Rational Prediction
Author(s): Wesley C. Salmon
Source: *The British Journal for the Philosophy of Science*, Vol. 32, No. 2 (Jun., 1981), pp. 115-125
Published by: Oxford University Press on behalf of The British Society for the Philosophy of Science
Stable URL: http://www.jstor.org/stable/687193
Accessed: 20/10/2010 09:38

Brit. J. Phil. Sci. 32 (1981), 115–125 *Printed in Great Britain*

Rational Prediction*

by WESLEY C. SALMON

A colleague, to whom I shall refer (quite accurately) as 'the friendly physicist,' recently recounted the following incident. While awaiting take-off on an airplane, he noticed a young boy sitting across the aisle holding onto a string to which was attached a helium-filled balloon. He endeavoured to pique the child's curiosity. 'If you keep holding the string just as you are now,' he asked, 'what do you think the balloon will do when the airplane accelerates before take off?' The question obviously had not crossed the youngster's mind before that moment, but after giving it a little thought, he expressed the opinion that the balloon would move toward the back of the cabin. 'I don't think so,' said the friendly physicist, 'I think it will move forward.' The child was now eager to see what would happen when the plane began to move. Several adults in the vicinity were, however, skeptical about the physicist's prediction; in fact, a stewardess offered to wager a miniature bottle of Scotch that he was mistaken. The friendly physicist was not unwilling and the bet was made. In due course, the airplane began to accelerate, and the balloon moved toward the front of the cabin. The child's curiosity was satisfied[1]; the theory—that all objects which are free to move will move toward the back of the cabin when the plane accelerates—was falsified; and the friendly physicist enjoyed a free drink.

I have related this anecdote to point out that there are at least three—probably more—legitimate reasons for making predictions. First, we are sometimes curious about future happenings, and we want to satisfy that curiosity without waiting for the events in question to transpire. To do so, we may make wild guesses, we may employ superstitious methods of prediction, we may appeal to common sense, or we may use more sophisticated scientific theories. Second, we sometimes make predictions for the sake of testing a theory. In the example at hand, the prediction regarding the motion of the balloon was a rather good test of the hypothesis that all objects free to move in the cabin will tend to move toward the rear when the airplane accelerates. The fact that objects heavier-than-air tend to fall toward the earth when they are unsupported, while objects lighter-than-

* A version of this paper was presented orally at the Symposium on the Philosophy of Sir Karl Popper, London School of Economics, 14–16 July 1980. This material is based upon work supported by the National Science Foundation (U.S.A.) under Grant No. SES-7809146.

[1] His curiosity regarding *what* would happen was satisfied, though not his curiosity as to *why*.

air (such as helium-filled balloons) tend to move in the opposite direction, suggests that the behaviour of a helium-filled balloon has a reasonable chance of falsifying the hypothesis about the behaviour of all material objects in the air-filled cabin of the accelerating airplane, if it is indeed false. Third, we sometimes find ourselves in situations in which some practical action is required, and the choice of an optimal decision depends upon predicting future occurrences. Although wagering is by no means the only such type of practical decision-making, it is a clear and comprehensible example. We all agree, I take it, that scientific theories often provide sound bases for practical prediction.

A central feature of Sir Karl Popper's philosophy is his thesis concerning the status of induction. Indeed, he begins his book, *Objective Knowledge*, with the statement:

> I think that I have solved a major philosophical problem: the problem of induction. . . . This solution has been extremely fruitful, and it has enabled me to solve a good number of other philosophical problems (Popper [1972], p. 1).

His solution, as is well known, involves a complete rejection of induction. This claim has been advanced in many of his writings spanning several decades, and it is reiterated in his Autobiography and in his 'Replies to My Critics' in *The Philosophy of Karl Popper* (Schlipp [1974]).

