

Logic of Discovery or Psychology of Research?¹

THOMAS S. KUHN
Princeton University

My object in these pages is to juxtapose the view of scientific development outlined in my book, *The Structure of Scientific Revolutions*, with the better known views of our chairman, Sir Karl Popper.² Ordinarily I should decline such an undertaking, for I am not so sanguine as Sir Karl about the utility of confrontations. Besides, I have admired his work for too long to turn critic easily at this date. Nevertheless, I am persuaded that for this occasion the attempt must be made. Even before my book was published two and a half years ago, I had begun to discover special and often puzzling characteristics of the relation between my views and his. That relation and the divergent reactions I have encountered to it suggest that a disciplined comparison of the two may produce peculiar enlightenment. Let me say why I think this could occur.

On almost all the occasions when we turn explicitly to the same problems, Sir Karl's view of science and my own are very nearly identical.³ We are both concerned with the dynamic process by which scientific knowledge is acquired rather than with the logical structure of the products of scientific research. Given that concern, both of us emphasize, as legitimate data, the facts and also the spirit of actual scientific life, and both of us turn often to history to find them. From this pool of shared data, we draw many of the same conclusions. Both of us reject the view that science progresses

¹ This paper was initially prepared at the invitation of P. A. Schilpp for his forthcoming volume, *The Philosophy of Karl R. Popper*, to be published by The Open Court Publishing Company, La Salle, Ill., in The Library of Living Philosophers. I am most grateful to both Professor Schilpp and the publishers for permission to print it as part of the proceedings of this symposium before its appearance in the volume for which it was first solicited.

² For purposes of the following discussion I have reviewed Sir Karl Popper's [1959], his [1963], and his [1957]. I have also occasionally referred to his original [1935] and his [1945]. My own [1962] provides a more extended account of many of the issues discussed below.

³ More than coincidence is presumably responsible for this extensive overlap. Though I had read none of Sir Karl's work before the appearance in 1959 of the English translation of his [1935] (by which time my book was in draft), I had repeatedly heard a number of his main ideas discussed. In particular, I had heard him discuss some of them as William James Lecturer at Harvard in the spring of 1950. These circumstances do not permit me to specify an intellectual debt to Sir Karl, but there must be one.

by accretion; both emphasize instead the revolutionary process by which an older theory is rejected and replaced by an incompatible new one¹; and both deeply underscore the role played in this process by the older theory's occasional failure to meet challenges posed by logic, experiment, or observation. Finally, Sir Karl and I are united in opposition to a number of classical positivism's most characteristic theses. We both emphasize, for example, the intimate and inevitable entanglement of scientific observation with scientific theory; we are correspondingly sceptical of efforts to produce any neutral observation language; and we both insist that scientists may properly aim to invent theories that *explain* observed phenomena and that do so in terms of *real* objects, whatever the latter phrase may mean.

That list, though it by no means exhausts the issues about which Sir Karl and I agree,² is already extensive enough to place us in the same minority among contemporary philosophers of science. Presumably that is why Sir Karl's followers have with some regularity provided my most sympathetic philosophical audience, one for which I continue to be grateful. But my gratitude is not unmixed. The same agreement that evokes the sympathy of this group too often misdirects its interest. Apparently Sir Karl's followers can often read much of my book as chapters from a late (and, for some, a drastic) revision of his classic, *The Logic of Scientific Discovery*. One of them asks whether the view of science outlined in my *Scientific Revolutions* has not long been common knowledge. A second, more charitably, isolates my originality as the demonstration that discoveries-of-fact have a life cycle very like that displayed by innovations-of-theory. Still others express general pleasure in the book but will discuss only the two comparatively secondary issues about which my disagreement with Sir Karl is most nearly explicit: my emphasis on the importance of deep commitment to tradition and my discontent with the implications of the term 'falsification'. All these men, in short, read my book through a quite special pair of spectacles, and there is another way to read it. The view through those spectacles is not wrong—my agreement with Sir Karl is real and substantial. Yet readers outside of the Popperian circle almost

¹ Elsewhere I use the term 'paradigm' rather than 'theory' to denote what is rejected and replaced during scientific revolutions. Some reasons for the change of term will emerge below.

² Underlining one additional area of agreement about which there has been much misunderstanding may further highlight what I take to be the real differences between Sir Karl's views and mine. We both insist that adherence to a tradition has an essential role in scientific development. He has written, for example, 'Quantitatively and qualitatively by far the most important source of our knowledge—apart from inborn knowledge—is tradition' (Popper [1963], p. 27). Even more to the point, as early as 1948 Sir Karl wrote, 'I do not think that we could ever free ourselves entirely from the bonds of tradition. The so-called freeing is really only a change from one tradition to another' ([1963], p. 122).

invariably fail even to notice that the agreement exists, and it is these readers who most often recognize (not necessarily with sympathy) what seem to me the central issues. I conclude that a gestalt switch divides readers of my book into two or more groups. What one of these sees as striking parallelism is virtually invisible to the others. The desire to understand how this can be so motivates the present comparison of my view with Sir Karl's.

The comparison must not, however, be a mere point by point juxtaposition. What demands attention is not so much the peripheral area in which our occasional secondary disagreements are to be isolated but the central region in which we appear to agree. Sir Karl and I do appeal to the same data; to an uncommon extent we are seeing the same lines on the same paper; asked about those lines and those data, we often give virtually identical responses, or at least responses that inevitably seem identical in the isolation enforced by the question-and-answer mode. Nevertheless, experiences like those mentioned above convince me that our intentions are often quite different when we say the same things. Though the lines are the same, the figures which emerge from them are not. That is why I call what separates us a gestalt switch rather than a disagreement and also why I am at once perplexed and intrigued about how best to explore the separation. How am I to persuade Sir Karl, who knows everything I know about scientific development and who has somewhere or other said it, that what he calls a duck can be seen as a rabbit? How am I to show him what it would be like to wear my spectacles when he has already learned to look at everything I can point to through his own?

In this situation a change in strategy is called for, and the following suggests itself. Reading over once more a number of Sir Karl's principal books and essays, I encounter again a series of recurrent phrases which, though I understand them and do not quite disagree, are locutions that I could never have used in the same places. Undoubtedly they are most often intended as metaphors applied rhetorically to situations for which Sir Karl has elsewhere provided unexceptionable descriptions. Nevertheless, for present purposes these metaphors, which strike me as patently inappropriate, may prove more useful than straightforward descriptions. They may that is, be symptomatic of contextual differences that a careful literal expression hides. If that is so, then these locutions may function not as the lines-on-paper but as the rabbit-ear, the shawl, or the ribbon-at-the-throat which one isolates when teaching a friend to transform his way of seeing a gestalt diagram. That, at least, is my hope for them. I have four such differences of locutions in mind and shall treat them *seriatim*.

