedure like the adoption of a hypothesis be qualified as appropriate or rational except in consideration of the objectives of scientific inquiry, i.e. in consideration of the contribution that the adoption of the hypothesis is likely to make to furthering the objectives of inquiry? And that contribution is expressed in the epistemic utilities or values assigned to the possible scientific consequences of the adoption.

On the simple construal I sketched a moment ago, the epistemic value of accepting a hypothesis would depend only on its truth value and on the proportion of new information it adds to the set K of previously accepted hypotheses. But, as I showed in my article, that interpretation, combined with the rule of maximizing expectable utility, yields an inductive acceptance rule that is intuitively unreasonable; indeed, I should add, it leads again to the lottery paradox.

Yet, I believe that there is something fundamentally right about the idea of epistemic value, and that the failure of the utility measure just considered may be attributable to a too narrow construal of the objectives of basic research. Science is interested not only in questions of truth and informational content, but also in the simplicity of the total system of accepted hypotheses, in its explanatory and predictive powers, and other factors, all of which a theory of inductive acceptance would have to take into account.

6. THE PRAGMATIST CONSTRUAL OF THEORY CHOICE IN SCIENCE

Just those factors have in fact been given considerable prominence in what I think may be viewed as the latest major turn in the evolution of the problem of induction. This turn is one important aspect of the historic-pragmatist construal of scientific inquiry of which Thomas Kuhn is the leading proponent.

The problem of induction arises here in an especially comprehensive and momentous context, namely in regard to scientific revolutions, in which an established theory is eventually abandoned in favor of a new rival. In his account of the kind of choice between rival theories that occurs in this context, Kuhn considers the question whether there are general precise criteria of preferability for theories that determine the outcome of the choice.

To this end, he surveys a variety of characteristics that are widely acknowledged as desirable characteristics of scientific hypotheses and theories; let me call them *desiderata* for short. Kuhn holds that theory choice in science is definitely influenced by a comparison of rival theories in regard to the extent to which they fulfill such desiderata; and he further discusses the question whether those desiderata admit of precise objective definitions which might make it possible to construct explicit criteria of preferability (we might say: epistemic utility) that are objective in the sense that different scientists, provided with the same relevant information, would arrive at the same decision as to which of two competing theories, if any, is to be accepted.

One familiar desideratum is wide scope of application: it is highly desirable that a theory should account for many quite different kinds of phenomena. Newton's theory is a good example: it covers phenomena as diverse as free fall, the swinging of pendulums, the motions of planets, comets, double stars, artificial satellites, and so on.

But what exactly should be taken as the principle of individuation for *kinds* of phenomena? Why, for example, should all instances of free fall on the earth count as one kind in this context: why should not the fall of metal bodies count as a kind of event different from the fall of bodies of wood or of glass? Perhaps it might be said that irrespective of their composition, physical bodies fall in accordance with one and the same law, which can be (approximatively) derived from Newton's laws. But then, all the diverse kinds of events that Newton's theory as a whole accounts for behave in accordance with the basic laws of that theory: so, should they not count as belonging to one and the same kind of event; and where would this leave the putative diversity of the phenomena falling within the range of Newton's theory? For these and other reasons, it is highly questionable whether a satisfactory precise explication of the notion of the range or the scope of a theory can be constructed.

A related desideratum recently proposed in Laurence Laudan's book, *Progress and its Problems*, is that a theory should solve a large number of important problems. This idea is certainly plausible and reflects scientific preference, but there are no clear, unambiguous ways of differentiating and counting the problems solved by a theory, and of assessing their importance; the reasons are closely related to those just considered.

Compatibility, or more loosely, "fitting together", with well-established theories in neighboring fields also counts as a desideratum. For example, a hypothesis sometimes suggested for telepathic phenomena construes these as brought about by a special kind of radiation emitted by the sender of a telepathic message; but this idea has been objected to on the ground that it does not fit together with a certain characteristic of radiation theories in physics, namely, that the energy transmitted by the energy source to a given receiver decreases with the square of the distance of the receiver, whereas according to evidence claimed to describe telepathic communication, no such decrease in strength or clarity of the message is generally observed.

But again, no clear criterion is in sight for this notion of "fitting together" with well-established neighboring theories – a notion which plainly is stronger than just logical compatibility.

Similar difficulties face the attempt to formulate precise and general explications of concepts invoked in other desiderata affecting theory choice, among them the simplicity of a theory, the closeness of fit between experimental data and the corresponding implications of a theory, and the power of a theory to predict novel phenomena.

Thus, Kuhn holds, there are no precise criteria for the comparison of the merits of competing theories in regard to any of these desiderata. And even if such sharp criteria were available, there would remain the problem of combining them all into one precise overall criterion which would determine which, if either, of two competing theories to accept. But while scientists do agree about the importance of the various desiderata, there is no unanimity at all as to the relative weights that are to be assigned to them in the overall comparison of theories.

The prospects, then, of formulating precise explications of those desiderata that might provide an objective characterization of epistemic utility seem very dim indeed.

In fact, Kuhn's view implies a radical rejection of that conception of the problem of induction which calls, first of all, for the formulation of precise rules of acceptance and, secondly, for a justification of those rules.

Kuhn offers instead a fundamentally different, historic-pragmatist conception of scientific theory choice. He holds, to put it briefly, that the choice between competing theories lies in the hands of the specialists in the field. These specialists are all committed to what Kuhn calls the values of simplicity, accuracy, large scope, etc. in theory choice. And even though they do not share precise criteria of application for those concepts, and even though their preference is determined in part also by idiosyncratic factors, there does in fact, in the course of a controversy over the merits of two competing theories, eventually emerge a consensus leading to the acceptance of one of the competitors.