For some time it has seemed to me that the crucial test of an anti-inductivist philosophy of science would be its capacity to deal with the predictive aspects of scientific knowledge. In a paper (Salmon [1968]) presented at the 1965 International Colloquium on Philosophy of Science at Bedford College, London, I attempted to offer a severe challenge to Popper's views concerning induction by posing what I took to be a serious dilemma: on Popper's account, science either embodies essential inductive aspects or else science is lacking in predictive content.[1] In the published proceedings of the Bedford College Colloquium (Lakatos [1968]), J. W. N. Watkins contributed an answer to my critique. He denied that scientific reasoning is inductively infected, and he argued that it can, nevertheless, provide a basis for rational prediction. In the replies to his critics (Schilpp [1974], pp. 1028–30), Popper acknowledges that I have understood his views 'fairly well,' and he endorses Watkins's response. I take this as evidence that we have located a genuine disagreement—one which is reasonably free from purely verbal disputes or out-and-out misrepresentations—regarding Popper's anti-inductivist stand. The question involves what Popper calls 'the pragmatic problem of induction'. It is this issue which I want to pursue in the present paper; it concerns the problem of rational prediction. Although the issue may appear to be rather narrow, it seems to me to have pivotal importance with regard to the assessment of Popper's deductivism.

[1] Similar themes were developed in Salmon [1967], chap. II, §3.

Let me attempt to formulate the basic difficulty as I see it. In its very simplest terms, Popper's account of scientific knowledge involves generalisations and their observational tests. If we find a *bona fide* counterexample to a generalisation, we can say that it has been deductively refuted. To be sure, as Popper explicitly acknowledges, there may be difficulties in some cases in determining whether certain observations constitute genuine counterexamples to a generalisation, but that does not undermine the claim that a genuine counterexample yields a deductive refutation. According to Popper, negative instances provide rational grounds for rejecting generalisations. If, however, we make observations and perform tests, but no negative instance is found, all we can say deductively is that the generalisation in question has not been refuted. In particular, positive instances do not provide confirmation or inductive support for any such unrefuted generalisation. At this stage, I claim, we have no basis for rational prediction. Taken in themselves, our observation reports refer to past events, and consequently they have no predictive content. They say nothing about future events. If, however, we take a general statement as a premise, and conjoin to it some appropriate observation statements about past or present events, we may be able to deduce a conclusion which says something about future occurrences and which, thereby, has predictive content. Popper himself gives this account (Schilpp [1974], p. 1030) of *the logic of prediction.*

The problem of rational prediction concerns the status of the general premise in such an argument. One may claim, as Popper does, that we ought not to use a generalisation which has actually been refuted as a premise in a predictive argument of this sort, for we are justified in regarding it as false. We ought not to employ premises which are known to be false if we hope to deduce true predictions. The exclusion of refuted generalisations does not, however, tell us what general premise should be employed. Typically there will be an infinite array of generalisations which are compatible with the available observational evidence, and which are therefore, as yet, unrefuted. If we were free to choose arbitrarily from among all the unrefuted alternatives, we could predict anything whatever. If there were no rational basis for choosing from among all of the unrefuted alternatives, then, as I think Popper would agree, there would be no such thing as rational prediction. We are not in this unfortunate situation, Popper contends, for we do have grounds for preferring one unrefuted generalisation to another:

> My *solution* of the logical problem of induction was that we may have *preferences* for certain of the competing conjectures; that is, for those which are highly informative and which so far have stood up to eliminative criticism (Schilpp [1974], p. 1024).

Popper's concept of corroboration is designed to measure the manner in which conjectures have stood up to severe criticism, including severe

testing. This, I take it, is the crucial thesis—that *there is a rational basis for preferring one unrefuted generalisation to another for use in a predictive argument.* If that is correct, then Popper can legitimately claim to have solved the problem of rational prediction.