I

Among the most fundamental issues on which Sir Karl and I agree is our insistence that an analysis of the development of scientific knowledge must take account of the way science has actually been practiced. That being so, a few of his recurrent generalizations startle me. One of these provides the opening sentences of the first chapter of the *Logic of Scientific Discovery*: 'A scientist', writes Sir Karl, 'whether theorist or experimenter, puts forward statements, or systems of statements, and tests them step by step. In the field of the empirical sciences, more particularly, he constructs hypotheses, or systems of theories, and tests them against experience by observation and experiment.'¹ The statement is virtually a cliché, yet in application it presents three problems. It is ambiguous in its failure to specify which of two sorts of 'statements' or 'theories' are being tested. That ambiguity can, it is true, be eliminated by reference to other passages in Sir Karl's writings, but the generalization that results is historically mistaken. Furthermore, the mistake proves important, for the unambiguous form of the description misses just that characteristic of scientific practice which most nearly distinguishes the sciences from other creative pursuits.

There is one sort of 'statement' or 'hypothesis' that scientists do repeatedly subject to systematic test. I have in mind statements of an individual's best guesses about the proper way to connect his own research problem with the corpus of accepted scientific knowledge. He may, for example, conjecture that a given chemical unknown contains the salt of a rare earth, that the obesity of his experimental rats is due to a specified component in their diet, or that a newly discovered spectral pattern is to be understood as an effect of nuclear spin. In each case, the next steps in his research are intended to try out or test the conjecture or hypothesis. If it passes enough or stringent enough tests, the scientist has made a discovery or has at least resolved the puzzle he had been set. If not, he must either abandon the puzzle entirely or attempt to solve it with the aid of some other hypothesis. Many research problems, though by no means all, take this form. Tests of this sort are a standard component of what I have elsewhere labelled 'normal science' or 'normal research', an enterprise which accounts for the overwhelming majority of the work done in basic science. In no usual sense, however, are such tests directed to current theory. On the contrary, when engaged with a normal research problem, the scientist must *premise* current theory as the rules of his game. His object is to solve a puzzle, preferably one at which others have failed, and current theory is required to

¹ Popper [1959], p. 27.

define that puzzle and to guarantee that, given sufficient brilliance, it can be solved.¹ Of course the practitioner of such an enterprise must often test the conjectural puzzle solution that his ingenuity suggests. But only his personal conjecture is tested. If it fails the test, only his own ability not the corpus of current science is impugned. In short, though tests occur frequently in normal science, these tests are of a peculiar sort, for in the final analysis it is the individual scientist rather than current theory which is tested.

This is not, however, the sort of test Sir Karl has in mind. He is above all concerned with the procedures through which science grows, and he is convinced that 'growth' occurs not primarily by accretion but by the revolutionary overthrow of an accepted theory and its replacement by a better one.² (The subsumption under 'growth' of 'repeated overthrow' is itself a linguistic oddity whose *raison d'être* may become more visible as we proceed.) Taking this view, the tests which Sir Karl emphasizes are those which were performed to explore the limitations of accepted theory or to subject a current theory to maximum strain. Among his favourite examples, all of them startling and destructive in their outcome, are Lavoisier's experiments on calcination, the eclipse expedition of 1919, and the recent experiments on parity conservation.³ All, of course, are classic tests, but in using them to characterize scientific activity Sir Karl misses something terribly important about them. Episodes like these are very rare in the development of science. When they occur, they are generally called forth either by a prior crisis in the relevant field (Lavoisier's experiments or Lee and Yang's⁴) or by the existence of a theory which competes with the existing canons of research (Einstein's general relativity). These are, however, aspects of or occasions for what I have elsewhere called 'extraordinary research', an enterprise in which scientists do display

¹ For an extended discussion of normal science, the activity which practitioners are trained to carry on, see my [1962], pp. 23-42, and 135-42. It is important to notice that when I describe the scientist as a puzzle solver and Sir Karl describes him as a problem solver (e.g. in his [1963], pp. 67, 222), the similarity of our terms disguises a fundamental divergence. Sir Karl writes (the italics are his), 'Admittedly, our expectations, and thus our theories, may precede, historically, even our problems. *Yet science starts only with problems.* Problems crop up especially when we are disappointed in our expectations, or when our theories involve us in difficulties, in contradictions'. I use the term 'puzzle' in order to emphasize that the difficulties which *ordinarily* confront even the very best scientists are, like crossword puzzles or chess puzzles, challenges only to his ingenuity. *He* is in difficulty, not current theory. My point is almost the converse of Sir Karl's.

² Cf. Popper [1963], pp. 129, 215 and 221, for particularly forceful statements of this position.

³ For example, Popper [1963], p. 220.

⁴ For the work on calcination see, Guerlac [1961]. For the background of the parity experiments see, Hafner and Presswood [1965].

very many of the characteristics Sir Karl emphasizes, but one which, at least in the past, has arisen only intermittently and under quite special circumstances in any scientific speciality.¹

I suggest then that Sir Karl has characterized the entire scientific enterprise in terms that apply only to its occasional revolutionary parts. His emphasis is natural and common: the exploits of a Copernicus or Einstein make better reading than those of a Brahe or Lorentz; Sir Karl would not be the first if he mistook what I call normal science for an intrinsically uninteresting enterprise. Nevertheless, neither science nor the development of knowledge is likely to be understood if research is viewed exclusively through the revolutions it occasionally produces. For example, though testing of basic commitments occurs only in extraordinary science, it is normal science that discloses both the points to test and the manner of testing. Or again, it is for the normal, not the extraordinary practice of science that professionals are trained; if they are nevertheless eminently successful in displacing and replacing the theories on which normal practice depends, that is an oddity which must be explained. Finally, and this is for now my main point, a careful look at the scientific enterprise suggests that it is normal science, in which Sir Karl's sort of testing does not occur, rather than extraordinary science which most nearly distinguishes science from other enterprises. If a demarcation criterion exists (we must not, I think, seek a sharp or decisive one), it may lie just in that part of science which Sir Karl ignores.

