From Conventionalism to Scientific Metaphysics

11

Brian Ellis

Contents

Introduction

By the end of the first decade of the twentieth century, the most urgent problem for philosophers of science appeared to be that of reconciling their philosophies with the astonishing discoveries in space-time theory and electromagnetism. Albert Einstein had written his remarkable paper on the electrodynamics of moving bodies, better known as the Special Theory of Relativity (STR); Max Planck had introduced his counterintuitive quantum hypothesis to explain the empirical laws of black body radiation; and Einstein had used Planck's theory of atomic resonators to explain the photoelectric effect. The world of physics, which until then had seemed so solid and well ordered, was shaken and in some disarray. For these developments appeared not to be reconcilable with the Newtonian worldview that had, until then, dominated the scientific image of reality.

B. Ellis
School of Communication, Arts and Critical Enquiry, La Trobe University, Melbourne, VIC, Australia
e-mail: brian.ellis8@bigpond.com; B.Ellis@latrobe.edu.au

G. Oppy, N.N. Trakakis (eds.), *History of Philosophy in Australia and New Zealand*,
DOI 10.1007/978-94-007-6958-8_12,

Many of the leading physicists at this time were conventionalists. They were inspired by the writings of Ernst Mach, Pierre Duhem, Henri Poincaré, and Einstein himself to think of successful theories as being nothing more than the intellectual constructions of scientists—constructions that proved to be more or less useful for organising and systematising the results of experiments. But, as such, they argued, they can have no special claim to be representative of the world itself or to describe it as it truly is. In defending this view, Duhem argued that to claim any more for a theory would always be to take a step into metaphysics. To explain, he said, is to 'strip reality of appearances, covering it like a veil, in order to reveal the bare reality itself' (1954, p. 7). But scientists cannot do this, he said, without abandoning their chosen profession. They can only observe, record, and make mathematico-logical models of reality and seek to bring all of these facts and artifacts together into a coherent system. But they cannot explain anything, he said. That is the function of the metaphysician. And, in the eyes of these philosophers, metaphysics is, at best, just an idle pursuit.

Twentieth-century philosophy of science has been dominated by the consequences of this upheaval and the issue of what science can or should aim to do. The moderate conventionalism of Mach and Poincaré required philosophers to distinguish carefully between empirical facts and conventions, presumably so that they could see more clearly what must be preserved and what may be varied, in any future theory. The anti-metaphysical stance that these same philosophers took led to the more radical philosophical programs of the Vienna Circle, whose purposes were (i) to define the empirical contents of our scientific laws and theories and (ii) to rid the sciences of metaphysics. The first of these aims was to be achieved by logico-empirical analysis, i.e., by representing the laws and theories of science as universal propositions in a first-order predicate language, in which the variables and constants range over observables and the predicates are all observational. Thus, they became known as 'the logical empiricists' or, sometimes, as 'the logical positivists'. The second of these aims was to be achieved by means of the principle of verifiability, which was offered as a criterion for distinguishing between the meaningful and the meaningless.

Philosophers of science pursued these programs vigorously in the first half of the twentieth century. But in the second half, they reacted strongly against them. Conventionalism was overcome by epistemic holism, verificationism was largely replaced by falsificationism, logical empiricism fell to scientific realism, logico-empirical analysis was replaced by 'possible worlds' analyses, and finally, the demand for metaphysical explanations has become respectable again, but perhaps in a way that it never was before. This chapter traces the history of these movements in the second half of the twentieth century, as seen from the perspective of one who has been involved in all of them. We shall see that the advent of scientific realism was a significant turning point in philosophy in Australia. It effectively ended the dominance of ordinary language philosophy in Australia and shifted the emphasis away from questions of meaning to questions of being.

Conventionalism

Duhem's anti-metaphysical, and ultimately anti-realist, view of the aim of physical theory did much to define the agenda for the philosophy of science in the first half of the twentieth century. It enabled the phlogiston and caloric theories that had been overthrown in the nineteenth century to be seen as premature attempts to model reality. The empirical data on which these theories were built were mostly sound, he argued. The fault lay in the concrete models of reality that were constructed to explain them. Therefore, he argued, we should not put any faith in such models. They were, he argued, simply aids for the construction of formal theories and should be abandoned once they have served their purpose.

The conventionalists of this period, including Mach and Ostwald, and many scientists of the time accepted this line of argument. But, as a student, almost half a century later, I was more inclined to accept Norman Campbell's compromise with process realism. In *Physics: The Elements* (1921), Campbell argued that a physical theory always has three parts: an abstract model structure, a dictionary, and an analogy. The abstract model structure was the formal part of the theory, within which all necessary deductions could be made. The dictionary linked elements of the abstract model structure with observable things or properties, thus enabling any deductions or calculations carried out within the abstract model to be interpreted. And the analogy is the notional basis for the construction of the model. In the case of a process theory, it is the physical process that is postulated to explain the physical states or processes that are to be explained. In the case of a non-process theory, it is a formal analogy of some kind that is suggestive of the abstract model that is to be used for the purposes of explanation. But unlike Duhem, Campbell argued that analogies 'are not "aids" to the establishment of theories; they are an utterly essential part of theories, without which theories would be completely valueless and unworthy of the name' (Campbell 1921, p. 129). This was Campbell's compromise with Duhem on the issue of scientific process realism. It was a position one could sensibly take without rejecting scientific process theories. For it left the question open as to why process analogies were manifestly as useful as they were.

Meanwhile, I was convinced that there was still much work to be done in the conventionalist program. If Mach, Poincaré, Einstein, and his great interpreter Hans Reichenbach were basically right in their analyses, a great deal of scientific theory must be seen as depending on theoretically untestable assumptions, which, because they were untestable, had to have the status of mere conventions. Mach had argued, however, that what is true as a matter of fact, as opposed to what is true only by convention, is not always clear. Truth by convention often masquerades as factual. Poincaré (1952), for example, had argued that Euclid's axioms were 'neither synthetic *a priori* intuitions nor experimental facts' (p. 50). They are, he said, just conventions. He also argued that the law of inertia and other laws of mechanics are, in reality, only conventions. Reichenbach had argued that even 'the geometrical form of a body is no absolute datum of experience' but, he said, is dependent on the conventions involved in measuring space (Reichenbach 1958, p. 18). Einstein himself had argued that the principle that the one-way speed of light is the same in all directions is not, as it had always been thought to be, empirically testable, but is true only in virtue of the conventions for measuring space and time. As a student, I found all this pretty heady stuff—much more exciting than the linguistic philosophy that was all the rage in Oxford when I was there.

Conventionalism is, of course, a kind of positivism. For conventionalists would certainly have agreed that for a proposition to tell us anything about the world, its truth or falsity would have to make some observable difference to the world. I certainly thought that. I also thought that if a proposition were true by convention, then this could only be because it was definitional in nature, or a consequence of definitions, or otherwise part of a logical or mathematical system, such as an axiom or theorem. Such propositions were, as we conventionalists used to say, 'factually empty'. But we did not think these factually empty propositions were of no importance. After all, we supposed the propositions of mathematics to be all factually empty, and we had no desire to eliminate mathematics from science. We just thought it was important to distinguish the formal or conventional truths of science from the factual ones, because the conventions could be changed by changing the formal bases of our theories, and/or our coordinating definitions, but the factual truths that we adhered to were empirically certified, and so could not be changed, without sufficient empirical warrant to overthrow that certification.

But this attitude was contrary to the logical positivism of the Vienna Circle. For these philosophers had a different program and employed a different methodology. The Vienna Circle philosophers cast their mission as being the elimination of metaphysics from science. They saw themselves as warriors engaged in a crusade to rid science of the scourge of metaphysics and thus to establish science once and for all on a firm empirical foundation. Their chosen weapon in this crusade was the verificationist theory of meaning, and the banner under which they marched was the slogan: 'The meaning of a statement is the method of its verification'. As an adherent of the older conventionalist school of philosophy, I never had much time for this crusade or for its slogan. Nevertheless, this was the form of positivism that became best known to the English-speaking world. For this was the positivist theory that was most directly concerned with questions of meaning and hence was most in tune with the sort of linguistic philosophy that became fashionable after WWII. The one book on positivism that every student read when I was at Oxford was A.J. Ayer's *Language*, *Truth and Logic* 1936. The falsificationist theory of empirical significance developed by Sir Karl Popper in his *Logik der Forschung* in 1934 was much more plausible to me, as a conventionalist, than Ayer's verificationism. I also preferred his theory of the metaphysical as meaningful, but empirically vacuous. Unfortunately, I did not know the details of Popper's theory as a student, because Popper's book did not appear in English translation until 1959, when it was published as *The Logic of Scientific Discovery*.

Upon my appointment to a lectureship in the Department of History and Philosophy of Science (then called 'History and Methods of Science') at the University of Melbourne, conventionalism became my passion. I worked diligently on the conventionalist program, convinced that Mach's project of distinguishing fact from convention in science was about the most important thing that one could do in the philosophy of science. It was interesting for all of the reasons that I found Mach's, Poincaré's, and Reichenbach's works interesting. And it was important, because whatever is true only by convention must be subject to change. New conventions could obviously yield new insights, as Einstein's STR had so clearly demonstrated.