If we are going to talk about preference among generalisations, then we have to be quite explicit about the purpose for which the generalisation is to be used. In this context, we are discussing prediction, so the preference must be in relation to predictive capability. As Popper rightly insists, any generalisation we choose will have predictive import in the sense that it will make statements about future events—more precisely, in a predictive argument as characterised above, it yields conclusions about future occurrences. But since all of the various unrefuted generalisations have predictive content in that sense, we must still ask on what basis the predictive content of one conjecture is rationally preferable to that of another conjecture.

At this stage of the discussion, it is important to recall the point of the opening story, namely, that predictions are made for various purposes. Thus, even if we agree that we want to select a generalisation for predictive purposes, we must still specify what type of prediction is involved. Popper explicitly acknowledges (Schilpp [1974], pp. 1024–5) that there are two types of preference, 'the theoretician's preference', and that of 'the man of practical action'. As I understand Popper's view, the theoretician is interested in formulating bold conjectures which have high content and in subjecting them to severe tests. Insofar as the theoretician is mainly interested in explanations of known phenomena, he may not be much involved in making any sorts of predictions. I suppose we might distinguish the theoretician's explanatory preference from the theoretician's predictive preference, recognising that there is bound to be a close connection between preferences of these two kinds. When the theoretician is actually involved *qua theoretician* in making predictions, the purpose is to devise (and, perhaps, to instruct the experimentalist on how to conduct) a severe test. The purpose of predictions made in this theoretical context is to gain information which is useful in the evaluation of scientific theories. If the chief value of the scientific theories is explanatory, then it is not at all clear that a primary desideratum of the predictive argument is to arrive at a true prediction. As Popper has emphasised, and as all of us know, a false prediction can be valuable, since the realisation (on the basis of observation) that it is false can be highly informative.

Having noted the distinction between the theoretical preference and the practical preference, let us now focus attention upon the kind of preference which is pertinent to the practical context, with special attention to the kinds of predictions which play a role in practical decision-making. As I have remarked above, Popper claims that for theoretical purposes we prefer theories which are highly corroborated to those which are less well

corroborated. I do not think this claim is unproblematic, but I do not propose arguing the matter here. My aim is to emphasise that, even if we are entirely justified in letting such considerations determine our theoretical preferences, it is by no means obvious that we are justified in using them as the basis for our preferences among generalisations which are to be used for prediction in the practical decision-making context. Popper and Watkins have maintained, however, that corroboration should play a crucial role in determining both theoretical preference and practical preference.

Since scientific theories are used for both theoretical and practical purposes—including prediction—and since, according to Popper, theory preference is based upon corroboration, I had mistakenly inferred (prior to 1968) that the appraisal of a theory in terms of corroboration must imply some attempt at an appraisal of the theory with respect to its future performance. If that were Popper's thesis, I had argued, then corroboration must involve some element of induction (or non-demonstrative inference of some sort), for past performance of the theory is taken to constitute a basis for some sort of claim about future performance. However, I have since been informed by Watkins [1968], Settle [1970], and Popper [1974] that I had misconstrued Popper's view. Statements about the corroboration of theories are no more than appraisals of their past performances; corroboration statements hold no predictions with respect to future performance. If they did, they would be inductive (as I had claimed); but they are not inductive, so they cannot be predictive.

This view of corroboration holds serious difficulties. Watkins and Popper agree, I take it, that statements which report observations of past and present events do not, in and of themselves, have any predictive content. Moreover, they maintain, statements about the corroboration of conjectures do not, in and of themselves, have any predictive content. Conjectures, hypotheses, theories, generalisations—call them what you will—do have predictive content. The problem is that there are many such statements, rich in predictive content, which make incompatible predictive claims when conjoined with true statements about past and present occurrences. The fact that a general statement has predictive content does not mean that what it says is true. In order to make a prediction, one must choose a conjecture which has predictive content to serve as a premise in a predictive argument. In order to make a *rational* prediction, it seems to me, one must make a *rational* choice of a premise for such an argument. But from our observational evidence and from the statements about the corroboration of a given conjecture, no predictive appraisal follows. Given two conjectures which, in a particular situation, will lead to incompatible predictions, and given the corroboration ratings of these two hypotheses, *nothing follows* about their comparative predictive capacities. Thus, it seems to me, corroboration—the ground for

theoretical preference—furnishes no rational basis for preference of one conjecture to another *for purposes of practical prediction.* I am not complaining that we are not told *for sure* that one will make a correct prediction and that the other will not. I am complaining that no rational basis whatever has been furnished for a preference of *this* type.