In one of his most evocative essays, Sir Karl traces the origin of 'the tradition of critical discussion [which] represents the only practicable way of expanding our knowledge' to the Greek philosophers between Thales and Plato, the men who, as he sees it, encouraged critical discussion both between schools and within individual schools.² The accompanying description of Presocratic discourse is most apt, but what is described does not at all resemble science. Rather it is the tradition of claims, counter-claims, and debates over fundamentals which, except perhaps during the Middle Ages, have characterized philosophy and much of social science ever since. Already by the Hellenistic period mathematics, astronomy, statics and the geometric parts of optics had abandoned this mode of discourse in favour of puzzle solving. Other sciences, in increasing numbers, have undergone the same transition since. In a sense, to turn Sir Karl's view on its head, it is precisely the abandonment of critical discourse that marks the transition to a science. Once a field has made that transition, critical discourse recurs only at moments of crisis when the bases of the

¹ The point is argued at length in my [1962], pp. 52-97.

² Popper [1963], chapter 5, especially pp. 148-52.

field are again in jeopardy.¹ Only when they must choose between competing theories do scientists behave like philosophers. That, I think, is why Sir Karl's brilliant description of the reasons for the choice between metaphysical systems so closely resembles my description of the reasons for choosing between scientific theories.² In neither choice, as I shall shortly try to show, can testing play a quite decisive role.

There is, however, good reason why testing has seemed to do so, and in exploring it Sir Karl's duck may at last become my rabbit. No puzzle-solving enterprise can exist unless its practitioners share criteria which, for that group and for that time, determine when a particular puzzle has been solved. The same criteria necessarily determine failure to achieve a solution, and anyone who chooses may view that failure as the failure of a theory to pass a test. Normally, as I have already insisted, it is not viewed that way. Only the practitioner is blamed, not his tools. But under the special circumstances which induce a crisis in the profession (e.g. gross failure, or repeated failure by the most brilliant professionals) the group's opinion may change. A failure that had previously been personal may then come to seem the failure of a theory under test. Thereafter, because the test arose from a puzzle and thus carried settled criteria of solution, it proves both more severe and harder to evade than the tests available within a tradition whose normal mode is critical discourse rather than puzzle solving.

In a sense, therefore, severity of test-criteria is simply one side of the coin whose other face is a puzzle-solving tradition. That is why Sir Karl's line of demarcation and my own so frequently coincide. That coincidence is, however, only in their *outcome*; the *process* of applying them is very different, and it isolates distinct aspects of the activity about which the decision—science or non-science—is to be made. Examining the vexing cases, for example, psychoanalysis or Marxist historiography, for which Sir Karl tells us his criterion was initially designed,³ I concur that they cannot now properly be labelled 'science'. But I reach that conclusion by a route far surer and more direct than his. One brief example may suggest that of the two criteria, testing and puzzle solving, the latter is at once the less equivocal and the more fundamental.

To avoid irrelevant contemporary controversies, I consider astrology rather than, say, psychoanalysis. Astrology is Sir Karl's most frequently cited example of a 'pseudo-science'.⁴ He says: 'By making their interpretations and prophecies sufficiently vague they [astrologers] were able to

¹ Though I was not then seeking a demarcation criterion, just these points are argued at length in my [1962], pp. 10-22 and 87-90.

² Cf. Popper [1963], pp. 192-200, with my [1962], pp. 143-58. ³ Popper [1963], p. 34.

⁴ The index to Popper [1963] has eight entries under the heading 'astrology as a typical pseudo science'.

explain away anything that might have been a refutation of the theory had the theory and the prophecies been more precise. In order to escape falsification they destroyed the testability of the theory.¹ Those generalizations catch something of the spirit of the astrological enterprise. But taken at all literally, as they must be if they are to provide a demarcation criterion, they are impossible to support. The history of astrology during the centuries when it was intellectually reputable records many predictions that categorically failed.² Not even astrology's most convinced and vehement exponents doubted the recurrence of such failures. Astrology cannot be barred from the sciences because of the form in which its predictions were cast.

Nor can it be barred because of the way its practitioners explained failure. Astrologers pointed out, for example, that, unlike general predictions about, say, an individual's propensities or a natural calamity, the forecast of an individual's future was an immensely complex task, demanding the utmost skill, and extremely sensitive to minor errors in relevant data. The configuration of the stars and eight planets was constantly changing; the astronomical tables used to compute the configuration at an individual's birth were notoriously imperfect; few men knew the instant of their birth with the requisite precision.³ No wonder, then, that forecasts often failed. Only after astrology itself became implausible did these arguments come to seem question-begging.⁴ Similar arguments are regularly used today when explaining, for example, failures in medicine or meteorology. In times of trouble they are also deployed in the exact sciences, fields like physics, chemistry, and astronomy.⁵ There was nothing unscientific about the astrologer's explanation of failure.

Nevertheless, astrology was not a science. Instead it was a craft, one of the practical arts, with close resemblances to engineering, meteorology, and medicine as these fields were practised until little more than a century ago. The parallels to an older medicine and to contemporary psychoanalysis are, I think, particularly close. In each of these fields shared theory was adequate only to establish the plausibility of the discipline and to provide a rationale for the various craft-rules which governed practice. These rules had proved their use in the past, but no practitioner supposed they were sufficient to prevent recurrent failure. A more articulated theory and more powerful rules were desired, but it would have been absurd to

¹ Popper [1963], p. 37.

² For examples see, Thorndike [1923-58], 5, pp. 225 ff.; 6, pp. 71, 101, 114.

³ For reiterated explanations of failure see, *ibid.* 1, pp. 11 and 514 f.; 4, 368; 5, 279.

⁴ A perceptive account of some reasons for astrology's loss of plausibility is included in Stahlman [1956]. For an explanation of astrology's previous appeal see, Thorndike [1955].