At the time, I was working pretty much on my own in this area. But I learned a lot from Douglas Gasking. Gasking was a student of Wittgenstein whose writings were a model of clarity and whose methodology was thorough and persuasive. In his essay 'Mathematics and the World', Gasking (1940) defended the conventionalist thesis that one could get along quite well with an arithmetic in which

$$(a\mathrm{x}^{*}b) = ((a+2)\times(b+2))/4,$$

and hence that

$$4\mathrm{x}^{*}6 = 12,$$

provided that one used different techniques for counting and measuring. And what I learned from this paper was a methodology of testing conventionalist claims: if you think that p is conventional, then to prove your point you must be able to show that for some q that is incompatible with p, a theory in which q is presupposed is no less viable empirically than one in which p is presupposed. This was the test that I used in all of my papers on conventionalism written while I was still working in the University of Melbourne's Department of History and Philosophy of Science. It was also the test that I used in my book *Basic Concepts of Measurement* (Ellis 1966). In this period, I never willingly accepted a conventionalist claim, unless I thought I could show that another, empirically no less viable, convention could be adopted in its place.[1]

It needs to be stressed that the conventionalist claims that have been made over the years are not necessarily analyticity claims, although every analyticity claim is ultimately a conventionalist one. For, as every conventionalist since Mach has argued, all conventions worthy of note depend on the existence of abstract theories linked to reality through coordinative definitions, i.e., propositions linking the abstract terms of the theory (e.g., numbers, addition operations, functions) with observables (e.g., spatial or temporal coincidences, meter readings). Adolf Grünbaum calls them 'Riemannian' conventions. These are the propositions that are contained in what Campbell called the 'dictionary' of a theory. But as Campbell had argued in 1921, a pure abstract theory T tells us nothing about the world. It is just an abstract logico-mathematical system. To be informative, he said, there needs to be a dictionary D that links the theoretical terms of T, or perhaps certain functions of these terms, with observables. In the case of analytic propositions of the most trivial kind (such as 'A bachelor is an unmarried man), there is no theory

[1]John Fox (2007) has recently convinced me that I did make one serious error of this kind. My 'dinch' scale for the measurement of length leads to inconsistencies, given the way in which I proposed to use it.

T involved, and the convention to name an object or property in one way rather than another is, from a scientific point of view, completely arbitrary. There might be aesthetic reasons or reasons of convenience for naming things as we do, but no good scientific ones. Grünbaum calls all such conventions 'trivial semantic' ones.

In the 1950s, while I was working on conventionalism, many of my colleagues in Philosophy Departments around the country were working on the analytic-synthetic distinction, which had apparently been demolished in W.V.O. Quine's (1953) *From a Logical Point of View*. But I never felt obliged to defend the concept of analyticity against Quine, despite my commitment to conventionalism. For none of the propositions that I argued were conventional could possibly be mistaken for the trivial semantic ones that are thought to be analytic. As Grünbaum and I and most other philosophers of science at the time understood it, analyticity was a problem for 'ordinary language' philosophers, not one for conventionalists.[2] It was not a problem for us, because we were studying scientific practice, not ordinary language. We were concerned with the possibilities of defending alternative $T+D$ (Theory + Dictionary) combinations to account for the same ranges of facts as existing theories. We would, almost all of us, have said that if T_1+D_1 and T_2+D_2 could both adequately explain the same set of facts about the world, and could not in principle be separated experimentally, then it is conventional in the nontrivial Riemannian sense that we should accept T_1+D_1 rather than T_2+D_2, or conversely. And, echoing Reichenbach, what most of us would have said is that 'there is no truth of the matter' whether T_1 or T_2. T_1 might be said to be true given D_1 or T_2 given D_2. But in the absence of the required coordinating definitions, there is no truth to be found.

In developing my conventionalist theories, I worked mostly on my own. For there were very few other philosophers in Australia engaged in the same program. George Schlesinger, who was a graduate student of mine in the late 1950s, was one with whom I could talk about conventionalism, and, among other things, we did some good work together on Moritz Schlick's bizarre claim that there is no fact of the matter whether the universe and everything in it either did or did not double in size overnight. Schlesinger and I thought that there clearly was a fact of the matter in this case, and we set about to prove it. We both argued (Ellis 1963, Schlesinger 1964) that there would be a great many observable consequences of such an occurrence. We argued that even in a Newtonian world, in which space-time would be Euclidean, quantities that vary nonlinearly with length would be differentially affected. And we all know that there are many such quantities. But Grünbaum, the world's most revered defender of conventionalism, would have none of it, and a humorous, but not very enlightening, debate followed in the journals on what became known as 'the nocturnal doubling hypothesis'. Grünbaum (1964, 1967) defended Schlick. Schlesinger (1964, 1967) argued against him.

[2] I remember Grünbaum saying to me once something to the effect that 'The analytic-synthetic distinction [i.e., the distinction between what is true in virtue of the meanings of words, and what is not] is one thing, and may well be untenable, but the fact-convention distinction is another, and is absolutely fundamental'.

In 1962–1963, I spent 8 months of my study leave in the Philosophy Department at the University of Pittsburgh, where I was required to teach two courses—one graduate and the other undergraduate. I also worked closely with the philosophers in the Andrew Mellon Center for the Philosophy of Science, where Adolf Grünbaum was the Director and Nicholas Rescher the Deputy Director. As a result of these arrangements, I was thoroughly involved with both teaching and research in philosophy of science in Pittsburgh and found myself working closely with other members of staff in this and other areas. The Department had as good a group of graduate students as you could possibly wish to have. As an added bonus, George Schlesinger was a Postdoctoral Fellow at the centre while I was there. I saw a lot of Bruce Aune, who was a fellow staff member involved in the graduate program, and Brian Skyrms, Ernie Sosa, and Kent Wilson were three of the graduate students that I remember well.

While in Pittsburgh, I wrote a long paper 'On the Origin and Nature of Newton's Laws of Motion' for Robert Colodny's (1965) book, *Beyond the Edge of Certainty*, and defended my conventionalism concerning the law of inertia against all comers. I also finished a paper that I had begun in Melbourne, called 'Universal and Differential Forces', in which I signalled that I no longer accepted some of the more outlandish conventionalist claims that had been made by Reichenbach and others. There was a difference, I thought, between the kind of geometrical conventionalism that Reichenbach and Grünbaum defended and the kind of conventionalist program that Schlesinger and I were pursuing. But, at that time, I had not appreciated just how deep this rift really was. The first clear symptom of this was the seemingly absurd dispute over the nocturnal doubling hypothesis. As I recall, this dispute did not surface while we were in Pittsburgh.[3] It broke with the publication in 1964 of Schlesinger's 'It is False that Overnight Everything has Doubled in Size'. But, even then, I did not understand its full significance. I thought that Schlesinger was obviously right. But I also thought that this was just one of those little in-house disputes that one can expect to find in any philosophical movement. I was wrong, however: it went much deeper than that.

On my return from Pittsburgh, I completed work on my manuscript, *Basic Concepts of Measurement*, and saw it through to publication. At about this time, I began work on the philosophical foundations of STR, which I considered to be a topic that every conventionalist would have to examine at some stage. I was familiar with the views of the Pittsburgh school, but I was not at all sure what my own attitude would be to some of Einstein's more extravagant conventionality claims. I was particularly interested in Einstein's claim that there is really no way of determining whether the speed of light in any one direction in space is the same or different from the one-way speed of light in any other direction. Consequently, we are free to adopt it as a convention that the one-way speed of light is the same in all directions. This thesis is known as that of the conventionality of distant simultaneity. In 1965, Peter Bowman, a graduate student from America, came to work

[3]My 1963 paper was not published until November of that year.

with me on the philosophy of space and time, and, in due course, we began to work on this topic. In the following year, I was fortunate to have some very able young philosophers of science in my honours and graduate classes, including John Fox, Greg Hunt, Robert Pargetter, and Barbara Marsh. The paper that Bowman and I eventually published owes a lot to the contributions they all made.

From Conventionalism to Holism

In the 1920s, in the early days of positivism, philosophers were given to making startling pronouncements, which they defended brilliantly by narrow geometrical conventionalist arguments. Consider, for example, the following propositions:

1. The universe and its contents did not double in size overnight.
2. The sun is a roughly spherical object many times the diameter of the earth.
3. The one-way speed of light in a vacuum is the same in all directions.

These would all appear to be straightforwardly true propositions. But, according to Schlick and Reichenbach, there is no truth of the matter concerning any of them. The last of these claims that the one-way speed of light is the same in all directions was held with great conviction to be a mere convention. On this, my conventionalist colleagues had the authority of Einstein himself. It is a fact, Einstein said, that the average speed of light (in vacuo) over an out-and-back path is always the same. But there is no fact of the matter whether the one-way speed of light is always the same. The one-way speed of light, Einstein argued, depends on our definition of simultaneity, which in turn depends on what we assume the one-way speed of light to be. The standard definition makes the speed of light in a vacuum the same in all directions, but, he thought, other definitions that make the speed of light a function of direction could equally well have been chosen.

The arguments for these conventionality claims all have the same form. Each argument points to some preceding definition or definitions, which, it is said, would have to be accepted before any measurements of shape, size, distance, or speed could begin. To measure shape, size, or distance, for example, we must have criteria for determining whether one thing is, or is not, the same size as another, where these things are at different places or exist at different times. But such criteria cannot be established experimentally, because they would have to be accepted before any relevant experimentation could begin. We would need criteria for comparing lengths or time intervals in order to judge whether any proposed new criterion for establishing these relationships is satisfactory.

I do not propose to go into details concerning these arguments. But if you keep asking yourself the question, ‘How could we possibly establish that this is equal in length (or distance) to that, when the objects concerned are not close together in space and time (so that the relationship between them could be directly observed)?’, you will quickly get the idea. For you will soon find yourself arguing in circles and getting nowhere. You will find, for example, that measurements of distance depend on measurements of length, which depend on assumptions about what is fixed in length, which cannot be checked without making further assumptions about lengths or distances. Likewise, if you follow the same lines of questioning, you will find that measurements of speed depend on measurements of distances and travelling times, which depend on assumptions about clocks, which cannot be checked without making other assumptions about distances, clocks, or speeds. Ultimately, say the old-fashioned conventionalists, you cannot break out of any of these circles. You will have to make some decisions somewhere about what you will count as being the same in length, or ticking at the same rate, or occurring at the same time. That is, you will have to make a number of stipulations about these things, and, in the final analysis, the stipulations you make will have to be made on grounds such as those of descriptive simplicity or convenience. Truth does not come into it.