In his reply to my Bedford College paper, Watkins acknowledges that there is an important distinction between theoretical and practical preferences, and he further acknowledges that the two kinds of appraisal may have quite different bases:

> Now our methods of hypothesis-selection in practical life should be well-suited to our practical aims, just as our methods of hypothesis-selection in theoretical science should be well suited to our theoretical aims; and the two kinds of method may very well yield different answers in a particular case (Lakatos [1968], p. 65).

He goes on to explain quite correctly how utility considerations may bear upon the practical situation. Then he considers the case in which utility does not play a decisive role:

> Now suppose that, for a particular agent, the mutually incompatible hypotheses h_1 and h_2 are on a par utility-wise, and that in the situation in which he finds himself, he has *got* to act since 'inaction' would itself be one mode of action. Then if h_1 is the only alternative to h_2 before him, he *has* to choose one of them. Then it would be rational for him to choose the better corroborated one, the one which has withstood the more severe criticism, since he has nothing else to go on (*ibid.*, pp. 65–6).

Watkins offers no further argument for supposing that corroboration provides a rational basis for *practical* preference. Moreover, the hint of an argument which he does supply appeals to a false premise. The agent does have other things 'to go on.' He could decide between the two hypotheses by the flip of a coin. He could count the numbers of characters in each of the two hypotheses in the particular formulation given, and choose the one which has fewer. He could choose the hypothesis which comes first lexicographically in the given formulation. What Watkins is suggesting, it seems to me, is not that the agent has 'nothing else to go on,' but rather, that he has no other *rational* basis for preference. But such an argument would be patently question-begging. Even if all other bases for choice were irrational, it would not follow that the one cited by Watkins is *ipso facto* rational. Indeed, if we take seriously Popper's statement, 'I regarded (and I still regard) the degree of corroboration of a theory merely as a critical report on the quality of past performance: *it could not be used to predict future performance*' (Schilpp [1974], p. 82), it is hard to see how corroboration can supply a rational basis for preference of a theory *for purposes of practical prediction.*

Whether my criticism of Popper's position is correct or incorrect, the issue I am raising has fundamental importance. For if it should turn out that Popper could not provide a tenable account of rational prediction—

given his persistent emphasis upon objectivity and rationality—then we could hardly credit his claim to have solved the problem of induction. Moreover, in his replies to his critics, Popper acknowledges the issue. With the comment, 'Our corroboration statements have no predictive import, although they motivate and justify our *preference* for some theory over another,' (Schilpp [1974], pp. 1029–30), he endorses the answer Watkins had furnished. Since I am not attempting to deal with the psychological problem of induction, I shall not dispute the claim that corroboration may *motivate* the preference of one theory to another. What I want to see is how corroboration could *justify* such a preference. Unless we can find a satisfactory answer to that question, it appears to me that we have no viable theory of *rational* prediction, and no adequate solution to the problem of induction.

In *Objective Knowledge*, Popper offers an answer to the basic question which seems closely related to that of Watkins:

> . . . *a pragmatic belief in the results of science* is not irrational, because there is nothing more 'rational' than the method of critical discussion, which is the method of science. And although it would be irrational to accept any of its results as certain, there is nothing 'better' when it comes to practical action: there is no alternative method which might be said to be more rational (p. 27).