⁵ Cf. my [1962], pp. 66-76.

abandon a plausible and badly needed discipline with a tradition of limited success simply because these desiderata were not yet at hand. In their absence, however, neither the astrologer nor the doctor could do research. Though they had rules to apply, they had no puzzles to solve and therefore no science to practise.¹

Compare the situations of the astronomer and the astrologer. If an astronomer's prediction failed and his calculations checked, he could hope to set the situation right. Perhaps the data were at fault: old observations could be re-examined and new measurements made, tasks which posed a host of calculational and instrumental puzzles. Or perhaps theory needed adjustment, either by the manipulation of epicycles, eccentrics, equants, etc., or by more fundamental reforms of astronomical technique. For more than a millennium these were the theoretical and mathematical puzzles around which, together with their instrumental counterparts, the astronomical research tradition was constituted. The astrologer, by contrast, had no such puzzles. The occurrence of failures could be explained, but particular failures did not give rise to research puzzles, for no man, however skilled, could make use of them in a constructive attempt to revise the astrological tradition. There were too many possible sources of difficulty, most of them beyond the astrologer's knowledge, control, or responsibility. Individual failures were correspondingly uninformative, and they did not reflect on the competence of the prognosticator in the eyes of his professional compeers.² Though astronomy and astrology were regularly practised by the same people, including Ptolemy, Kepler, and Tycho Brahe, there was never an astrological equivalent of the puzzle-solving astronomical tradition. And without puzzles, able first to challenge and then to attest the ingenuity of the individual practitioner, astrology could

¹ This formulation suggests that Sir Karl's criterion of demarcation might be saved by a minor restatement entirely in keeping with his apparent intent. For a field to be a science its conclusions must be *logically derivable* from *shared premises*. On this view astrology is to be barred not because its forecasts were not testable but because only the most general and least testable ones could be derived from accepted theory. Since any field that did satisfy this condition *might* support a puzzle solving tradition, the suggestion is clearly helpful. It comes close to supplying a sufficient condition for a field's being a science. But in this form, at least, it is not even quite a sufficient condition, and it is surely not a necessary one. It would, for example, admit surveying and navigation as sciences, and it would bar taxonomy, historical geology, and the theory of evolution. The conclusions of a science may be both precise and binding without being fully derivable by logic from accepted premises. Cf. my [1962], pp. 35-51, and also the discussion in Section III, *below*.

² This is not to suggest that astrologers did not criticize each other. On the contrary, like practitioners of philosophy and some social sciences, they belonged to a variety of different schools, and the inter-school strife was sometimes bitter. But these debates ordinarily revolved about the *implausibility* of the particular theory employed by one or another school. Failures of individual predictions played very little role. Compare Thorndike [1923-58], 5, p. 233.

not have become a science even if the stars had, in fact, controlled human destiny.

In short, though astrologers made testable predictions and recognized that these predictions sometimes failed, they did not and could not engage in the sorts of activities that normally characterize all recognized sciences. Sir Karl is right to exclude astrology from the sciences, but his over-concentration on science's occasional revolutions prevents his seeing the surest reason for doing so.

That fact, in turn, may explain another oddity of Sir Karl's historiography. Though he repeatedly underlines the role of tests in the replacement of scientific theories, he is also constrained to recognize that many theories, for example the Ptolemaic, were replaced before they had in fact been tested.¹ On some occasions, at least, tests are not requisite to the revolutions through which science advances. But that is not true of puzzles. Though the theories Sir Karl cites had not been put to the test before their displacement, none of these was replaced before it had ceased adequately to support a puzzle-solving tradition. The state of astronomy was a scandal in the early sixteenth century. Most astronomers nevertheless felt that normal adjustments of a basically Ptolemaic model would set the situation right. In this sense the theory had not failed a test. But a few astronomers, Copernicus among them, felt that the difficulties must lie in the Ptolemaic approach itself rather than in the particular versions of Ptolemaic theory so far developed, and the results of that conviction are already recorded. The situation is typical.² With or without tests, a puzzle-solving tradition can prepare the way for its own displacement. To rely on testing as the mark of a science is to miss what scientists mostly do and, with it, the most characteristic feature of their enterprise.

II

With the background supplied by the preceding remarks we can quickly discover the occasion and consequences of another of Sir Karl's favourite locutions. The preface to *Conjectures and Refutations* opens with the sentence: 'The essays and lectures of which this book is composed, are variations upon one very simple theme—the thesis that *we can learn from our mistakes*.' The emphasis is Sir Karl's; the thesis recurs in his writing from an early date³; taken in isolation, it inevitably commands assent. Everyone

¹ Cf. Popper [1963], p. 246.

² Cf. my [1962], pp. 77–87.

³ The quotation is from Popper [1963], p. vii, in a preface dated 1962. Earlier Sir Karl had equated 'learning from our mistakes' with 'learning by trial and error' ([1963], p. 216), and the trial-and-error formulation dates from at least 1937 ([1963], p. 312) and is in spirit older than that. Much of what is said below about Sir Karl's notion of 'mistake' applies equally to his concept of 'error'.

can and does learn from his mistakes; isolating and correcting them is an essential technique in teaching children. Sir Karl's rhetoric has roots in everyday experience. Nevertheless, in the contexts for which he invokes this familiar imperative, its applications seems decisively askew. I am not sure a mistake has been made, at least not a mistake to learn from.

One need not confront the deeper philosophical problems presented by mistakes to see what is presently at issue. It is a mistake to add three plus three and get five, or to conclude from 'All men are mortal' to 'All mortals are men'. For different reasons, it is a mistake to say, 'He is my sister', or to report the presence of a strong electric field when test charges fail to indicate it. Presumably there are still other sorts of mistakes, but all the normal ones are likely to share the following characteristics. A mistake is made, or is committed, at a specifiable time and place by a particular individual. That individual has failed to obey some established rule of logic, or of language, or of the relations between one of these and experience. Or he may instead have failed to recognize the consequences of a particular choice among the alternatives which the rules allow him. The individual can learn from his mistake only because the group whose practice embodies these rules can isolate the individual's failure in applying them. In short, the sorts of mistakes to which Sir Karl's imperative most obviously applies are in individual's failure of understanding or of recognition within an activity governed by pre-established rules. In the sciences such mistakes occur most frequently and perhaps exclusively within the practice of normal puzzle-solving research.

That is not, however, where Sir Karl seeks them, for his concept of science obscures even the existence of normal research. Instead, he looks to the extraordinary or revolutionary episodes in scientific development. The mistakes to which he points are not usually acts at all but rather out-of-date scientific theories: Ptolemaic astronomy, the phlogiston theory, or Newtonian dynamics, and 'learning from our mistakes' is, correspondingly, what occurs when a scientific community rejects one of these theories and replaces it with another.¹ If this does not immediately seem an odd usage,

¹ Popper [1963], pp. 215 and 220. In these pages Sir Karl outlines and illustrates his thesis that science grows through revolutions. He does not, in the process, ever juxtapose the term 'mistake' with the name of an out-of-date scientific theory, presumably because his sound historic instinct inhibits so gross an anachronism. Yet the anachronism is fundamental to Sir Karl's rhetoric, which does repeatedly provide clues to more substantial differences between us. Unless out-of-date theories are mistakes, there is no way to reconcile, say, the opening paragraph of Sir Karl's preface ([1963], p. vii: 'learn from our mistakes'; 'our often mistaken attempts to solve our problems'; 'tests which may help us in the discovery of our mistakes') with the view ([1963], p. 215) that 'the growth of scientific knowledge . . . [consists in] the repeated overthrow of scientific theories and their replacement by better or more satisfactory ones'.

that is mainly because it appeals to the residual inductivist in us all. Believing that valid theories are the product of correct inductions from facts, the inductivist must also hold that a false theory is the result of a mistake in induction. In principle, at least, he is prepared to answer the questions: what mistake was made, what rule broken, when and by whom, in arriving at, say, the Ptolemaic system? To the man for whom those are sensible questions and to him alone, Sir Karl's locution presents no problems.