The common assumption of all of these arguments is that our spatial and temporal concepts are purely comparative, i.e., they depend entirely on the procedures we use for comparing these quantities directly. However, I was beginning to think that this assumption must be false. Length and time interval are two of the most basic physical concepts, and there are few physical laws that do not involve one or other of them. Consequently, changes of length or time interval will affect behaviour in a whole lot of different ways and will not only affect the results obtained by direct comparisons of length or time interval. When one object expands relative to another, the effect can be established directly by measurement. But the change of relative size that would be noticeable is not the only effect. There will be hundreds of other effects, depending on how the change of size is brought about. Consequently, even if we could not observe any changes of relative size, it should be perfectly obvious to us that a change of size has occurred. There are, consequently, many good reasons for believing that there has been no catastrophic expansion of the universe overnight and no good reason for believing it has. And, as George Schlesinger and I argued, the hypothesis that the universe and everything in it has doubled in size overnight is not meaningless: it is simply false.

Similarly, it is undeniable that the sun is a roughly spherical object that is many times the diameter of the earth. There may be some ways of measuring what might be called the ‘shapes’ and ‘sizes’ of things, and of comparing them at different places, that yield a different result. But these new ‘shapes’ and ‘sizes’ would not be the ones we need, or could plausibly use, for describing the world. The price of adopting a system of conventions for measuring length (or what will now be called ‘length), according to which the sun is not roughly spherical or is smaller than the earth, would therefore be very great. If, given a new definition of ‘length’, such an undoubted truth as the proposition that the sun is roughly spherical and very much bigger than the earth must be considered to be false, then that definition of ‘length must be unsatisfactory. For this is surely a fact about the world, if anything is.

The realisation, evident in my ‘Universal and Differential Forces’ (1963), that there are such tight constraints as these on what definitions are acceptable was a turning point in my philosophical thinking. For it led to my abandonment of the old conventionalist program and the adoption of a more sophisticated philosophical position. If such constraints exist, I argued, then the definitions we accept are by no means arbitrary. We may, for theoretical purposes, try to define terms that are already in common scientific use. But our acceptance of these definitions must always be tentative, and we should be willing to abandon them, if the price of persevering with them is too high. Any proposed definitions of terms like 'length' or 'time interval', which are deeply involved in our theoretical understanding of the world, must pass some very severe tests. If a proposed definition would force us to deny what is obviously true, according to accepted theories about the nature of reality, then this definition must be rejected. So, definitions turn out to have a theoretical status not significantly different from other hypotheses. They can be shown to be unsatisfactory, if they can be shown to have clearly unacceptable consequences.

The old conventionalism was based on the belief that there is a clear distinction between what is true as a matter of fact and what is true by definition or convention. But I no longer believed that there was any such clear distinction. So the old conventionalist program of sorting the empirical facts from the conventions in science had to be abandoned. I still thought it was important to be clear about why we should accept or reject the propositions we do. But from now on I would expect there to be a spectrum, ranging from arbitrary definitions at one extreme to hard empirical data at the other. In between, I supposed, there would be the accepted body of scientific theories and hypotheses, for which the evidence would be just more or less compelling.

This change of perspective had many consequences, which I could hardly begin to think through. Firstly, if the body of scientific knowledge is a complex integrated structure of laws and theories that cannot be analysed into propositions that are true by definition or convention (which signal how we are proposing to use language to describe the world) and propositions that are true as a matter of fact (and so, presumably, correspond to reality in some way), then what is the concept of truth that is needed for science? This is a question that I would later take up and return to several times, before reaching an answer that I could feel reasonably happy with.

Secondly, if scientific concepts, like mass, length, charge, and time interval, are not normally definable, except in ways that are already consistent with the laws and theories we accept, then what does this tell us about these concepts? It implies, for one thing, that many of our most important scientific concepts are defined implicitly by their roles in the laws and theories in which they occur. And, if any of these concepts should be defined explicitly, then this explicit definition has no special status. It is, at best, just a tentative agreement to axiomatise or formalise the system in one way, rather than in another. But it remains as open to adjustment in the light of experience as any of a number of other propositions that have not been declared to be true by definition. The stance that I was forced to adopt, therefore, implies a kind of holism about scientific knowledge and understanding. As Quine (1953, p. 42) remarked in his important paper 'Two Dogmas of Empiricism':

> The totality of our so-called knowledge or beliefs, from the most casual matters of geography and history to the profoundest laws of atomic physics or even of pure mathematics and logic, is a man-made fabric which impinges on experience only along the edges. Or, to change the figure, total science is like a field of force whose boundary conditions are experience. A conflict with experience at the periphery occasions adjustments in the interior of the field... But the total field is so underdetermined by its boundary conditions, experience, that there is much latitude of choice as to what statements to re-evaluate in the light of any single contrary experience.

In 1960, Douglas Gasking, in one of his more Wittgensteinian papers, argued that many of our kind concepts are ‘cluster’ concepts. For many recognisable kinds of things, such as games, exist but have no defining characteristics, i.e., no sets of characteristics that would distinguish them from things of other kinds. They are, he argued, defined only by the overlapping clusters of characteristics by which they might be identified. In 1962, Hilary Putnam argued that many of the quantitative concepts of science are also cluster concepts of a sort. Where two or more different kinds of procedures for measuring a quantity exist, and these procedures (when properly carried out) are guaranteed by the laws of nature to yield the same results, no one of them can be singled out as defining the measure of the quantity. For this would be tantamount to making one of the laws involving the quantity true by definition, and all of the others true only as matters of fact (which would be arbitrary). Putnam (1962) argued that where a quantity might equally well be defined in any of a number of different ways, depending on which law is chosen to define it, what we have is a ‘law cluster’ concept.

The movement towards Quinean holism, and hence away from the empiricist distinction between facts and conventions, was certainly in the air by 1962 and was gathering strength. So I cannot claim any great originality for my belated discovery of this basic flaw in the foundations of conventionalism. In fact, as a committed conventionalist until the mid-1960s, I was rather slow off the mark. Perhaps this was because measurement theory, on which I had been chiefly engaged for many years, is one of the few areas in which the distinction often seems to be both clear and justified. But Quine’s attack on the assumptions of conventionalism had hardly touched the philosophy of science establishment in America, which remained as wedded to the empirical fact-convention distinction as it had ever been. And this was the source of our disagreement with Grünbaum about the nocturnal doubling hypothesis. It would also prove to be the source of the much more virulent disagreement with the American philosophy of science establishment that arose later about conventionality in distant simultaneity.

For my honours and graduate class in philosophy of science in 1966, we decided to look at the alleged conventionality of distant simultaneity. This particular conventionality claim had been made originally by Einstein and was argued for at length by Reichenbach. The thesis was very widely accepted by philosophers of science, and it had become something of a cornerstone of conventionalist theory. According to Reichenbach’s analysis, there are no empirically establishable facts about the one-way speed of light other than those that are already implied by the fact that its speed (*in vacuo*) over any out-and-back path is always the same. For there is no way of defining distant simultaneity that does not already depend on what assumptions we make about the one-way speed of light. This was argued specifically by Reichenbach. But I was sceptical. If it is not a matter of fact that the one-way speed of light is the same in all directions, but is just a consequence of a decision to define distant simultaneity in such a way as to make it so, then we should be able to vary that decision and construct what is plausibly still a reasonable definition of simultaneity that makes the one-way speed of light a function of, say, direction. That this could be done had never been demonstrated.

We proceeded, in accordance with normal procedure, to consider the possibility of constructing a version of the STR that would be equivalent to the STR empirically, but based on a different convention regarding the one-way speed of light. We put certain constraints on the definitions of distant simultaneity that would be acceptable, some formal and some empirical. Firstly, we argued, it should be consistent with all of the known facts about light signals: the average speed of light over any out-and-back path must be a constant c. Secondly, it must be causally consistent: signals should not be able to arrive at their destinations before they are sent. Thirdly, it must satisfy certain formal requirements. In particular, the relationship of simultaneity that is defined must be formally an equivalence relationship, i.e., one that is reflexive, symmetrical, and transitive. The standard signal definition of synchrony, according to which the one-way speed of light between any two objects A and B is at rest in a given inertial system, clearly satisfies all of these requirements. The question is: Are there any others? To simplify the question, we chose to consider whether it would be possible to construct a formal definition of synchrony which made the one-way speed of light a continuous function of direction—one that is symmetrical about the X-axis of a rectangular coordinate system. We proved that there is indeed a way of doing just this.

Let $e_\theta = {}^1\!/_2\left({}^c/_{c_\theta}\right)$, where c_θ is the speed of light in the direction θ from the X-axis. Then, demonstrably, the various requirements on a non-standard signal synchrony relationship in a given inertial system are satisfied, if

$$e_\theta = (e_\theta - {}^1\!/_2)\cos\theta + {}^1\!/_2,$$

where $(0 \leq e_0 \leq 1)$.

We called this 'the distribution law for light velocities' (Ellis and Bowman 1967). So, we concluded, the conventionality thesis passes the first test. Next, we considered whether we could use this definition in the standard way to derive a non-standard version of the STR. But we were able to prove that this would require a sacrifice. In deriving the Lorentz transformation equations from the standard signal definition of synchrony, it is normally assumed that (1) a uniform straight line motion in any one inertial system always appears as a uniform straight line motion from the perspective of any other inertial system (we called this the 'principle of linearity') and (2) the velocity of A with respect to B as measured by B must always be minus the velocity of B with respect to A as measured by A (we called this the 'principle of reciprocity of relative velocities'). John McPhee, a Melbourne mathematician who helped us on this project, proved that if these two assumptions are added as requirements on a non-standard signal synchrony definition, then no such definition is possible. The only possible one that would preserve both linearity and reciprocity is the standard one. Moreover, if we were prepared to pay this price, then we should be faced with other difficulties. For acceptance of any non-standard definition of signal synchrony would require us to postulate the existence of universal forces to explain what we must then suppose to be the odd behaviour of slowly transported clocks.