This response appears to miss the point. The question is not whether other methods—*e.g.*, astrology or numerology—provide more rational approaches to prediction than does the scientific method. The question is whether the scientific approach provides a more rational basis for prediction, for purposes of practical action, than do these other methods. The position of the Humean skeptic would be, I should think, that none of these methods can be shown either more or less rational than any of the others. But if every method is equally lacking in rational justification, then there is no method which can be said to furnish a rational basis for prediction, for any prediction will be just as unfounded rationally as any other. If the Humean skeptic were right, we could offer the following parallel claim. A pragmatic belief in the predictions found in Chinese fortune cookies is not irrational, for there is nothing more rational. . . .

In his replies to his critics, Popper again addressed the problem, and he came more firmly to grips with it:

> But every action presupposes a set of expectations, that is, of theories about the world. Which theory shall the man of action choose? Is there such a thing as a *rational choice*?
>
> This leads us to the *pragmatic problems of induction*, which to start with, we might formulate thus:
>
> (*a*) Upon which theory should we rely for practical action, from a rational point of view?
>
> (*b*) Which theory should we prefer for practical action, from a rational point of view?

My answer to (*a*) is: from a rational point of view, we should not 'rely' on any theory, for no theory has been shown to be true, or can be shown to be true (or 'reliable').

My answer to (*b*) is: we should *prefer* the best tested theory as a basis for action.

In other words, there is no 'absolute reliance'; but since we *have* to choose, it will be 'rational' to choose the best tested theory. This will be 'rational' in the most obvious sense of the word known to me: the best tested theory is the one which, in the light of our *critical discussion*, appears to be the best so far; and I do not know of anything more 'rational' than a well-conducted critical discussion (Schilpp [1974], p. 1025).

Let us not be seduced by honeyed words. If we wish to claim that a theory 'appears to be the best so far,' we must ask, 'Best for what purpose—theoretical explanation or practical prediction?' Since it is 'the best tested theory' and it has been subjected to 'critical discussion,' then, in the light of the many statements by Popper *et al.* about the lack of predictive import of corroboration, we must conclude, I believe, that the answer is, 'Best for theoretical explanation.' Perhaps I am being unduly obtuse, but I cannot see that any reason has been provided for supposing that such a theory is best *for practical prediction.*

I must confess to the feeling that we have been 'given the run-around'. We begin by asking how science can possibly do without induction. We are told that the aim of science is to arrive at the best explanatory theories we can find. When we ask how to tell whether one theory is better than another, we are told that it depends upon their comparative ability to stand up to severe testing and critical discussion. When we ask whether this mode of evaluation does not contain some inductive aspect, we are assured that the evaluation is made wholly in terms of their comparative success up to now; but since this evaluation is made entirely in terms of past performance, it escapes inductive contamination because it lacks predictive import. When we then ask how to select theories for purposes of rational prediction, we are told that we should prefer the theory which is 'best tested' and which 'in the light of our *critical discussion*, appears to be the best so far', even though we have been explicitly assured that testing and critical discussion have no predictive import. Popper tells us, 'I do not know of anything more "rational" than a well-conducted critical discussion.' I fail to see how it could be rational to judge theories *for purposes of prediction* in terms of a criterion which is emphatically claimed to be lacking in predictive import.[1]

Fearing that the point of the preceding argument may have been missed, Popper attempts another formulation:

Let us forget momentarily about what theories we 'use' or 'choose' or 'base our practical actions on', and consider only the resulting *proposal* or *decision* (to do X;

[1] The argument advanced in this paragraph bears a strong resemblance, I think, to one developed in Grünbaum [1976]; see especially p. 246.

not to do X; to do nothing; or so on). Such a proposal can, we hope, be rationally criticized; and if we are rational agents we will want it to survive, if possible, the most testing criticism we can muster. *But such criticism will freely make use of the best tested scientific theories in our possession.* Consequently any proposal that ignores these theories (where they are relevant, I need hardly add) will collapse under criticism. Should any proposal remain, it will be rational to adopt it.