But neither Sir Karl nor I is an inductivist. We do not believe that there are rules for inducing correct theories from facts, or even that theories, correct or incorrect, are induced at all. Instead we view them as imaginative posits, invented in one piece for application to nature. And though we point out that such posits can and usually do at last encounter puzzles they cannot solve, we also recognize that those troublesome confrontations rarely occur for some time after a theory has been both invented and accepted. In our view, then, no mistake was made in arriving at the Ptolemaic system, and it is therefore difficult for me to understand what Sir Karl has in mind when he calls that system, or any other out-of-date theory, a mistake. At most one may wish to say that a theory which was not previously a mistake has become one or that a scientist has made the mistake of clinging to a theory for too long. And even these locutions, of which at least the first is extremely awkward, do not return us to the sense of mistake with which we are most familiar. Those mistakes are the normal ones which a Ptolemaic (or a Copernican) astronomer makes within his system, perhaps in observation, calculation, or the analysis of data. They are, that is, the sort of mistake which can be isolated and then at once corrected, leaving the original system intact. In Sir Karl's sense, on the other hand, a mistake infects an entire system and can be corrected only by replacing the system as a whole. No locutions and no similarities can disguise these fundamental differences, nor can it hide the fact that before infection set in the system had the full integrity of what we now call sound knowledge.

Quite possibly Sir Karl's sense of 'mistake' can be salvaged, but a successful salvage operation must deprive it of certain still current implications. Like the term 'testing', 'mistake' has been borrowed from normal science, where its use is reasonably clear, and applied to revolutionary episodes, where its application is at best problematic. That transfer creates, or at least reinforces, the prevalent impression that whole theories can be judged by the same sort of criteria that one employs when judging a theory's individual research applications. The discovery of applicable criteria then becomes a primary desideratum for many people. That Sir

Karl should be among them is strange, for the search runs counter to the most original and fruitful thrust in his philosophy of science. But I can understand his methodological writings since the *Logik der Forschung* in no other way. I shall now suggest that he has, despite explicit disclaimers, consistently sought evaluation procedures which can be applied to theories with the apodictic assurance characteristic of the techniques by which one identifies mistakes in arithmetic, logic, or measurement. I fear that he is pursuing a will-o'-the-wisp born from the same conjunction of normal and extraordinary science which made tests seem so fundamental a feature of the sciences.

III

In his *Logik der Forschung*, Sir Karl underlined the asymmetry of a generalization and its negation in their relation to empirical evidence. A scientific theory cannot be shown to apply successfully to all its possible instances, but it can be shown to be unsuccessful in particular applications. Emphasis upon that logical truism and its implications seems to me a forward step from which there must be no retreat. The same asymmetry plays a fundamental role in my *Structure of Scientific Revolutions*, where a theory's failure to provide rules that identify solvable puzzles is viewed as the source of professional crises which often result in the theory's being replaced. My point is very close to Sir Karl's, and I may well have taken it from what I had heard of his work.

But Sir Karl describes as 'falsification' or 'refutation' what happens when a theory fails in an attempted application, and these are the first of a series of related locutions that again strike me as extremely odd. Both 'falsification' and 'refutation' are antonyms of 'proof'. They are drawn principally from logic and from formal mathematics; the chains of argument to which they apply end with a 'Q.E.D.'; invoking these terms implies the ability to compel assent from any member of the relevant professional community. No member of this audience, however, still needs to be told that, where a whole theory or often even a scientific law is at stake, arguments are seldom so apodictic. All experiments can be challenged, either as to their relevance or their accuracy. All theories can be modified by a variety of *ad hoc* adjustments without ceasing to be, in their main lines, the same theories. It is important, furthermore, that this should be so, for it is often by challenging observations or adjusting theories that scientific knowledge grows. Challenges and adjustments are a standard part of normal research in empirical science, and adjustments, at least, play a dominant role in informal mathematics as well. Dr Lakatos's brilliant analysis of the permissible rejoinders to mathematical refutations

provides the most telling arguments I know against a naive falsificationist position.¹

Sir Karl is not, of course, a naive falsificationist. He knows all that has just been said and has emphasized it from the beginning of his career. Very early in his *Logic of Scientific Discovery*, for example, he writes: 'In point of fact, no conclusive disproof of a theory can ever be produced; for it is always possible to say that the experimental results are not reliable or that the discrepancies which are asserted to exist between the experimental results and the theory are only apparent and that they will disappear with the advance of our understanding.'² Statements like these display one more parallel between Sir Karl's view of science and my own, but what we make of them could scarcely be more different. For my view they are fundamental, both as evidence and as source. For Sir Karl's, in contrast, they are an essential qualification which threatens the integrity of his basic position. Having barred conclusive disproof, he has provided no substitute for it, and the relation he does employ remains that of logical falsification. Though he is not a naive falsificationist, Sir Karl may, I suggest, legitimately be treated as one.

If his concern were exclusively with demarcation, the problems posed by the unavailability of conclusive disproofs would be less severe and perhaps eliminable. Demarcation might, that is, be achieved by an exclusively syntactic criterion.³ Sir Karl's view would then be, and perhaps is, that a theory is scientific if and only if *observation statements*—particularly the negations of singular existential statements—can be logically deduced from it, perhaps in conjunction with stated background knowledge. The difficulties (to which I shall shortly turn) in deciding whether the outcome of a particular laboratory operation justifies asserting a particular observation statement would then be irrelevant. Perhaps, though the basis for doing so is less apparent, the equally grave difficulties in deciding whether an observation statement deduced from an approximate (e.g. mathematically manageable) version of the theory should be considered consequences of the theory itself could be eliminated in the same way. Problems like these would belong not to the syntactics but to the pragmatics or semantics of the language in which the theory was cast, and they would therefore have no role in determining its status as a science. To be scientific a theory need be falsifiable only by an observation statement not by actual observation. The relation between statements, unlike that between

¹ Lakatos [1963-4].