(They would be found to get out of synchrony as they are moved apart but would come back into synchrony as they are brought back together again).

Thus, we discovered, after we had already done most of the work of testing Einstein's conventionality claim, that his original justifying reason for making this claim in the first place is simply false. For there is, contrary to what Einstein and Reichenbach say, a way of synchronising widely separated clocks without making any prior assumptions about the one-way speed of light. Given the STR, it is demonstrable that clocks can, in principle, always be synchronised (to any desired degree of precision) in any inertial system just by moving a standard clock around sufficiently slowly and synchronising all other clocks in the reference frame with this standard. The method is known as that of 'slow clock transport'. There is no dispute about this: it is a clear and unequivocal prediction of the STR that this can be done. And, as we later discovered, the method had already been described by P.W. Bridgman in his little book, *A Sophisticate's Primer of Relativity* (1962). There is, therefore, a perfectly good way of measuring the one-way speed of light empirically, viz., by Römer's method, using Jupiter's moons as clocks, and measuring the times of successive occultations and re-emergences (Jupiter's moons being, effectively, slowly transported clocks). Moreover, this method enables us to measure the one-way speed of light in many different directions in space (depending on the direction of Jupiter from earth). And the empirical finding is that, to a high degree of precision, the one-way speed of light is the same in all of these different directions.

We published these results in 1967 in *Philosophy of Science*. The paper produced a strong reaction. Most of the March 1969 issue of *Philosophy of Science* was given over to what can only be described as a concerted attack on our paper by three of the leading philosophers of science in America (Grünbaum, et al. 1969), viz., Adolf Grünbaum, Wesley Salmon, and Bas van Fraassen, all of whom were in Pittsburgh at the time. Their papers made up the bulk of an 81-page 'Panel discussion of simultaneity by slow clock transport in the special and general theories of relativity'. Included in this panel discussion was also a paper by the physicist Allen Janis, which dealt with the possibility of using slow clock transport as a way of synchronising clocks in non-inertial frames, i.e., systems of a kind that can only be described adequately using the apparatus of the General Theory of Relativity. This was not an attack on our paper, however, since the more general question is not one that we considered, and was not relevant to the point we were making.

I did not mind the attack. In fact, I thought it was all good fun. John Fox and I were the two people most concerned with the issue who were still working in Melbourne, Peter Bowman having returned to America. It was obvious to us that the Pittsburgh Panel had misunderstood, and systematically misrepresented, our philosophical position. And for a long time I wondered why. The panel had no quarrel with our understanding of the STR, or with what we took to be its factual basis, or with any of our proofs concerning the theory's implications. The main argument was with our claim that there are 'good physical reasons' for preferring a definition of simultaneity that makes the one-way speed of light always the same. But what is wrong with that? Isn't the existence of a coincidence of at least two logically independent, isotropic, and formally satisfactory criteria for distant simultaneity, and the absence of any comparable reasons for adopting any other possible criterion, good reason enough? And isn't it an empirical fact, one establishable by observation and experiment, that such a coincidence of isotropic criteria exists? If so, then surely there are good physical reasons for adopting such a criterion. What was all the fuss about?

The fuss was all about conventionalism itself. The dispute about the conventionality of distant simultaneity was not just one about the status of Einstein's definition; it was about a core doctrine of the conventionalist program. If this conventionality claim were agreed to be lacking in substance, as Bowman and I had argued, then the program of conventionalism must itself be discarded as lacking in substance. The trouble, although I was not fully aware of it at the time, was that I had ceased to think as a conventionalist. I spoke and understood the language of conventionalism, but I was thinking as a Quinean holist.

My new way of thinking is well illustrated by the following passage taken from my reply to the Pittsburgh Panel:

> There is no foundation of hard empirical fact in science, only a choice between competing theories and conceptual frameworks, which, at any given time, seem adequate for the description and prediction of events. Any theoretical statement which occurs in any theory may come to be rejected if a better or more promising theory comes to replace it, and it is simply irrelevant whether the statement in question is, relative to some particular axiomatisation of the theory, definitional or not. The conventional-empirical distinction, as it has come to be used, has been a plague on the philosophy of science since the rise of Positivism. (Ellis 1971, p. 199)

Fox and I completed our replies and sent them to *Philosophy of Science*. But the editor of the journal rejected them. It was not because they were not up to standard, he explained, but because 'too much journal space had already been occupied by the issue'.

Scientific Epistemology

In 1962 Thomas Kuhn published *The Structure of Scientific Revolutions*, a book that had a profound effect on our understanding of scientific method. For the logical empiricists, the method of science was thought to be essentially inductive. Hence, for them, the problems of induction, and of inductive logic, were the main ones in scientific epistemology. But there was another very different view of scientific method that had not, as yet, had much impact in the English-speaking world. This was Sir Karl Popper's anti-inductivist methodology of conjectures and refutations. Popper did not share the logical empiricists' view that verifiability is the hallmark of empirical significance. So, he had no interest in trying to show that the laws and theories of science are true or even probable. On the contrary, he argued that what distinguishes science from non-science is the falsifiability of its laws and theories. So, he argued, scientists should not be seeking to confirm their theories. Rather, they should always expose them to falsification as much as possible in their pursuit of knowledge. If their laws and theories are corroborated, i.e., pass severe enough tests, then they may be provisionally accepted. Otherwise, they must be rejected.

But Kuhn's book challenged both methodologies. Normal science, i.e., the kind of scientific work that most scientists are involved in most of the time, involves commitment to a program that satisfies the broad parameters of what he calls a 'research paradigm'. And the aim of scientists working within such a paradigm is to show what it can do. Their aim is not to refute the main tenets of the theoretical stance they have taken (as Popper thought it should be), but to articulate the position, with the aim of showing how, consistently with these tenets, it can be adapted to deal with the outstanding problems of the area. So, Kuhn's methodology of normal science was not Popperian. It was not inductivist either, although, naturally, the more successful a research program was at handling the empirical data, and solving the problems that fell within its ambit, the more highly it was regarded. Empiricists thought that science required a theoretically neutral observation language as a foundation for their work. But normal science operated under no such constraints. It was research that proceeded from an overall position that interpreted the data, defined the main problems of the area, and explained how one should go about trying to solve them and what would constitute a satisfactory solution. Within this theoretical framework, normal science was discovery oriented, but the soundness of the framework was seen to depend on its problem-solving ability, not on the inductive evidence for it or on the severity of the tests that it had passed.

Kuhn argued that scientific revolutions are paradigm shifts that are normally brought about by paradigm failures. When a research program gets into difficulties, he argued, scientists working in the area begin to explore other ways of conceptualising the data and thinking about its problems. And, when they begin to do this, he said, the science enters an abnormal phase. New ideas are thrown around, and the tenets of the old program are cast into doubt. Of course, one cannot create a new paradigm overnight. One has to work at it and gather colleagues around one to develop new ideas. And, typically, one will see the development of different schools of thought, each seeking better ways of understanding and researching the troubled area. The methodology of abnormal science is thus very different from that of normal science. It is much more philosophical and reflective, and the results are much harder to evaluate, because, in the area affected, there is no longer any general agreement about how the data should be read, or what it shows, and there may not even be agreement about what the main problems are, or what would constitute solving them. Consequently, defenders of different paradigms will often misunderstand and talk past each other. Where this happens, Kuhn argued, the problem may be one of incommensurability, i.e., the different perspectives on the world may be so different that they do not have even a common observation language.

But the Popperians at the London School of Economics (LSE) did not take this attack on their position lying down, and Imré Lakatos and his colleagues took up the challenge. Lakatos had earlier extended Popper's methodology into the field of mathematics in his seminal papers of (1963–1964) entitled *Proofs and Refutations* and showed clearly how counterexamples to alleged proofs could (and did) lead to the development of new concepts and theories. But despite the attractiveness of Popper's scientific methodology, and Lakatos' extension of it into the field of mathematics, Popperians really had no answer to Kuhn's methodology of normal science. It was clearly conservative of core doctrines, in a way that Popper's own methodology of conjectures and refutations was not. If the experimental evidence seems to contradict the theory that is being used to make predictions, about all that one can say is that something is wrong somewhere. It might be in the observations that are being made, or in the theory of the instrumentation involved, or in any of the many subsidiary hypotheses made in the design of the experiment, or in the more frustrating cases, it may be supposed that there must be unknown forces (e.g., due to dark matter) or extraordinary processes (e.g., that of global inflation) occurring, whose mechanisms are as yet unknown. Lakatos was, in fact, very familiar with some of the many strategies that can be used to deal with counterevidence, as he demonstrated in *Proofs and Refutations*, and some of them, such as those of 'monster barring' and 'monster adjustment', are often referred to in philosophical literature on subjects other than the philosophy of mathematics. Lakatos' (1970) considered reply to Kuhn's critique is to be found in his major paper on the methodology of scientific research programs.

In 1972 I was fortunate enough to be able to spend a period of study leave in London and work with the exciting group of philosophers of science there, particularly those at the LSE. I found myself torn between the Popperians and the Kuhnians. I liked Popper's forthright anti-inductivism, but was enough of a historian to think that the methodology of science was not so rigidly falsificationist, or tied to the project of increasing the empirical content of our theories, as Popper had supposed. I thought it was to increase our understanding of the world, although I have to admit that I did not have a very clear idea of what that involved. The 'conjectures' part of Popper's methodology might well be the mechanism of growth, I thought. But the 'refutations' part of it was manifestly inadequate. It would be more accurate to speak of a methodology of 'conjectures, articulation, development, and testing'. But somehow it does not have quite the same ring to it. Kuhn, on the other hand, was historically well informed and honest in his reporting of scientific methodology. But Kuhn's theory provided no simple answer to the question of method. As a Quinean holist, I was inclined to think that this was as it should be and that the important features of scientific belief systems are their explanatory power, elegance, rational coherence, and general compatibility with observational and experimental evidence, not how they were arrived at. Presumably, scientific method would have to be one that was guided by our epistemic values and which allowed a good deal of latitude in the making and development of scientific hypotheses.

