Part I

Cognitive decision making

1

Must the scientist make value judgements?*†

> The scientific man has above all things to strive at self-elimination in his judgments, to provide an argument which is true for each individual mind as for his own.
>
> KARL PEARSON

Two assumptions implicit in Pearson's characterization of 'the scientific man' have been called into question in recent years: (*a*) At least one major goal of the scientist *qua* scientist is to make judgements – i.e., to accept or reject hypotheses – and to justify his judgements. (*b*) The scientific inquirer is prohibited by the canons of scientific inference from taking his attitudes, preferences, temperament, and values into account when assessing the correctness of his inferences.

One currently held view affirms (*a*) but denies (*b*). This position maintains that the scientist does and, indeed, must make value judgements when choosing between hypotheses. The other position upholds the value-neutrality thesis (*b*) at the expense of the claim that scientific inference issues in the acceptance and rejection of hypotheses (*a*). According to this view, a scientific inquiry does not terminate with the replacement of doubt by belief but with the assignment of probabilities or degrees of confirmation to hypotheses relative to the available evidence.

In this paper, a critical examination of these conflicting conceptions of scientific inference will be undertaken; the *prima facie* tenability of the claim that scientists can, do, and ought to accept and reject hypotheses in accordance with the value-neutrality thesis will be defended; and some indication will be given of the kind of question that must be answered before this plausible view can be converted into a coherent and adequate theory of the relation of values to scientific inference.

1

The tenability of the value-neutrality thesis has been questioned by C. W. Churchman (1948, ch. xv) and R. B. Braithwaite (1955, pp. 250–4) at least insofar as it applies to statistical inference. However, the most explicit

* I wish to acknowledge my debt to Sidney Morgenbesser, whose critical comments in conversation have greatly influenced my thinking on this question, and to Mortimer Kadish and John McLellan, whose reactions to earlier drafts of this paper have helped shape the final result.

† Reprinted from *The Journal of Philosophy* **57** (1960), pp. 345–57.

and sweeping attack against the value-neutrality thesis is to be found in an article by Richard Rudner (1953), who argues that the scientist must make value judgements in drawing any kind of non-deductive inference.

> Now I take it that no analysis of what constitutes the method of science would be satisfactory unless it comprised some assertion to the effect that the scientist as scientist accepts or rejects hypotheses.
>
> But if this is so then clearly the scientist as scientist does make value judgements. For, since no scientific hypothesis is ever completely verified, 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 regarding the evidence and 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. Thus, to take a crude but easily manageable example, if the hypothesis under consideration were to the effect that a toxic ingredient of a drug was not present in lethal quantity, we would require a relatively high degree of confirmation or confidence before accepting the hypothesis – for the consequences of making a mistake here are exceedingly grave by our moral standards. On the other hand, if, say, our hypothesis stated that, on the basis of a sample, a certain lot of machine stamped belt buckles was not defective, the degree of confidence we should require would be relatively not so high. *How sure we need to be before we accept a hypothesis will depend on how serious a mistake would be* (Rudner, 1953, p. 2).

Rudner's claim is not that values play a role in the scientist's selection of research problems, nor is he arguing that scientists often let their attitudes, values, and temperaments influence their conclusions. These points are relevant to the psychology and sociology of inquiry but not to its logic. Rudner is making an assertion about the requirements imposed upon the inquirer who embraces the goals and the canons of scientific inference.[1] He contends that the scientist in his capacity as a scientist *must* make value judgements even if it is psychologically possible for him to avoid doing so. His argument for this conclusion can be summarized in the following series of statements:

(1) The scientist *qua* scientist accepts or rejects hypotheses.

(2) No amount of evidence ever completely confirms or disconfirms any (empirical) hypothesis but only renders it more or less probable.

(3) As a consequence of (1) and (2), the scientist must decide how high the probability of a hypothesis relative to the evidence must be before he is warranted in accepting it.

(4) The decision required in (3) is a function of how important it will be if a mistake is made in accepting or rejecting a hypothesis.

1 The canons of scientific inference can be construed to be normative principles. The value-neutrality thesis does not deny this but does insist that given an initial commitment to these principles, the scientist need not and should not let his values, attitudes, and temperament influence his inferences any further. It is this claim that Rudner appears to deny.

The need for assigning minimum probabilities for accepting and rejecting hypotheses (3) is a deductive consequence of the claim that scientists accept and reject hypotheses (1) and the corrigibility of empirical hypotheses (2). Since (2) is a cardinal tenet of an empiricist philosophy of science and will not be questioned in this paper, the first part of Rudner's argument reduces to the correct claim that if (1) is true (3) is true.