This seems to me all far from tautological. Indeed, it might well be challenged by challenging the italicized sentence in the last paragraph. Why, it might be asked, does rational criticism make use of the best tested although highly unreliable theories? The answer, however, is exactly the same as before. Deciding to criticize a practical proposal from the standpoint of modern medicine (rather than, say, phrenological terms) is itself a kind of 'practical' decision (anyway it may have practical consequences). Thus the rational decision is always: adopt critical methods which have themselves withstood severe criticism (Schilpp [1974], pp. 1025–6).

I have quoted Popper *in extenso* to try to be quite sure not to misunderstand his answer. The italicised sentence in the first paragraph raises precisely the question which seems to me crucial. In the second paragraph, Popper admits the legitimacy of the question, and he offers an answer. When he says, 'The answer . . . is exactly the same as before . . . the rational decision is always: adopt critical methods which have themselves withstood severe criticism,' he seems to be saying that we should adopt his methodological recommendations, because they have 'withstood severe criticism'. But his answer is inappropriate in this context because our aim is precisely to subject his philosophical views, in the best Popperian spirit, to severe criticism.

In my reply to Watkins, I said, 'Watkins acknowledges . . . that corroboration does have predictive import in practical decision making' (Lakatos [1968], p. 97). Popper has objected to this way of putting the matter:

> [*O*]*ur theories do have predictive import.* Our corroboration statements have no predictive import, although they motivate and justify our *preference* for some theory or other (Schilpp]1974], pp. 1029–30).

Let us grant that corroboration statements have no predictive *content*—indeed, that they are analytic, as Watkins remarks (Lakatos [1968], p. 63)—and that theories are the kinds of statements which do have predictive *content*. It does not follow, as Popper has claimed, that corroboration has no predictive *import*. The distinction between predictive content and predictive import is no mere verbal quibble; a fundamental substantive point is at issue. Statements whose consequences refer to future occurrences may be said to have predictive content; rules, imperatives, and directives are totally lacking in predictive content because they do not entail any statements at all. Nevertheless, an imperative—such as 'No smoking, please'—may have considerable predictive import, for it may effectively achieve the goal of preventing the occurrence of smoking in a particular room in the immediate future.

Since corroboration, in some cases at least, provides the basis for

deciding which theory (with its predictive content) is to be used for the purpose of making practical predictions, it seems to me that corroboration, even if it is lacking in predictive content, does have enormous predictive import. Perhaps this point can be put more clearly in the following way. *Statements* assessing the corroboration of theories have no predictive *content*, as Popper, Watkins *et al.* maintain. The *directive*—to choose more highly corroborated theories in preference to theories which are less well corroborated for purposes of practical prediction—has considerable predictive *import*. The problem, which it seems to me the anti-inductivists have failed to solve, is how to vindicate this directive for making predictions.[1] Without some sort of vindication for this directive, the problem of rational prediction remains unresolved.

I have wondered why it would seem evident to Popper that corroboration, as he construes it, should provide a guide to rational prediction. In his Autobiography, he gives what appear to be indications of an answer.

> I regarded (and I still regard) the degree of corroboration of a theory merely as a critical report on the quality of past performance: *it could not be used to predict future performance*. . . . When faced with the *need to act*, on one theory or another, the rational choice was to act on that theory—if there was one—which so far had stood up to criticism better than its competitors had: there is no better idea of rationality than that of a readiness to accept criticism. Accordingly, the degree of corroboration of a theory was a rational guide to practice (Schilpp [1974], p. 82).

A further elaboration of the theme informs us that

> . . . when we think we have found an approximation to the truth in the form of a scientific theory which has stood up to criticism and to tests better than its competitors, we shall, as realists, accept it as a basis for practical action, simply because we have nothing better (or nearer to the truth) (*ibid.*, pp. 120–1).