At the time of my visit to LSE, I was working on a paper I called the 'Epistemic Foundations of Logic'. In fact, by 1971, I had drafted a book with this title and sent copies of it to colleagues at home and overseas for comment. This project was based on the following assumptions: (1) logic is, or ought to be, part of the general theory of rational coherence. For the logic of the truth and falsity claims that can be made in a given formal language L is just the set of necessary and sufficient conditions for the rational coherence of any subset of such claims. (2) The logic of subjective probability is, likewise, part of the general theory of rational coherence. For the logic of the subjective probability claims that can be made in a given formal language L is just the set of necessary and sufficient conditions for the rational coherence of any subset of such claims. (3) The logic of truth and falsity claims is derivable from that of subjective probability simply by restricting the range of possible subjective probability values to 1 and 0. It is demonstrable, for example, that the set of all valid formulae of the propositional calculus is the set of all propositional formulae Z such that $P(Z) = 1$ is a theorem of the probability calculus. I called this 'the logical correspondence principle'.

My aim was to use the logical correspondence principle, and the full apparatus of the probability calculus, to derive a much more comprehensive system of logic than any that had so far been developed. The classical propositional calculus corresponds to just the absolute fragment of the probability calculus. But what was needed, I thought, was a propositional calculus with an 'if' connective that corresponds to the 'given' operator in an enhanced probability theory. For it seemed to me that this would be a much better way of representing conditionals formally than the usual one using the material conditional. But there were two major problems to overcome. (1) $P(q/p)$ is undefined, if $P(p) = 0$. That is, there was a problem of how to deal with counterfactual conditionals. (2) $P(r/(q/p))$ and $P((r/q)/p)$ are undefined in the probability calculus. That is, nested conditionals are undefined in the probability calculus. But, as far as I could see, the probability calculus worked well as a system of logic, provided that there were no counterfactual or nested conditionals. And, I could see no reason why an augmented probability calculus with counterfactual and nested conditionals could not be developed. It was my project to do just that.

My manuscript 'Epistemic Foundations of Logic' was never published. Robert Stalnaker (1968) had done much better than I had in developing a logic of the required kind with strong conditionals. I suppose that Stalnaker's theory could also have been used to derive a probability calculus with counterfactual and nested conditionals, although I have never seen this idea explored. Such a logic was needed, I argued, because the material conditional was a manifestly inadequate representation of 'if…then', especially within the context of a probability claim. The only plausible representation of 'if', I argued, was the conditionalisation operator '/'. But my own logical system based on this identification, which I had hoped would prove to be infinitely many valued, proved not to be. I had a proof that it was at least 4-valued. But that did not count for much. David Lewis wrote to me with his famous triviality proof and showed that it was at most 4-valued. So, in the end, I just gave up. I had been gazumped in doing what I had hoped to do and proven wrong about what I had done. Possible worlds semantics was indeed a powerful tool, which Lewis and Stalnaker used to considerable advantage in constructing their modal and conditional logics. But I did not believe in merely possible worlds, and at the time I had nothing to put in their place. I did not even have a theory of why possible worlds semantics worked as well as they did. I now think I know. But it was not until I had written *Rational Belief Systems* that I could identify clearly the equivalents of possible worlds in my meta-logical framework. The equivalent of a merely possible world in my meta-logic is just a world that would correspond to a kind of rationally completed belief system, if everything believed to be true in it was indeed true (Ellis 1979, Chap. 3).

Scientific Realism

In 1963 Smart published his important book *Philosophy and Scientific Realism*. This book is important in the history of Australian philosophy for a number of reasons. Firstly, it represented a significant change in emphasis in Smart's own philosophy, from one of conceptual clarification to one of seeking a more comprehensive understanding of the world. When he first arrived in Australia in 1950, he brought the Oxford conception of philosophy with him and so tended to see philosophy as a form of intellectual therapy, as a clearing up of muddles created by common misunderstandings of ordinary language. Something of this same attitude also existed in Melbourne's Department of Philosophy, where the influence of Wittgenstein was supreme. Sydney had long had a very different tradition, due to the charismatic influence of John Anderson, who was a realist of sorts. Smart's book on scientific realism (1963) was not in the Andersonian tradition, but it was a clear break with the Cambridge/Cambridge one of ordinary language philosophy and was a significant attempt to elaborate a new worldview. Secondly, the book bridged the gap between Cambridge philosophy and Sydney realism and helped to end the dominance of Melbourne philosophy. It also did much to define the nature of Australian philosophy. Following the publication of this book, Australian philosophy was often referred to overseas as 'Australian materialism', and the scientific realism that characterised it was often thought of as a 'down to earth', 'no nonsense' sort of philosophy that was based on a layman's understanding of the science of our times, which is what *Philosophy and Scientific Realism* was.

The publication of Smart's book also marked the beginning of an Australia-wide shift away from the philosophical method that was characterised by Richard Rorty (1967) as 'the linguistic turn'. For few of the arguments in this book were ones that could possibly have been defended in ordinary language philosophy or by arguments that depended primarily on semantic analyses. On the contrary, arguments from considerations of meaning would seem to count decisively against Smart's basic thesis of mind/brain identity. The book was concerned primarily with what there is, not with what our linguistic practices may presuppose there is. In what follows, I wish to say something about how this played out and led to the kind of scientific metaphysics that now dominates the philosophy of science in Australia.

For my part, I found Smart's scientific realism to be compatible with my generally physicalist outlook. I was no longer anti-realist—if, indeed, I ever was. My main concerns about Smart's scientific realism were (1) that the identity theory, i.e., the theory identifying sensations with brain processes, did not give an adequate account of the qualities of our sense experiences (the qualia) and (2) that his theory accepted theoretical entities too indiscriminately. I addressed the first of these two concerns in my essay on 'Physicalism and the Contents of Sense Experience' (Ellis 1975). My other main concern about Smart's scientific realism was its casualness in attributing reality to theoretical entities. Smart admitted that some of the theoretical entities of science, such as lines of force, are fictions. But, in general, his attitude appeared to be that we should believe in whatever the scientists do—at least in their capacity as scientists. While being sympathetic to this attitude, I thought we should be a bit more discriminating.

My starting point for developing a more discriminating scientific realism was the Maxwell-Bridgman theory of the real as that which may have several different kinds of effects. If, for example, a theoretical entity such as a field of force is postulated, then the question of whether it is real is just that of whether it is capable of manifesting itself in any way other than as a field of force. That is, does it have any kinds of properties other than those it has by definition? If not, then it is fictional, and, ontologically, it would be better to accept action at a distance (as Bridgman said in 1925). Or, we may ask, are Newtonian forces real? Must we admit them into our ontology just because scientists generally seem to believe in them? As one who had written extensively on the subject, I thought not. For Newtonian forces cannot do anything other than have the kinds of effects they are defined as having, and nothing other than a Newtonian force is capable of having just these kinds of effects. So such forces fail the Maxwell-Bridgman test.

In 1976 I published a paper entitled 'The Existence of Forces', in which I developed independent criteria for physical reality. In that paper, I argued that mass-energy appeared to be the defining characteristic of the physically real. A physical object, for example, is anything that has mass-energy. A physical event is any change of energy distribution in the universe. A physical causal process is any causally connected sequence of physical events. A physical property is any property that makes a difference to some physical causal process. But forces, as they are understood in Newtonian physics, are none of these things. For they do not have energy, do not, in virtue of their existence, involve any change of energy distribution in the universe, are not physical causal processes, and are not physical properties. So, if they are physical entities at all, then they are *sui generis* physical entities. I was more inclined to believe that Newtonian forces simply do not exist and argued that there were a number of independent reasons for holding this.

This theory of the physically real is the one I had in mind when I remarked rather cryptically in *Rational Belief Systems* that I was a 'scientific entity realist'. I did not believe, as most of my colleagues evidently did, that scientific laws are just empirical generalisations that are true in some naïve correspondence sense. For I thought that laws were mostly universal counterfactual conditionals, i.e., statements of the form: 'If anything were an X in circumstances of the kind K, then X would do Y'. But if laws really do have this form, then any analysis of their truth conditions will require reference to sets of possible worlds. And, since I did not believe in possible worlds other than this one, I did not believe that any such statements could be understood simply as descriptive of reality. Rather, I thought that, in some sense, the laws must be understood as describing the underlying structure of the world.

I did not, at that time, have a very clear idea of how the phrase 'the underlying structure of the world' should be understood. I was sure, however, that it could not reasonably be understood as referring only to the most convenient, or even the axiomatically most elegant, approximation to the truth. For if this were what laws of nature were, they would be less fundamental ontologically than the messy facts or crude empirical laws they allegedly explained. But, in any case, I did not think that I had to believe in this absurdity to be a scientific realist. I thought it would be enough if one just believed in the reality of those theoretical entities that passed the Maxwell-Bridgman test for reality (Bridgman, 1927). For the only respectable theoretical entities that failed this test were things such as numbers, sets, forces, geometrical points, perfectly reversible heat engines, and ideal incompressible fluids in steady flow in uniform gravitational fields. And these all seemed to me to be things that no good scientific realist ever seriously believed in.