Rudner's rejection of the value-neutrality thesis cannot be justified, however, on the basis of (3) alone. He must show that the assignment of minimum probabilities is a function of the importance of making mistakes (4). But (4) cannot be obtained from (3) without further argument.[2] Rudner attempts to fill the gap by citing illustrations from quality control and appealing to current theories of statistical inference. He believes that the problem of choosing how to act in the face of uncertainty, which is the fundamental problem of quality control, is typical of all scientific inquiry and concludes from this that the importance of making mistakes must be taken into account in all scientific inference.

This argument seems to rest upon certain assumptions adopted more or less explicitly by Rudner and Churchman (1956, p. 248). These assumptions involve the notion of acting on the basis of a hypothesis relative to an objective. To say 'X acts on the basis of H relative to some objective P' is to assert that X carries out action A where A is the best procedure[3] to follow relative to P, given that H is true. The Rudner–Churchman assumptions can now be stated as follows:

(5) To choose to accept a hypothesis H as true (or to believe that H is true) is equivalent to choosing to act on the basis of H relative to some specific objective P.

(6) The degree of confirmation that a hypothesis H must have before one is warranted in choosing to act on the basis of H relative to an objective P is a function of the seriousness of the error relative to P resulting from basing the action on the wrong hypothesis.

Assumption (6) is a version of a principle adopted by Pearson, Neyman, and Wald in their theories of statistical inference. The plausibility of Rudner's argument from quality control (where the problem is how to act on the basis of hypotheses) to (4) is due largely to the reasonableness of this presupposition. However, (6) without (5) will not yield (4).

2 Actually Rudner's version of (4) is stronger than mine. According to Rudner, the importance of making a mistake can be construed in 'a typically ethical sense.' In order to simplify the discussion, this rider will be dropped. The importance of making a mistake will be understood to be a function of the values, attitudes, preferences, and temperament of the investigator or group whose interests he serves regardless of the ethical character of these values, etc. Understood in this sense, (4) is still incompatible with the value-neutrality thesis.

3 Perhaps 'A is believed by X to be the best procedure' should replace 'A is the best procedure'. The following discussion does not, however, demand a choice between these two definitions.

Unlike (6), (5) cannot be justified by an appeal to the authority of the statisticians. Not only are these authorities fallible, but some of them have been non-committal regarding the acceptability of (5).[4] Substantial grounds can be offered for praising this exercise of caution.

2

An interesting case against the tenability of (5) has been made by Richard Jeffrey. Jeffrey considers the problem of deciding whether a given batch of polio vaccine is free from active polio virus. The seriousness of the consequences of mistakenly accepting the hypothesis would seem to demand that we confirm the hypothesis to a far higher degree before accepting it than would be the case if we were interested in the quality of a batch of roller skate bearings.

> But what determines these consequences? There is nothing in the hypothesis, 'This vaccine is free from active polio virus,' to tell us what the vaccine is *for*, or what would happen if the statement were accepted when false. One naturally assumes that the vaccine is intended for inoculating children, but for all we know from the hypothesis it might be intended for inoculating pet monkeys. One's confidence in the hypothesis might well be high enough to warrant inoculation of monkeys but not of children (Jeffrey, 1956, p. 242).

Jeffrey's point can be reformulated as follows: Action on the basis of a hypothesis H is always relative to an objective P. Consequently if accepting H is identical with acting on the basis of H (5), accepting H in an 'open-ended' situation[5] where there is no specific objective is impossible. But accepting H is possible in open-ended situations, for it is compatible with different and even conflicting objectives. Hence, (5) must be rejected.

In a reply to Jeffrey's paper, Churchman compares Jeffrey's open-ended decision problems to situations that occur in production. Suppose that a manufacturer wishes to place on the market a certain product (rope) that has many different uses. Churchman points out that procedures are available to the manufacturer in terms of which he can single out needs that his product should be designed to meet. He contends that similar procedures must be employed if we are to accept and reject hypotheses intelligibly.

4 'The terms "accepting" and "rejecting" a statistical hypothesis are very convenient and are well established. It is important, however, to keep their exact meaning in mind and to discard various additional implications which may be suggested by intuition. Thus, to accept a hypothesis H means only to decide to take action A rather than action B. This does not mean that we necessarily believe that the hypothesis H is true. Also if the application of a rule of inductive behavior "rejects" H, this means only that the rule prescribes action B and does not imply that we believe that H is false.' (Neyman, 1950, 259–260.) In this passage, Neyman does identify accepting a hypothesis H with acting on H. However, he refuses to identify accepting H with believing that H. In effect, therefore, he suspends judgement regarding the truth of (5).