² Popper [1959], p. 50.

³ Though my point is somewhat different, I owe my recognition of the need to confront this issue to C. G. Hempel's strictures on those who misinterpret Sir Karl by attributing to him a belief in absolute rather than relative falsification. See his [1965], p. 45. I am also indebted to Professor Hempel for a close and perceptive critique of this paper in draft.

a statement and an observation, could be the conclusive disproof familiar from logic and mathematics.

For reasons suggested above (p. 9, footnote 1) and elaborated immediately below, I doubt that scientific theories can without decisive change be cast in a form which permits the purely syntactic judgements which this version of Sir Karl's criterion requires. But even if they could, these reconstructed theories would provide a basis only for his demarcation criterion, not for the logic of knowledge so closely associated with it. The latter has, however, been Sir Karl's most persistent concern, and his notion of it is quite precise. 'The logic of knowledge . . .,' he writes, 'consists solely in investigating the methods employed in those systematic tests to which every new idea must be subjected if it is to be seriously entertained.'¹ From this investigation, he continues, result methodological rules or conventions like the following: 'Once a hypothesis has been proposed and tested, and has proved its mettle, it may not be allowed to drop out without "good reason". A "good reason" may be, for instance . . . the falsification of one of the consequences of the hypothesis.'²

Rules like these, and with them the entire logical enterprise described above, are no longer simply syntactic in their import. They require that both the epistemological investigator and the research scientist be able to relate sentences derived from a theory not to other sentences but to actual observations and experiments. This is the context in which Sir Karl's term 'falsification' must function, and Sir Karl is entirely silent about how it can do so. What is falsification if it is not conclusive disproof? Under what circumstances does the *logic* of knowledge require a scientist to abandon a previously accepted theory when confronted, not with statements about experiments, but with experiments themselves? Pending clarification of these questions, I am not clear that what Sir Karl has given us is a logic of knowledge at all. In my conclusion I shall suggest that, though equally valuable, it is something else entirely. Rather than a logic, Sir Karl has provided an ideology; rather than methodological rules, he has supplied procedural maxims.

That conclusion must, however, be postponed until after a last deeper look at the source of the difficulties with Sir Karl's notion of falsification. It presupposes, as I have already suggested, that a theory is cast, or can without distortion be recast, in a form which permits scientists to classify each conceivable event as either a confirming instance, a falsifying instance, or irrelevant to the theory. That is obviously required if a general law is to be falsifiable: to test the generalization $(x) \phi(x)$ by applying it to the constant a , we must be able to tell whether or not a lies within the

¹ Popper [1959], p. 31.

² Popper [1959], pp. 53 f.

range of the variable x and whether or not $\phi(a)$. The same presupposition is even more apparent in Sir Karl's recently elaborated measure of verisimilitude. It requires that we first produce the class of all logical consequences of the theory and then choose from among these, with the aid of background knowledge, the classes of all true and of all false consequences.¹ At least, we must do this if the criterion of verisimilitude is to result in a *method* of theory choice. None of these tasks can, however, be accomplished unless the theory is fully articulated logically and unless the terms through which it attaches to nature are sufficiently defined to determine their applicability in each possible case. In practice, however, no scientific theory satisfies these rigorous demands, and many people have argued that a theory would cease to be useful in research if it did so.² I have myself elsewhere introduced the term 'paradigm' to underscore the dependence of scientific research upon concrete examples that bridge what would otherwise be gaps in the specification of the content and application of scientific theories. The relevant arguments cannot be repeated here. But a brief example, though it will temporarily alter my mode of discourse, may be even more useful.

My example takes the form of a constructed epitome of some elementary scientific knowledge. That knowledge concerns swans, and to isolate its presently relevant characteristics I shall ask three questions about it: (a) How much can one know about swans without introducing explicit generalizations like 'All swans are white'? (b) Under what circumstances and with what consequences are such generalizations worth adding to what was known without them? (c) Under what circumstances are generalizations rejected once they have been made? In raising these questions my object is to suggest that, though logic is a powerful and ultimately an essential tool of scientific enquiry, one can have sound knowledge in forms to which logic can scarcely be applied. Simultaneously, I shall suggest that logical articulation is not a value for its own sake, but is to be undertaken only when and to the extent that circumstances demand it.

Imagine that you have been shown and can remember ten birds which have authoritatively been identified as swans; that you have a similar acquaintance with ducks, geese, pigeons, doves, gulls, etc.; and that you are informed that each of these types constitutes a natural family. A natural family you already know as an observed cluster of like objects,

¹ Popper [1963], pp. 233-5. Notice also, at the foot of the last of these pages, that Sir Karl's comparison of the relative verisimilitude of two theories depends upon there being 'no revolutionary changes in our background knowledge', an assumption which he nowhere argues and which is hard to reconcile with his conception of scientific change by revolutions.

² Braithwaite [1953], pp. 50-87, especially p. 76, and my [1962], pp. 97-101.

sufficiently important and sufficiently discrete to command a generic name. More precisely, though here I introduce more simplification than the concept requires, a natural family is a class whose members resemble each other more closely than they resemble the members of other natural families.¹ The experience of generations has to date confirmed that all observed objects fall into one or another natural family. It has, that is, shown that the entire population of the world can always be divided (though not once and for all) into perceptually discontinuous categories. In the perceptual spaces between these categories there are believed to be no objects at all.

What you have learned about swans from exposure to paradigms is very much like what children first learn about dogs and cats, tables and chairs, mothers and fathers. Its precise scope and content are, of course, impossible to specify, but it is sound knowledge nonetheless. Derived from observation, it can be infirmed by further observation, and it meanwhile provides a basis for rational action. Seeing a bird much like the swans you already know, you may reasonably presume that it will require the same food as the others and will breed with them. Provided swans are a natural family, no bird which closely resembles them on sight should display radically different characteristics on closer acquaintance. Of course you may have been misinformed about the natural integrity of the swan family. But that can be discovered from experience, for example, by the discovery of a number of animals (note that more than one is required) whose characteristics bridge the gap between swans and, say, geese by barely perceptible intervals.² Until that does occur, however, you will know a great deal about swans though you will not be altogether sure what you know or what a swan is.

Suppose now that all the swans you have actually observed are white. Should you embrace the generalization, 'All swans are white'? Doing so will change what you know very little; that change will be of use only in the unlikely event that you meet a non-white bird which otherwise resembles a swan; by making the change you increase the risk that the swan

¹ Note that the resemblance between members of a natural family is here a learned relationship and one which can be unlearned. Contemplate the old saw, 'To an occidental, all chinamen look alike'. That example also highlights the most drastic of the simplifications introduced at this point. A fuller discussion would have to allow for hierarchies of natural families with resemblance relations between families at the higher levels.