The major challenges to scientific realism that surfaced in the early 1980s created some heated discussion. Laudan's historical argument (Laudan, 1981) that the laws and theories of the mature sciences are probably not true, and, in many cases, not even approximately so, created problems for those scientific realists who had based their case for realism on the success of science in making the world more predictable. But I will not have anything much to say about this dispute, for it does not deal with the main issue. The main issue is how scientific theories are to be understood. Should they be understood primarily as more or less useful instruments for prediction? Or should they be understood as attempts to describe the underlying reality, on which what is seen to happen in the world ultimately depends? Predictive success does provide an argument for realism, because the simplest explanation for it would be just that the proposed laws and theories are true of the underlying structures of the world. But realism concerning our understanding of the aims of science could survive on quite a modest degree of predictive success. For all that is required for the belief that the scientific picture of reality is the most rational one to accept is that it should be better at predicting what will happen than any other picture. And this, I am sure, has been the case for centuries.

Van Fraassen's philosophical challenge to scientific realism was more to the point, although it was, essentially, just a revival of Duhem's empiricism. Duhem (1914/1954) thought that the aim of science is necessarily limited to describing the world as it appears to us, i.e., to what Kant calls 'the empirical world', and to synthesising our knowledge of it, e.g., by making logico-mathematical models of the observable things or processes. Like Van Fraassen, Duhem believed in the existence of a transcendental reality, i.e., a world beyond appearances that is the real world. But, like many before him, he argued that it is not the task of science to reveal the nature of this reality. That is the task of metaphysics, he said.

Van Fraassen (1980) thought the same. Science, he stated, can never do anything more than 'give us theories which are empirically adequate', and, therefore, acceptance of a theory should never do anything more than 'involve as belief that it is empirically adequate' (p. 12). Van Fraassen's theory was clearly at odds with the philosophies of many of the later positivists, for he did not deny the existence of a transcendental world that science is incapable of describing. In an essay entitled 'What Science Aims To Do', written for a volume of essays on Van Fraassen's constructive empiricism, I argued for the pragmatic thesis that 'Science aims to provide the best possible scientific account of natural phenomena, and acceptance of a scientific theory involves the belief that it belongs to such an account' (Churchland et al. 1985, p. 169).

In retrospect, I think that Duhem, Van Fraassen, and I were all wrong about what science could tell us about the world. Duhem and Van Fraassen were both wrong in thinking that science itself could tell us nothing about the world, other than what accounts of reality are empirically adequate. Science does much more than this: it selects and endorses the best of the empirically most adequate accounts. And the best of these accounts are, for reasons I gave in (Ellis, 1957), normally process theories. I argued then that if the explanations that these theories offer are sound, then they tell us more than their non-process equivalents. They do so by making it possible to establish links between theories that would otherwise not be linked. And, by establishing these links, they increase the connectivity of our knowledge in ways that non-process theories generally cannot match. As I said in my contribution to the Van Fraassen volume, realistic process theories increase the field of evidence for a theory, e.g., by allowing cross-theoretic identifications. As a result, evidence for or against one theory may become evidence for or against another that is linked to it theoretically. To illustrate, the null result of the Michelson-Morley experiment, which was designed to detect differences in the speed of light in different directions, counted decisively against Newton's absolute theory of space and time. But Newton's theory itself had *nothing whatever* to say about the speed of light. It was a system of dynamics for corporeal bodies. The relevance of the Michelson-Morley experiment was due to certain cross-theoretical linkages (viz., between the Newtonian concept of absolute space and the nineteenth-century one of the luminiferous ether).

But I too was wrong about what science can tell us about the world. I had thought that the envisaged scientific worldview would include all of the knowledge that it was possible to have about reality, and so I left no place in my epistemology for any kind knowledge of reality other than scientific knowledge. I did not believe, for instance, that our understanding of the world could be increased by metaphysical speculation. On the contrary, I thought that all such speculation was pretentious. But I no longer think that this is so. On the contrary, I would now say, metaphysics has the same kind of role in improving our knowledge and understanding of the world as scientific theorising. Its methods are not those of the empirical sciences. Metaphysical inquiries can, nevertheless, increase the connectedness of our knowledge and hence contribute usefully to the project of seeking the truth (in both the epistemic and metaphysical senses of this word).

Scientific Metaphysics

Science is limited by what scientists are able to do. In practice, it is restricted by lack of resources, failure to make the required observations, the intellectual limitations of scientists, and in other ways. But let us imagine a world in which all such limitations have been overcome, as if by magic, and let us call the theory of the natural world that science would ideally deliver in such a world 'the scientific worldview'. Then, plausibly, this worldview has some claim to be considered the true one—the one at which we should aim. For, by definition, this is the view of reality that would rationally be accepted on the basis of the best and most comprehensive set of observations that human beings could possibly make. Nevertheless, most philosophers would probably say that even this ideal scientific worldview might not be true. There might be parts of reality that we cannot ever know about. Or, we might, either by accident or design, be systematically deceived about the nature of reality. Or, perhaps, we are just not biologically programmed in the right sort of way to discover the objective truth about the world—even in ideal circumstances. We can, no doubt, discover by scientific investigation many of the things that we (i.e., we human beings) ought rationally to believe and rule out a great many things that it would ultimately be irrational for us to believe. So, even if there are limits to what it is possible for scientists to discover, the aim of discovering all and only those things that it would, in ideal circumstances, be rational for us to believe about the world would seem to be a plausible objective of scientific inquiry.

For many years I assumed that these doubts about the limits of science were not well founded and reflected badly on the metaphysical theory of truth that gave rise to them. Consequently, I accepted the pragmatist theory that identifies truth with what it would ideally be rational to believe and called myself an 'internal realist', as others before me had done. I embraced this position, because the empiricist in me identified science with rational inquiry about the nature of reality. I did not believe that there was any other kind of rational inquiry about reality that could take over where science left off or that scientific knowledge was essentially limited in any way. There might be a theory of science, a logic of science, or an inquiry into the language of science, or into the various kinds of concepts employed in science. But these inquiries were not, I thought, continuations of the scientific quest to understand the nature of reality. They were just metascientific inquiries, i.e., inquiries about the nature of scientific inquiry, which philosophers of science were at least as well equipped as anyone else to undertake. The idea that one could continue the inquiry into the nature of reality by rational means that were not essentially scientific was one that struck me as preposterous. It would be better to abandon the concept of truth as a metaphysical correspondence relationship. I was reinforced in this stance by my conviction that what I called 'the metaphysical concept' was not required for any of the purposes of logic. In *Rational Belief Systems*, I demonstrated that the standard deductive logics, including all of the quantified, modal, and conditional ones, could all be founded adequately in a theory of rationality. No semantic concept of truth is required for this purpose, just some more or less self-evident principles of rationality based on a conception of truth as epistemic rightness.

In retrospect, I now recognise that it was a mistake to abandon the metaphysical concept of truth. For there are important questions about the nature of reality that cannot, even in principle, be resolved by the methods of science. But I did not see this at the time. Analytic philosophy, which dominated twentieth-century philosophy of science, was largely the attempt to explain the nature and structure of scientific laws, theories, and explanations and to specify empirically adequate truth conditions for the various kinds of claims that scientists make in expressing their conclusions. These inquiries were not scientific ones. They were just attempts to understand better the work that scientists do and the nature of their achievements. This does not necessarily make them metaphysical inquiries of the sort that I thought were pretentious. But they did presuppose theories of knowledge and understanding that were not themselves scientific findings. They depended, for example, on Frege's conception of logic as the theory of truth preservation. According to Frege's theory, arguments are valid if and only if there is no possible world in which their premises are true and their conclusions false. Therefore, it was argued, if we are to understand any statement sufficiently for all of the purposes of logic, it is necessary to know the truth conditions of its premises and conclusion in all possible worlds.

These inquiries also depended on the acceptance of certain paradigms of knowledge and understanding. In the early days of logical positivism, basic observation statements, e.g., the statement that the object A has the characteristic C (where A and C are both directly observable), were held to be both transparently clear and knowable. Therefore, any statement of truth conditions acceptable to the generation or so of philosophers of science involved in the positivist program of analysis would have to have been a simple truth function of basic observation statements. The statements of the analysans were required to be both formally adequate and empirically ascertainable, i.e., ones that could in principle be discovered to be true or false directly by observation. In practice, however, such analyses were rarely attempted. Usually, it was thought to be enough if formulae for producing such analyses could be specified. How, for example, was a statement of the form 'A causes B' to be analysed? What, in general, are the empirically adequate truth conditions for such statements? Most philosophers of science were convinced by Hume's arguments that no such statements could ever be accepted as truly basic, i.e., as atomic propositions. Therefore, the attempt had to be made to discover their empirically adequate truth conditions.

But propositions attributing causal connections were not the only problematic ones for the logical positivists. Counterfactual conditionals describing the ways in which ideal objects would behave in certain possible circumstances were also problematic. So also were propositions assigning causal powers, capacities, or objective probabilities to things. And what were logical positivists to make of the physical necessities and possibilities that evidently do exist in nature? According to the Second Law of Thermodynamics, it is impossible to build a perpetual motion machine (of the second kind). What are the empirical truth conditions for this statement? Indeed, what are the truth conditions for 'X is a law of nature'? A great deal of work went into trying to answer these questions. However, answers acceptable to the logical positivists were not to be found. Consequently, their program of empirical analysis was largely abandoned in favour of one of semantic analysis, which was much less demanding. Semantic analysts still looked for formally adequate truth conditions, but they abandoned the requirement of empirical adequacy that logical positivists had formerly insisted upon.