5 This expression is due to Churchman (1956, 248).

In this sense, it is certainly meaningless to talk of *the* acceptance of the hypothesis about the freedom of a vaccine from active polio virus, provided the information has a number of different uses. Even within one business organization one can readily point out that the many uses of information imply many different criteria for the 'acceptance' or 'rejection' of hypotheses (1956, p. 248–9).

Churchman's argument seems to be this: *If* 'accepting a hypothesis H' is understood in a sense that makes (5) true, then open-ended decision problems involving the acceptance or rejection of hypotheses can be treated like open-ended production problems. The solvability and, hence, the intelligibility of such problems requires the elimination of the open-endedness.

This true observation does not meet, however, the major point of Jeffrey's objection. Jeffrey's argument attempts to show that in *one* sense of 'accepting a hypothesis' to accept a hypothesis in an open-ended situation is perfectly meaningful and consistent. Consequently, in *that* sense, (5) does not hold.

An easy but cheap victory might be gained at Jeffrey's expense by pointing out that wherever a scientist does not appear to have an objective in mind, nonetheless, one can always be specified – namely, the objective of accepting true answers to questions as true. Accepting a hypothesis H would then be equivalent to acting on the basis of H relative to that objective.

Resorting to this strategy would be to miss the point of the discussion. To say that accepting a hypothesis is the same as acting on the basis of H in order to obtain true answers is tantamount to asserting that accepting H is equivalent to accepting H. One could not conclude from this alone that the problem of deciding what to believe is on all fours with decision problems in quality control – at least with respect to the value-neutrality thesis. In the latter kind of problem, the objectives are 'practical'; in the former, they are 'theoretical'.

In order to avoid misunderstanding, therefore, an open-ended decision problem will be understood to be a decision problem for which no practical objective has been specified.[6] Consequently, the issue at stake in the debate between Jeffrey and Churchman is whether there is any sense in which a person can meaningfully and consistently be said to accept a hypothesis as true without having a practical objective. The following considerations are offered in favor of an affirmative answer to this question.

(*i*) Many apparently intelligible questions are raised and answered in the sciences for which practical objectives are difficult to specify. What practical

6 By a 'theoretical' objective, I shall understand any objective that is concerned with selecting true hypotheses from a given list. A practical objective is one that is not theoretical. This dichotomy overlooks distinctions between ethical, practical, and aesthetic objectives by grouping them together. It also treats many objectives as practical that might legitimately be held to be theoretical. The purpose of the twofold partition of objectives, however, is to avoid a trivial interpretation of (5) while permitting Churchman and Rudner as much leeway as possible in their interpretation of this assumption.

objectives are at stake when an investigator is deciding whether to accept or reject the principle of parity, the hypothesis of an expanding universe, or the claim that Galileo never conducted the Leaning Tower experiment? One could try to show that appearances are deceiving and that practical objectives are always the goals of such decision problems. However, this would be difficult to prove. Furthermore, it would not follow that appearances *must* be deceiving and that practical objectives *must* be operative. Indeed, the cases just cited would normally be considered to be problems of deciding what to accept as or believe to be true regardless of whether practical objectives are involved. This seems to indicate that there is a sense of 'accepting a hypothesis' which is meaningfully applied to choices in open-ended situations.

(*ii*) Even in the case of decision problems where practical objectives are involved, it often seems appropriate to distinguish between acting on the basis of a hypothesis relative to that objective and accepting the hypothesis as true. Suppose that an investor in oil stocks knows that if a certain oil company whose stocks are selling at a low price strikes oil at a certain location, the price of the stock will increase one hundredfold. The investor might buy stock in the company while suspending judgement as to the eventual discovery of oil. Here is a case where one would normally say that a person has acted on the basis of a hypothesis and perhaps was justified in doing so without accepting the hypothesis as true or being warranted in so accepting it.

One could reply by saying that the investor refused to accept the hypothesis because such acceptance would have been tantamount to acting on the basis of the hypothesis of an oil strike relative to some practical objective other than making a profitable investment. However, such an objective would not always be easy to find. Furthermore, the situation would normally be considered a case of action without belief regardless of whether the existence of a practical objective could be shown or not.