² This experience would not necessitate the abandonment of either the category 'swans' or the category 'geese', but it would necessitate the introduction of an *arbitrary* boundary between them. The families 'swans' and 'geese' would no longer be natural families, and you could exclude nothing about the character of a new swan-like bird that was not also true of geese. Empty perceptual space is essential if family membership is to have cognitive content.

family will prove not to be a natural family after all. Under those circumstances you are likely to refrain from generalizing unless there are special reasons for doing so. Perhaps, for example, you must describe swans to men who cannot be directly exposed to paradigms. Without superhuman caution both on your part and on that of your readers, your description will acquire the force of a generalization; this is often the problem of the taxonomist. Or perhaps you have discovered some grey birds that look otherwise like swans but eat different food and have an unfortunate disposition. You may then generalize to avoid a behavioural mistake. Or you may have a more theoretical reason for thinking the generalization worthwhile. For example, you may have observed that the members of other natural families share colouration. Specifying this fact in a form which permits the application of powerful logical techniques to what you know may enable you to learn more about the animal colour in general or about animal breeding.

Now, having made the generalization, what will you do if you encounter a black bird that looks otherwise like a swan? Almost the same things, I suggest, as if you had not previously committed yourself to the generalization at all. You will examine the bird with care, externally and perhaps internally as well, to find other characteristics that distinguish this specimen from your paradigms. That examination will be particularly long and thorough if you have theoretical reasons for believing that colour characterizes natural families or if you are deeply ego involved with the generalization. Very likely the examination will disclose other differentiae, and you will announce the discovery of a new natural family. Or you may fail to find such differentiae and may then announce that a black swan has been found. Observation cannot, however, force you to that falsifying conclusion, and you would occasionally be the loser if it could do so. Theoretical considerations may suggest that colour alone is sufficient to demarcate a natural family: the bird is not a swan because it is black. Or you may simply postpone the issue pending the discovery and examination of other specimens. Only if you have previously committed yourself to a full definition of 'swan', one which will specify its applicability to every conceivable object, can you be logically *forced* to rescind your generalization.¹ And why should you have offered such a definition? It could serve no cognitive function and would

¹ Further evidence for the unnaturalness of any such definition is provided by the following question. Should 'whiteness' be included as a defining characteristic of swans? If so, the generalization 'All swans are white' is immune to experience. But if 'whiteness' is excluded from the definition, then some other characteristic must be included for which 'whiteness' might have substituted. Decisions about which characteristics are to be parts of a definition and which are to be available for the statement of general laws are often arbitrary and, in practice, are seldom made. Knowledge is not usually articulated in that way.

expose you to tremendous risks.¹ Risks, of course, are often worth taking, but to say more than one knows solely for the sake of risk is foolhardy.

I suggest that scientific knowledge, though logically more articulate and far more complex, is of this sort. The books and teachers from whom it is acquired present concrete examples together with a multitude of theoretical generalizations. Both are essential carriers of knowledge, and it is therefore Pickwickian to seek a methodological criterion that supposes the scientist can specify in advance whether each imaginable instance fits or would falsify his theory. The criteria at his disposal, explicit and implicit, are sufficient to answer that question only for the cases that clearly do fit or that are clearly irrelevant. These are the cases he expects, the ones for which his knowledge was designed. Confronted with the unexpected, he must always do more research in order further to articulate his theory in the area that has just become problematic. He may then reject it in favour of another and for good reason. But no exclusively logical criteria can entirely dictate the conclusion he must draw.

IV

Almost everything said so far rings changes on a single theme. The criteria with which scientists determine the validity of an articulation or an application of existing theory are not by themselves sufficient to determine the choice between competing theories. Sir Karl has erred by transferring selected characteristics of everyday research to the occasional revolutionary episodes in which scientific advance is most obvious and by thereafter ignoring the everyday enterprise entirely. In particular, he has sought to solve the problem of theory choice during revolutions by logical criteria that are applicable in full only when a theory can already be presupposed. That is the largest part of my thesis in this paper, and it could be the entire thesis if I were content to leave altogether open the questions that have been raised. How do the scientists make the choice between competing theories? How are we to understand the way in which science does progress?

Let me at once be clear that having opened that Pandora's box, I shall close it quickly. There is too much about these questions that I do not understand and must not pretend to. But I believe I see the directions in which answers to them must be sought, and I shall conclude with an attempt briefly to mark the trail. Near its end we shall once more encounter a set of Sir Karl's characteristic locutions.

¹ This incompleteness of definitions is often called 'open texture' or 'vagueness of meaning', but those phrases seem decisively askew. Perhaps the definitions are incomplete, but nothing is wrong with the meanings. That is the way meanings behave!

I must first ask what it is that still requires explanation. Not that scientists discover the truth about nature, nor that they approach ever closer to the truth. Unless, as one of my critics suggests,¹ we simply define the approach to truth as the result of what scientists do, we cannot recognize progress towards that goal. Rather we must explain why science—our surest example of sound knowledge—progresses as it does, and we must first find out how, in fact, it does progress.

Surprisingly little is yet known about the answer to that descriptive question. A vast amount of thoughtful empirical investigation is still required. With the passage of time, scientific theories taken as a group are obviously more and more articulated. In the process, they are matched to nature at an increasing number of points and with increasing precision. Or again, the number of subject matters to which the puzzle-solving approach can be applied clearly grows with time. There is a continuing proliferation of scientific specialities, partly by an extension of the boundaries of science and partly by the subdivision of existing fields.

Those generalizations are, however, only a beginning. We know, for example, almost nothing about what a group of scientists will sacrifice in order to achieve the gains that a new theory invariably offers. My own impression, though it is no more than that, is that a scientific community will seldom or never embrace a new theory unless it solves all or almost all the quantitative, numerical puzzles that have been treated by its predecessor.² They will, on the other hand, occasionally sacrifice explanatory power, however reluctantly, sometimes leaving previously resolved questions open and sometimes declaring them altogether unscientific.³ Turning to another area, we know little about historical changes in the unity of the sciences. Despite occasional spectacular successes, communication across the boundaries between scientific specialties becomes worse and worse. Does the number of incompatible viewpoints employed by the increasing number of communities of specialists grow with time? Unity of the sciences is clearly a value for scientists, but for what will they give it up? Or again, though the bulk of scientific knowledge clearly increases with time, what are we to say about ignorance? The problems solved during the last thirty years did not exist as open questions a century ago. In any age, the scientific knowledge already at hand virtually exhausts what there is to know, leaving visible puzzles only at the horizon of existing knowledge. Is it not possible, or perhaps even likely, that contemporary scientists know less of what there is to know about their world than the scientists of the eighteenth century knew of theirs? Scientific theories, it must be remembered, attach

¹ Hawkins [1963].