To accept the semantic analyses of modals and conditionals of the sorts that have been widely used in philosophical logic since the 1960s, it seemed necessary to accept that the truth or falsity of such propositions depends not only on what there is in the actual world but also on what exists in other possible worlds and on how these other possible worlds are related to the actual one. It is possible to think of this theory as just a formal model that happens to be useful for developing logics of modals and conditionals. But to do so would be to deprive the semantic theory of any explanatory power. If the model has no basis in reality, why be guided by it? David Lewis and his many followers in Australia all boldly accepted realism about possible worlds. That is, they embraced the idea that the actual world—the one we happen to inhabit—is just one of an infinity of possible worlds, all of which are real. Moreover, they accepted that these possible worlds all exist necessarily. And, having accepted this incredible thesis, they had to suppose that the truth or falsity of modal and counterfactual conditional propositions couldn't in general depend only on what exists in the actual world. In most cases, they are required to say that the truth or falsity of a modal or conditional depends on what there is in other possible worlds and on how these other worlds are related to this one.

For me, there are no real possible worlds other than the actual one. But there are many more or less rational belief systems concerning it, and, using the theory of rational belief systems, we can easily construct ideally rational belief systems that have just the kinds of properties we seek. For example, if we wish to consider an ideally rational belief system that is as much like our own limited one as we can make it, but in which p is accepted as true, then we may easily do so, even if we ourselves believe that not-p. And, we may then use the theory of rational belief systems to determine whether q could rationally be denied in such a system. If not, then, according to the theory, the conditional '$p \Rightarrow q$' must rationally be accepted by us as true. In general, to found a satisfactory propositional logic or predicate calculus, or to introduce modal operators and conditional connectives into a logical system, all one needs to do is develop appropriate axiom systems for rational belief systems on languages that have the relevant structures, and, effectively, to define the connectives and operators by the acceptability conditions for propositions that include them. Thus, acceptability conditions can replace truth conditions in the foundations of logic, and ideally completed rational belief systems can replace possible worlds in the theories of modals and conditionals. The only price one has to pay for this is that one has to abandon the implausible idea that logic is the theory of truth preservation. It is not: it is, both intuitively and in reality, just a basic part of the theory of rational belief systems.

In developing the theory of rational belief systems, I approached the question of how people ought rationally to think about the world as I imagined a scientist would. I was fully aware that ordinary human belief systems are incomplete, messy, confused, and contradictory and that human reasoning is often fallacious. Indeed, it is not just fallacious in random ways, but systematically so, as cognitive psychologists have convincingly shown. Nevertheless, there appear to be certain underlying patterns of human thought and reasoning that are universal. And, I thought that these deep structures might be used to construct a model-theoretic ideal of human rationality. It is, after all, standard scientific practice to look for such patterns and, where possible, to use them like this for such purposes. The resulting scientific theory, I argued, is one that enables us to develop epistemological foundations for all of the standard logical systems. So, as a theory, it was highly successful. But despite the success of this project, I found myself becoming increasingly isolated philosophically. No one else, to my knowledge, ever accepted that the theory of rational belief systems provides an adequate foundation for standard logical theory. Yet, this thesis was not refuted in the literature or even much criticised. In fact, it was all but ignored. Philosophers went on believing in real but non-actual possible worlds or that someone would someday tell them what these theoretical entities really are, without them having to give up on the Fregean conception of logic as the theory of truth preservation. The main influence that the theory of rational belief systems had in philosophy was just that it served as a springboard for the development of theories of the dynamics of belief. Peter Forrest (1986) and Peter Gärdenfors (1984) led the way in this area.

I was also more or less alone in Australia in defending the theory of truth as epistemic rightness. Richard Rorty liked this theory and wrote to me after the publication of my paper (Ellis, 1970) 'Truth as a Mode of Evaluation' to congratulate me. But most Australians were wedded to the idea of truth as a semantic relationship, i.e., a relationship between words and the world. For this was the theory of truth they thought a realist would just have to accept. Nevertheless, I defended the evaluative theory (a) because the theory of rational belief systems evidently required a concept of truth as a mode of epistemic evaluation; (b) because the pragmatic contradictions in 'It is true, but I don't believe it' and 'I believe it, but it isn't true' are best explained this way; and (c) because I could not see that anything would be lost if we were all to use such a conception. Nevertheless, from the mid-1980s, I began to have serious doubts about the adequacy of the theory of truth that I had embraced. For it was, essentially, an intersubjectivist theory. If truth for me were just what I thought it was right for me to believe, then that would be purely subjectivist. But is the intersubjectivist theory that truth is what is true for us at the limit of experience really much better? It might be better than the best we can ever hope to achieve through scientific inquiry. But it is still not an objective concept of the sort demanded by Australian realists.

In my book *Truth and Objectivity* (Ellis, 1990), I made a final attempt to rescue the theory of truth as a mode of evaluation and hence to justify the concept of truth that I required for my theory of rational belief systems. The consensus is that I failed in this attempt, and in retrospect, I also think I failed. For I now think that there are two quite legitimate, but related, conceptions of truth with similar logics, just as there are two or more legitimate but related conceptions of probability (empirical, logical, rationalised subjective) that satisfy the axioms of the same probability calculus. The concept of truth as epistemic rightness is the one that is required for human belief systems and hence for logic and science. The metaphysical concept of truth is the one required for truthmaker theory. Consider John Fox's 'Truthmaker' axiom, viz.,

> If *p*, some *x* exists such that *x*'s existing necessitates *p*. (Fox 1987, p. 189)

or John Bigelow's supervenience thesis:

> [T]here is no difference in what is true without a corresponding difference in the inventory of what is; that what there is determines what is true; that truth is supervenient on being. (Fox 1987, p. 205)

These two theses are both very plausible. But neither is suggested nor even rendered plausible by the theory of truth that I had been defending. For my theory of truth has no obvious implications concerning existence. It is, for example, as readily applicable to the theorems of mathematics as it is to the fundamental laws of physics. For example, to decide the question of whether a given mathematical proposition is true in, say, Euclidean geometry, one only has to consider whether it has a sound Euclidean proof. One does not have to think about what exists in reality. If there is such a proof, then the proposition is true in my sense. No further argument. Whether and if so how it corresponds to reality are other matters.

Until about 1990, I had thought it was sufficient to argue for realism as an extension of the argument for physicalism. If you accept a scientific worldview, then you are bound to be a realist about most of the causal processes that are supposed to occur in nature. If the effects to be explained are real, which they undoubtedly are, then so must be their causes. The scientific worldview thus contains an ontology of its own, independently of any theory of truth. And, it certainly implies realism about all, or nearly all, of the sorts of things that scientific realists say they believe in. I called myself 'a scientific entity realist' in the early 1980s—mainly to distinguish myself from those who thought that scientific realism necessarily involves the belief that the established laws and theories of science are mostly true in a substantial correspondence sense of 'truth'. I was, for the most part, willing to accept that the established laws and theories of science were true in my weak evaluative sense of 'true'. But I was not willing to accept that the theoretical entities of abstract model theories (e.g., Newtonian forces, geometrical points, inertial frames, perfectly reversible heat engines) had anything like the same status as the atoms, molecules, and electromagnetic waves of the established causal process theories of physics and chemistry. But I could still be a scientific realist, I argued, because no theory of truth was required for realism about the established causes of things. And this, for me, was enough, because there was no plausible theory of truth, of which I was aware, that would imply realism about the theoretical entities of abstract model theories.

I further explored the idea of deriving one's ontology from the scientific worldview in my paper (Ellis, 1987) on 'The Ontology of Scientific Realism', which I wrote as my contribution to a volume of essays written in honour of Jack Smart. What sort of ontology, I wondered, do Smart's original arguments for scientific realism really imply? As well as realism about the causal mechanisms proposed in successful scientific theories, and hence about the theoretical entities postulated as being involved in these mechanisms, I argued that they also require realism about the causal powers that these things are supposed to have, about the spatiotemporal relationships that the parts of these mechanisms are supposed to bear to each other, and about the numerical relationships that are supposed to hold between various groups of elements occurring in these mechanisms. The more I thought about what acceptance of the scientific worldview implies for ontology, the richer my ontology became. So, I concluded that Smart's original arguments for scientific realism should have led him, as it eventually led me, to reject the austere Humean ontology that he persisted with throughout his career.

The ontology required for a scientific worldview appears to be a highly structured one. For one of the most striking facts about the world is the extraordinary dominance of natural kinds. Every different chemical substance (and there are hundreds of thousands of them) is a member of a natural kind: (a) each kind of chemical substance is categorically distinct from all others, and (b) each has its own essential properties and structures. Moreover, the chemical kinds all belong in a natural hierarchy, the more general ones having essences that are included in those of the more specific. Plausibly, the existence of this hierarchy of natural kinds is a significant fact about the world that should be reflected in the ontology of scientific realism. The world is evidently not just a physical world, as I had assumed in the 1970s, but a highly structured one. Perhaps the world itself is a member of a natural kind. John Bigelow, Caroline Lierse, and I published a joint paper on this topic in Bigelow et al. 1992.

Shortly after our collaboration on this paper, John Bigelow was appointed to a chair at Monash University, and I inherited Caroline Lierse as a graduate student. Caroline was enthusiastic about the kind of work I was doing on essentialism and natural kinds and was keen to collaborate on other projects in this area. The issue that interested me most at the time was the status of dispositional properties. Most philosophers in the Anglo-American tradition regarded dispositions as second-grade properties. They were second grade, it was argued, because they had ultimately to be grounded in categorical properties. But even a quick survey of the kinds of properties that have significant roles in the causal process theories of the sciences reveals that most of them are dispositional. Indeed, the most fundamental properties of objects would all appear to be dispositional. Massive bodies always appear to have certain gravitational and inertial powers and to manifest themselves to us in the ways in which they exercise them. So, we may ask: What is it that makes a body massive? The standard answer is that massive bodies all have the quantitative property we call 'mass' to some degree. But if we inquire further what gives a body mass, we may find ourselves without an answer. Is it, for example, the numbers of atoms of the various kinds that make up these bodies, multiplied by the masses of these atoms? No, it is not that. But even if it were, we should only have explained the masses of the bodies by reference to the masses of their constituents. But then how should we account for the masses of the most fundamental constituents? A causal power, like the mass of a body, can be dependent on the causal powers of its constituents. But a causal power can never be dependent on anything that does not have any causal powers. And if matter has an ultimate atomic structure, then we must eventually get down to things that have causal powers that do not depend on the causal powers of their parts. Perhaps causal power dependencies go all the way down—to the parts of the parts of the parts, and so on. Or must we say, as Hume would have said, that the causal powers of things are illusions due to regularities? I think the best answer is that causal powers such as mass are not illusions and that if the question ‘Why do things have mass?’ can eventually be answered, it will be because the causal powers of massive bodies can be shown to be dependent on other causal powers. Therefore, at the most fundamental level, there must be some irreducible causal powers.