(*iii*) There seems to be a sense in which it is possible for a person to believe in the truth of a hypothesis and nonetheless refuse to act on it. He may even be justified in proceeding in this fashion. The Sabin live virus polio vaccine serves as an illustration. The available evidence might warrant belief in the safety and effectiveness of the vaccine without justifying a program of mass inoculation.[7]

7 This claim might seem counterintuitive. There is a widely held view that if a person really believes in a hypothesis he should be ready to act on it. This 'put up or shut up' analysis may be understood in two ways: (*a*) belief in H implies acting on the basis of H, and (*b*) belief in H implies that one ought to act on the basis of H. R. M. Martin seems to adopt the former view (1959, p. 11). This version of the 'put up or shut up' analysis does not seem adequate to at least one familiar sense of 'accepting a hypothesis'. The very fact that people often think that one ought to act on a hypothesis if one believes it implies that one might not so act. Sense (*b*) of the 'put up or shut up' analysis seems more plausible. Nonetheless it yields results that themselves appear to be counterintuitive. If this thesis demands readiness to act relative to

(*iv*) A plausible case can be made for saying that even when a person is deciding how to act in order to realize a practical objective he will have to accept some statements as true in a sense that does not meet the conditions of (5). The evidence upon which he bases his decisions consists of statements which he accepts as true. He might have to accept the truth of statements asserting the degrees to which various hypotheses are confirmed relative to the available evidence. Finally, he will also have to accept the truth of statements that indicate the best actions relative to his objectives given the truth of various hypotheses.[8]

The considerations just advanced suggest that there is a familiar sense in which a person can meaningfully and consistently accept or reject a hypothesis in an open-ended situation. In that sense, (5) is false and Rudner's argument in favor of (4) and against the value-neutrality thesis fails.

This result need not in itself be fatal to the Churchman–Rudner position. Apologists for this view could admit the meaningfulness of this sense of 'accepting a hypothesis' and deny that the aim of the sciences is (or ought to be) to accept or reject hypotheses in that sense. They might contend that scientific inferences indicate how one ought to act on the basis of hypotheses but not what one ought to believe. The rejection of the value-neutrality thesis would flow quite naturally from this transmutation of scientific inquiry into a quest for normative principles. Oddly enough, however, it is Jeffrey, an apparent defender of the value-neutrality thesis, who denies that scientists accept and reject hypotheses.

3

Jeffrey proposes a conception of the aim and function of science also suggested by Carnap (1950, pp. 205–7) and Hempel (1949, p. 560). According to this view, a scientist does not, or at least should not, accept and reject hypotheses. Instead, he should content himself with assigning degrees of confirmation to hypotheses relative to the available evidence. Anyone who is confronted with a practical decision problem can go to the scientist to ascertain the degrees of confirmation of the relevant hypotheses. He can then utilize this information together with his own estimates of the seriousness of mistakes in order to decide upon a course of action.

One consequence of this view is that all non-deductive inference in science

any objective, then one would not be warranted in accepting a hypothesis as true unless the degree of confirmation approached certainty. For there is always the possibility that some objectives exist relative to which mistakes are so serious as to demand enormously high degrees of confirmation. Such a requirement seems unreasonable. On the other hand, if the objectives relative to which one should be ready to act are restricted in some way, it is difficult to see how the restrictions could be specified without destroying the initial plausibility of the 'put up or shut up' analysis.

8 I owe this observation to Mortimer Kadish.

consists in assigning degrees of confirmation to hypotheses relative to given evidence. Indeed, Carnap (1950, p. 206) defines inductive inference in this way. Hence, if Carnap is correct in maintaining that degrees of confirmation can be ascertained without consideration of values, the Carnap–Hempel–Jeffrey view supports the value-neutrality thesis.[9] However, the value neutrality thesis is upheld at the expense of the claim that scientists accept or reject hypotheses. In this respect, the Carnap–Hempel–Jeffrey view breaks as radically with tradition as does the Braithwaite–Churchman–Rudner position.

In his paper, Jeffrey offers an extremely clever argument to show that scientists can neither accept nor reject hypotheses.