Realism is a position to which Popper has adhered since the time of his earliest philosophical activity; near the beginning of his Autobiography he tells us that 'a realist who believes in an "external world" necessarily believes in the existence of a cosmos rather than a chaos; that is, in regularities' (*ibid.*, p. 14). Thus, I am led to conjecture, it may be that Popper's adherence to the thesis that corroboration can provide a basis for rational prediction rests ultimately upon his realism, which embodies a version of a principle of uniformity of nature. If this suggestion is correct, we can still legitimately wonder whether Popper's epistemology is as far from traditional inductivism as he would have us believe.

To conclude this discussion, I should like to recall the point of my opening anecdote. It seems to me incorrect to suppose that the only concern of *theoretical* science is to make bold explanatory conjectures which can be tested and criticised. It is a mistake, I believe, to suppose that all prediction,

[1] This felicitous reformulation was suggested by Abner Shimony (if I did not misunderstand him) in the discussion following my presentation at the Popper Symposium.

aside from that involved in the testing of theories, is confined to contexts in which practical action is at stake. Theoretical science furnishes both explanations and predictions. Some of these predictions have practical consequences and others do not. When, for example, scientists assembled the first man-made atomic pile under the West Stands at the University of Chicago, they had to make a prediction as to whether the nuclear chain-reaction they initiated could be controlled, or whether it would spread to surrounding materials and engulf the entire city—and perhaps the whole earth—in a nuclear holocaust. Their predictions had both theoretical and practical interest. Contemporary cosmologists, for another example, would like to *explain* certain features of our universe in terms of its origin in a 'big bang'; many of them are trying to *predict* whether it will end in a 'big crunch'. In this case, the predictive question seems motivated by pure intellectual curiosity, quite unattached to concerns regarding practical decision-making. Whether a helium-filled balloon will move forward in the cabin of an airplane when the airplane accelerates, whether a nuclear chain reaction—once initiated—will run out of control, and whether the universe will eventually return to a state of high density are all matters of legitimate scientific concern.

In this paper, I have attempted to argue that pure deductivism could not do justice to the problem of rational prediction in contexts of practical decision-making. If we ask whether Popperian deductivism can adequately account for scientific predictions of the more theoretical varieties, then I suspect that we would have to go through all of the preceding arguments once more. The net result would be, I think, that science is inevitably inductive in matters of intellectual curiosity as well as practical prediction. It *may* be possible to excise all inductive ingredients from science, but if the operation were successful, the patient (science), deprived of all predictive import, would die.

University of Arizona

REFERENCES

GRÜNBAUM, A. [1976]: 'Is Falsifiability the Touchstone of Scientific Rationality? Karl Popper versus Inductivism', in R. S. Cohen, P. K. Feyerabend and M. W. Wartofsky (*eds.*), *Essays in Memory of Imre Lakatos*, Dordrecht: D. Reidel Publishing Co., pp. 213–52.

LAKATOS, I. (*ed.*) [1968]: *The Problem of Inductive Logic*. Amsterdam: North-Holland Publishing Co.

POPPER, K. R. [1972]: *Objective Knowledge*. Oxford: Clarendon Press.

POPPER, K. R. [1974]: 'Autobiography', in Schilpp (*ed.*) [1974], pp. 3–181.

POPPER, K. R. [1974*a*]: 'Replies to my Critics', in Schilpp (*ed.*) [1974], pp. 961–1197.

SALMON, W. C. [1967]: *The Foundations of Scientific Inference*. Pittsburgh: University of Piitsburgh Press. Originally published in R. C. Colodny (*ed.*): *Mind and Cosmos*, Pittsburgh: University of Pittsburgh Press [1966], pp. 135–275.

SALMON, W. C. [1968]: 'The Justification of Inductive Rules of Inference', in Lakatos [1968], pp. 24–43.

SALMON, W. C. [1968*a*]: 'Reply', in Lakatos [1968], pp. 74–97.

SCHILPP, P. A. (*ed.*) [1974]: *The Philosophy of Karl Popper*. LaSalle, Ill.: Open Court.

WATKINS, J. W. N. [1968]: 'Non-inductive corroboration', in Lakatos [1968], pp. 61–6.