Accepting this conclusion, Lierse and I wrote a paper on dispositional properties, (Ellis, et al. 1994) i.e., properties, such as causal powers, that dispose their bearers to behave in certain ways or ranges of ways. We approached the subject believing that there are, in reality, two kinds of properties, dispositional and categorical. But we did not accept any of the theories of dispositional properties that were then currently on offer. Specifically, we argued against Armstrong’s strong categorialism, i.e., the thesis that all basic properties are categorical, and also against Shoemaker’s strong dispositionalism, according to which all genuine properties are dispositional. Our position was dispositionalist about causal powers, capacities, and propensities, but categorialist about spatiotemporal and numerical relations. We argued against the three theses concerning dispositions (Prior et al. 1982) that had been proposed and defended by Elizabeth Prior, Robert Pargetter, and Frank Jackson, and we defended the following more radical theses: (a) that there are real irreducible dispositional properties in nature, e.g., causal powers; (b) that causal powers necessarily dispose their bearers to produce effects of certain kinds in certain kinds of circumstances; (c) that such properties are among the essential properties of the natural kinds and are not necessarily grounded in other properties; and (d) that if P is a causal power that is an essential property of things of the natural kind K and L is the law of action of P, to the effect that things that possess P are necessarily disposed to have the effect E in the circumstances C, then it is metaphysically necessary that things of the kind K will act as L requires. We called our position ‘dispositional essentialism’.

In my most recent books (Ellis 2001, 2002), I have elaborated and defended an essentialist ontology, which builds upon this earlier work. In these two books, I argue that the world is a physical one that is structured into hierarchies of natural kinds. There are three categories of such kinds, I argued: substantive, dynamic, and tropic; and within each category, there are various genera and species. The substantive category includes all of the natural kinds of objects or substances; the dynamic category, all of the natural kinds of events or processes; and the tropic category, the natural kinds of tropes (i.e., property or relation instances) of the properties or relations that hold of or between things. At the summit of each category, I supposed there to be a global kind, i.e., a natural kind that includes all of the more specific natural kinds within the category. The global kind of substance, for example, would be the class of physical systems, while the global dynamic kind would include the whole category of physical events or processes. I then argue that the laws of nature may reasonably be identified as true descriptions of the essential properties of the natural kinds. Granted this, it follows that there must be natural hierarchies of laws of nature, with the global laws describing the essential properties of the global kinds and the more specific laws describing the essential properties of the more specific kinds. Thus, if all physical systems are essentially Lagrangian (i.e., obey Lagrange's Principle of Least Action), then the Principle of Least Action will be a universal law of nature. And, because of this, it will be metaphysically necessary, not contingent as most philosophers suppose. In a similar way, the essential properties of the more specific kinds will give rise to the more specific laws of nature, e.g., those relating to particular kinds of substances or particular kinds of fields. And these laws too must be metaphysically necessary.

This theory of the laws of nature has a number of profound implications. Firstly, it implies that the laws of nature are firmly grounded in the hierarchically structured physical world we see around us. They are not, as Hume thought and most other philosophers still think, imposed upon an intrinsically passive world, as if by God. They are grounded in the things that exist in nature. And, if you could somehow change what there is, you would thereby change the laws of nature. But it is metaphysically impossible to change the laws of nature without changing the world's ontology.

References

Ayer, A. (1936). *Language, truth and logic*. London: Gollancz.

Bigelow, J., Ellis, B., & Lierse, C. (1992). The world as one of a kind: Natural necessity and laws of nature. *The British Journal for the Philosophy of Science, 43*, 371–388.

Bridgman, P. (1927). *The logic of modern physics*. New York: Macmillan.

Bridgman, P. (1962). *A sophisticate's primer of relativity*. London: Routledge & Kegan Paul.

Campbell, N. (1921). *Physics: The elements*. New York: Dover. Republished 1957 as Foundations of Science: The Philosophy of Theory and Experiment.

Churchland, P., & Hooker, C. (Eds.). (1985). *Images of science: Essays on realism and empiricism, with a reply from Bas C. Van Fraassen*. Chicago: University of Chicago Press.

Colodny, R. (Ed.). (1965). *Beyond the edge of certainty: Essays in contemporary science and philosophy*. Englewood Cliffs: Prentice-Hall.

Duhem, P. (1954). *The aim and structure of physical theory*. Princeton: Princeton University Press. Trans. P. Wiener.

Ellis, B. (1957). A comparison of process and non-process theories in the physical sciences. *The British Journal for the Philosophy of Science, 8*, 45–56.

Ellis, B. (1963). Universal and differential forces. *The British Journal for the Philosophy of Science, 14*, 177–194.

Ellis, B. (1966). *Basic concepts of measurement*. Cambridge: Cambridge University Press.

Ellis, B. (1971). On conventionality and simultaneity: A reply. *Australasian Journal of Philosophy, 49*, 177–203.

Ellis, B. (1975). Physicalism and the contents of sense experience. In C.-Y. Cheng (Ed.), *Philosophical aspects of the mind-body problem* (pp. 64–77). Honolulu: University Press of Hawaii.

Ellis, B. (1979). *Rational belief systems*. Cambridge: Blackwell.

Ellis, B. (1980). Truth as a mode of evaluation. *Pacific Philosophical Quarterly, 1*, 85–99.

Ellis, B. (1985). What science aims to do. In P. Churchland & C. Hooker (Eds.), *Images of science: Essays on realism and empiricism, with a reply by Bas C. Van Fraassen* (pp. 166–193). Chicago: University of Chicago Press.

Ellis, B. (1987). The ontology of scientific realism. In P. Pettit, R. Sylvan, & J. Norman (Eds.), *Metaphysics and morality: Essays in honour of J. J. C. Smart* (pp. 50–70). Cambridge: Blackwell.

Ellis, B. (1990). *Truth and objectivity*. Cambridge: Blackwell.

Ellis, B. (2001). *Scientific essentialism*. Cambridge: Cambridge University Press.

Ellis, B. (2002). *The philosophy of nature: A guide to the new essentialism*. London: Acumen.

Ellis, B., & Bowman, P. (1967). Conventionality in distant simultaneity. *Philosophy in Science, 34*, 116–136.

Ellis, B., & Lierse, C. (1994). Dispositional essentialism. *Australasian Journal of Philosophy, 72*, 27–45.

Forrest, P. (1986). *The dynamics of belief*. Cambridge: Blackwell.

Fox, J. (1987). Truthmaker. *Australasian Journal of Philosophy, 65*, 188–207.

Fox, J. (2007). Why we shouldn't give Ellis a dinch. *Analysis, 67*, 301–303.

Gärdenfors, P. (1986). The epistemic importance of minimal changes of belief. *Australasian Journal of Philosophy, 62,* 136–157.

Gasking, D. (1940). Mathematics and the world. *Australasian Journal of Philosophy, 18*, 97–116.

Gasking, D. (1960). Clusters. *Australasian Journal of Philosophy, 38*, 1–36.

Grünbaum, A. (1964). Is a universal nocturnal expansion falsifiable or physically vacuous? *Philosophical Studies, 15*, 71–79.

Grünbaum, A. (1967). The denial of absolute space and the hypothesis of a universal nocturnal expansion. *Australasian Journal of Philosophy, 45*, 61–91.

Grünbaum, A., et al. (1969). Panel discussion of simultaneity by slow clock transport in the special and general theories of relativity. *Philosophy in Science, 36*, 1–81.

Kuhn, T. (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.

Lakatos, I. (1963–1964). Proofs and refutations. *British Journal for the Philosophy of Science*, *14*: 1–25, 120–39, 221–45, 269–342.

Lakatos, I. (1970). Falsification and the methodology of scientific research programs. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge* (pp. 91–196). Cambridge: Cambridge University Press.

Laudan, L. (1981). A confutation of convergent realism. *Philosophy in Science, 48*, 19–49.

Poincaré, H. (1952). *Science and hypothesis*. New York: Dover.

Popper, K. (1959). *The logic of scientific discovery*. New York: Basic Books.

Prior, E., Pargetter, R., & Jackson, F. (1982). Three theses about dispositions. *American Philosophical Quarterly, 19*, 251–257.

Putnam, H. (1962). The analytic and the synthetic. In H. Feigl & G. Maxwell (Eds.), *Scientific explanation, space, and time* (pp. 358–397). Minneapolis: Minnesota University Press.

Quine, W. (1953). *From a logical point of view: Nine logico-philosophical essays*. Cambridge: Harvard University Press.

Reichenbach, H. (1958). *The philosophy of space and time*. New York: Dover. Trans. M. Reichenbach and J. Freund.

Rorty, R. (Ed.). (1967). *The linguistic turn: Recent essays in philosophical method*. Chicago: University of Chicago Press.

Schlesinger, G. (1964). It is false that overnight everything has doubled in size. *Philosophical Studies, 15*, 65–71.
Schlesinger, G. (1967). What does the denial of absolute space mean? *Australasian Journal of Philosophy, 45*, 44–60.
Smart, J. (1963). *Philosophy and scientific realism*. London: Routledge & Kegan Paul.
Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in logical theory* (pp. 98–112). Oxford: Blackwell.
van Fraassen, B. (1980). *The scientific image*. Oxford: Clarendon.