> On the Churchman–Braithwaite–Rudner view it is the task of the scientist as such to accept and reject hypotheses in such a way as to maximize the expectation of good for, say a community for which he is acting. On the other hand, our conclusion is that if the scientist is to maximize good he should refrain from accepting or rejecting hypotheses, since he cannot possibly do so in such a way as to optimize every decision which may be made on the basis of those hypotheses. We note that this difficulty cannot be avoided by making acceptance relative to the most stringent possible set of utilities (even if there were some way of determining what that is) because then the choice would be wrong for all less stringent sets. One cannot, by accepting or rejecting the hypothesis about the polio vaccine, do justice both to the problem of the physician and the veterinarian. The conflict can be resolved if the scientist either contents himself with providing them both with a single probability for the hypothesis (whereupon each makes his own decision based on the utilities peculiar to his problem) or if the scientist takes on the job of making a separate decision as to the acceptability of the hypothesis in each case. In any event, we conclude that it is not the business of the scientist as such, least of all of the scientist who works with lawlike hypotheses, to accept or reject hypotheses (Jeffrey, 1956, p. 245).

Jeffrey's argument rests upon two lemmas: (*a*) if scientists accept and reject hypotheses (1), then they must make value judgements (4); and (*b*) if (1) is true, then (4) is false. The inevitable conclusion is that (1) is false – i.e., that the scientist neither accepts nor rejects hypotheses.

Jeffrey accepts (*a*) without any question as having been established by Rudner. His argument for (*b*) may be paraphrased as follows: Deciding whether to accept or reject a hypothesis is an open-ended decision problem – i.e., there is no practical objective in terms of which seriousness of error can be assessed. Hence, if a scientist decides to accept or reject a hypothesis, he cannot be taking the seriousness of error into account. Consequently, if (1) is true, (4) is false.

9 The difference between this view and the revised version of the Churchman–Rudner position suggested above is that the latter considers the scientist as a formulator of practical policy whereas the former considers him to be an adviser to the policy maker. This difference reflects itself in differing conceptions of non-deductive inference. According to the revised Churchman–Rudner view, the 'conclusion' of a non-deductive inference is a choice of a course of action. According to the Carnap–Hempel–Jeffrey view, the conclusion is an assignment of a degree of confirmation to a hypothesis.

In spite of its persuasive character, Jeffrey's argument breaks down at several points.

(*i*) Rudner's argument for lemma (*a*) has already been shown to hold only if accepting a hypothesis is understood to be meaning*less* in open-ended situations. On the other hand, Jeffrey's argument for (*b*) depends upon the understanding that accepting a hypothesis is meaning*ful* in such cases. Hence, Jeffrey is guilty of equivocation.

(*ii*) Jeffrey's argument from the truth of (1) to the falsity of (4) depends upon the assumption that the decision problem is an open-ended one. An open-ended decision problem has been understood to be one that lacks a *practical* objective. However, such problems may still have a theoretical objective. It is at least an open question whether such an objective can serve as a basis for ascertaining the seriousness of mistakes.

(*iii*) Even if theoretical objectives cannot function in this way, Jeffrey's inference from (1) to the negation of (4) can still be avoided. It has been argued that Jeffrey is correct in asserting and Churchman is wrong in denying that there is a sense of 'accepting a hypothesis' that is meaningful in open-ended situations. This does not mean, however, that this sense of 'accepting a hypothesis' is meaningful *only* in open-ended situations. A person may decide what to believe only in order to believe true statements. But he may wish to believe statements which are true and which have some other desirable characteristic such as simplicity, explanatory power, effectiveness as propaganda, or a consoling emotive connotation. And the sense in which he accepts a statement as true in attempting to realize one of these objectives will be the same sense in which he might accept statements as true in open-ended situations. Again, it is at least an open question whether a scientist *qua* scientist has such a practical objective in accepting and rejecting hypotheses and, hence, has a basis for determining the seriousness of mistakes.

The failure of Jeffrey's argument does not, of course, imply the falsity of his conclusion. Indeed, another argument can be offered for rejecting (1). Whatever may be the merits of the inference from (1) to (4), empiricists are committed to accepting the inference from (1) to (3) – i.e., the inference from the claim that scientists accept and reject hypotheses to the need for assigning minimum probabilities for such acceptance and rejection. How are such minimum probabilities to be assigned? If no plausible alternative to a procedure that takes the values of the investigator into account is available, then (1) entails the rejection of the value-neutrality thesis.

Defenders of the Carnap–Hempel–Jeffrey view might feel that we are in such a predicament. Not wishing to abandon the value-neutrality thesis, they reject the conception of the scientist as one who accepts and rejects hypotheses.[10] However, following this strategy is like crashing into Scylla in

10 Hempel (1949) comes closer to arguing in this way than either Carnap or Jeffrey.