And though no explicit rules of acceptance are countenanced by Kuhn, he nevertheless offers what amounts to a justification for theory choice thus construed as resulting from the efforts of a group of professional specialists. Kuhn argues here that scientific research behavior as exemplified in theory choice serves an essential function in improving scientific knowledge, and that in the absence of an alternate mode of behavior that would serve the same function, "scientists should behave essentially as they do if their concern is to improve scientific knowledge"('Reflections on my critics', in I. Lakatos and A. Musgrave (eds.), *Criticism and the Growth of Knowledge*, 1970, p. 237). Kuhn regards this consideration also as an argument in favor of the claim that scientific behavior is rational, and is indeed the best example we have of rationality.

I think that this conception faces several difficulties:

First, it is clear that Kuhn's pragmatist account does not specify what count as the essential aspects of scientific behavior, those which the aspiring scientist should emulate if he wants to improve scientific knowledge.

Second, the account offers no specific indication of what counts as "improving scientific knowledge".

Third, on Kuhn's view, scientific theory choice (and other facets of scientific behavior) are not effected by means of procedures that are deliberately adopted by the scientific community as a presumably optimal means for advancing scientific knowledge; and I would think that any kind of action, including scientific theory choice, can be called *rational* only if it can be causally linked to deliberation or reasoning aimed at achieving specific ends. Scientific theory choice as characterized by Kuhn would not be rational in this sense, but would rather be akin to what in anthropology are called *latently functional* behavior patterns, which serve a function they were never *chosen* to fulfill.

For these reasons I think that Kuhn's pragmatist account of theory choice is not entirely satisfactory. However, I consider it as an important and illuminating corrective for an approach, perhaps most elaborately developed by analytic empiricism, which seeks to "explicate" scientific procedures, including induction, as governed by explicit and precise rules. For it does seem highly unlikely that the more comprehensive modes of induction, especially those involved in theory change, can be characterized by general rules that would be acknowledged as binding by scientists and would be observed by them.

On the other hand, for certain narrow and specific inductive problems, such as measuring a quantity, or testing a statistical hypothesis, fairly precise rules can be stated that are acknowledged and observed in scientific practice.

7. SOME REMARKS ON JUSTIFICATION

And how might such *explicit* rules of inductive acceptance be *justified*? Let me refer here to the ideas of a thinker who would generally be thought an analytic philosopher, but whose views concerning justification show a suggestive and illuminating kinship to the pragmatist approach – namely, Nelson Goodman. To quote a passage from his *Fact, Fiction, and Forecast* (1955; p. 67):

> ...deductive inferences are justified by their conformity to valid general rules, and ... 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.

It seems to me that, somewhat surprisingly, this view of the justification of inductive procedures has a significant kinship both with the analytically oriented ideas of Carnap and with the pragmatist account of Kuhn and kindred thinkers.

Carnap holds that the reasons to be given in support, or justification, of the basic principles of inductive reasoning "are based upon our intuitive judgments concerning inductive validity" (P. A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*, p. 978) and surely, there is a close relationship between this idea and that of considering what particular inductive arguments we are willing to give up, and what general rules we are willing to amend.

Kuhn leaves the ultimate decisions not to "our" intuitions, to "our" dispositions to accept or amend, but to those of the specialists professionally concerned and equipped to make decisions about acceptability.

Goodman's construal of justification, however, as the quoted passage shows, pertains only to modes of reasoning that are governed by explicitly stated rules; it cannot be straightforwardly extended to scientific theory choice as characterized by Kuhn.

As for that procedure, we briefly considered a different justification, offered by Kuhn. It argues that theory choice as practiced in science serves to improve scientific knowledge and that in the absence of a better alternative scientists should go about it essentially in the manner they do, if their aim is the improvement of scientific knowledge.

We noted some difficulties facing this idea, among them the point that no independent characterization had been provided of the notion of "improving scientific knowledge", so that it remained unclear just how to assess the effectiveness of scientific or alternate modes of theory choice in the pursuit of improved knowledge.

But this particular difficulty can be avoided by a change of perspective which does make it possible to attribute rationality to theory choice effected by reference to desiderata. The imposition of desiderata may be regarded, at least schematically, as the use of a set of means aimed at the improvement of scientific knowledge. But instead of viewing such improvement as a research goal that must be characterizable independently of the desiderata, we might plausibly conceive the goal of scientific inquiry to *be* the development of theories that ever better satisfy the desiderata. On this construal, the desiderata are different constituents of the goal of science rather than conceptually independent means for its attainment, and it becomes a truism that replacing a theory by a competing one that better satisfies the desiderata will constitute an improvement of scientific knowledge and will thus be a rational procedure.

This schematic construal in no way disposes of the vagueness, stressed earlier, of the desiderata and their relative weights, and it offers no prospect of a precise explication, in the sense of analytic empiricism, of scientific theory choice. But vagueness need not prevent a set of concepts from being epistemologically illuminating in certain contexts. The idea here outlined shows, I think, that if scientific inquiry aims at theories which ever better satisfy certain desiderata, then no matter how the latter may be construed in detail, the choice among theories by reference to the chosen desiderata is, in its basic structure, a rational enterprise.

University of Pittsburgh

NOTE

* This material is based upon work supported by the National Science Foundation under Grant No. SES-8025399.