² Cf. Kuhn [1961].

³ Cf. Kuhn [1962], pp. 102–8.

to nature only here and there. Are the interstices between those points of attachment perhaps now larger and more numerous than ever before?

Until we can answer more questions like these, we shall not know quite what scientific progress is and cannot therefore quite hope to explain it. On the other hand, answers to those questions will very nearly provide the explanation sought. The two come almost together. Already it should be clear that the explanation must, in the final analysis, be psychological or sociological. It must, that is, be a description of a value system, an ideology, together with an analysis of the institutions through which that system is transmitted and enforced. Knowing what scientists value, we may hope to understand what problems they will undertake and what choices they will make in particular circumstances of conflict. I doubt that there is another sort of answer to be found.

What form that answer will take is, of course, another matter. At this point, too, my sense that I control my subject matter ends. But again, some sample generalizations will illustrate the sorts of answers which must be sought. For a scientist, the solution of a difficult conceptual or instrumental puzzle is a principal goal. His success in that endeavour is rewarded through recognition by other members of his professional group and by them alone. The practical merit of his solution is at best a secondary value, and the approval of men outside the specialist group is a negative value or none at all. These values, which do much to dictate the form of normal science, are also significant at times when a choice must be made between theories. A man trained as a puzzle-solver will wish to preserve as many as possible of the prior puzzle-solutions obtained by his group, and he will also wish to maximize the number of puzzles that can be solved. But even these values frequently conflict, and there are others which make the problem of choice still more difficult. It is just in this connection that a study of what scientists will give up would be most significant. Simplicity, precision, and congruence with the theories used in other specialties are all significant value for the scientists, but they do not all dictate the same choice nor will they all be applied in the same way. That being the case, it is also important that group unanimity be a paramount value, causing the group to minimize the occasions for conflict and to reunite quickly about a single set of rules for puzzle solving even at the price of subdividing the specialty or excluding a formerly productive member.¹

I do not suggest that these are the right answers to the problem of scientific progress, but only that they are the types of answers that must be sought. Can I hope that Sir Karl will join me in this view of the task still to be done? For some time I have assumed he would not, as a set of phrases

¹ Cf. my [1962], pp. 161–9.

that recurs in his work seems to bar the position to him. Again and again he has rejected 'the psychology of knowledge' or the 'subjective' and insisted that his concern was instead with the 'objective' or 'the logic of knowledge'.¹ The title of his most fundamental contribution to our field is *The Logic of Scientific Discovery*, and it is there that he most positively asserts that his concern is with the logical spurs to knowledge rather than with the psychological drives of individuals. Until very recently I have supposed that this view of the problem must bar the sort of solution I have advocated.

But now I am less certain, for there is another aspect of Sir Karl's work, not quite compatible with what precedes. When he rejects 'the psychology of knowledge', Sir Karl's explicit concern is only to deny the methodological relevance of an *individual's* source of inspiration or of an individual's sense of certainty. With that much I cannot disagree. It is, however, a long step from the rejection of the psychological idiosyncrasies of an individual to the rejection of the common elements induced by nurture and training in the psychological make-up of the licensed membership of a *scientific group*. One need not be dismissed with the other. And this, too, Sir Karl seems sometimes to recognize. Though he insists he is writing about the logic of knowledge, an essential role in his methodology is played by passages which I can only read as attempts to inculcate moral imperatives in the membership of the scientific group.

'Assume', Sir Karl writes, 'that we have deliberately made it our task to live in this unknown world of ours; to adjust ourselves to it as well as we can; . . . and to explain it, *if possible* (we need not assume that it is) and as far as possible, with help of laws and explanatory theories. *If we have made this our task, then there is no more rational procedure than the method of . . . conjecture and refutation*: of boldly proposing theories; of trying our best to show that these are erroneous; and of accepting them tentatively if our critical efforts are unsuccessful.'² We shall not, I suggest, understand the success of science without understanding the full force of rhetorically induced and professionally shared imperatives like these. Institutionalized and articulated further (and also somewhat differently) such maxims and values may explain the outcome of choices that could not have been dictated by logic and experiment alone. The fact that passages like these occupy a prominent place in Sir Karl's writing is therefore further evidence of the resemblance of our views. That he does not, I think, ever see them for the social-psychological imperatives that they are is further evidence of the gestalt switch that still divides us deeply.

¹ Popper [1959], pp. 22 and 31 f., 46; and [1963], p. 52.

² Popper [1963], p. 51. Italics in original.

REFERENCES

- Braithwaite [1953]: *Scientific Explanation*, 1953.
 Guerlac [1961]: *Lavoisier—The Crucial Year*, 1961.
 Hafner and Presswood [1965]: 'Strong Interference and Weak Interactions', *Science*, 149, pp. 503-10.
 Hawkins [1963]: Review of Kuhn's 'Structure of Scientific Revolutions', *American Journal of Physics*, 31.
 Hempel [1965]: *Aspects of Scientific Explanation*, 1965.
 Lakatos [1963-4]: 'Proofs and Refutations', *The British Journal for the Philosophy of Science*, 14, pp. 1-25, 120-39, 221-43, 296-342.
 Kuhn [1961]: 'The Function of Measurement in Modern Physical Science', *Isis*, 52, pp. 161-93.
 Kuhn [1962]: *The Structure of Scientific Revolutions*, 1962.
 Popper [1935]: *Logik der Forschung*, 1935.
 Popper [1945]: *The Open Society and its Enemies*, 2 vols, 1945.
 Popper [1957]: *The Poverty of Historicism*, 1957.
 Popper [1959]: *Logic of Scientific Discovery*, 1959.
 Popper [1963]: *Conjectures and Refutations*, 1963.
 Stahlman [1956]: 'Astrology in Colonial America: An Extended Query', *William and Mary Quarterly*, 13, pp. 551-63.
 Thorndike [1923-58]: *A History of Magic and Experimental Science*, 8 vols, 1923-58.
 Thorndike [1955]: 'The True Place of Astrology in the History of Science', *Isis*, 46, pp. 273-8.