

14

PHILOSOPHY OF SCIENCE

Stathis Psillos

Synoptic overview

Philosophy of science emerged as a distinctive part of philosophy in the twentieth century. It set its own agenda, the systematic study of the metaphysical and epistemological foundations of science, and acquired its own professional structure, departments and journals. Its defining moment was the meeting (and the clash) of two courses of events: the breakdown of the Kantian philosophical tradition and the crisis in the sciences and mathematics at the beginning of the century. The emergence of the new Frege–Russell logic (see “The birth of analytic philosophy,” Chapter 1), the arithmetization of geometry and the collapse of classical mechanics called into question the neat Kantian scheme of synthetic a priori principles (see “Kant in the twentieth century,” Chapter 4). But the thought that *some* a priori (framework) principles should be in place in order for science to be possible had still had a strong grip on the thinkers of the European continent. A heated intellectual debate started concerning the status of these a priori principles. The view that dominated the scene after the dust had settled was that the required framework principles were conventions. The seed of this thought was found in Poincaré’s writings, but in the hands of the logical positivists, it was fertilized with Frege’s conception of analyticity and Hilbert’s conception of implicit definitions. The consolidation of modern physics lent credence to the view that a priori principles can be revised; hence, a new conception of *relativized* a priori emerged. The linguistic turn in philosophy reoriented the subject matter of philosophical thinking about science to the *language* of science (see “The development of analytic philosophy: Wittgenstein and after,” Chapter 2). Formal-logical methods and conceptual analysis were taken to be the privileged philosophical tools. Not only, it was thought, do they clarify and perhaps solve (or dissolve) traditional philosophical problems; they also make philosophy rigorous and set it apart from empirical science. In the 1930s, philosophy of science became the logic of science; it became synonymous with anti-psychologism, anti-historicism, and anti-naturalism. At the same time, philosophy of science, in Vienna and elsewhere, was completing the project of the Enlightenment: the safeguarding of the objectivity and epistemic authority of science.

The havoc brought about by the Nazis liquidated most philosophical thinking on the Continent and many continental philosophers of science took refuge in the US. There, their thought came under the pressure of American pragmatism. Pragmatism's disdain for drawing sharp distinctions where perfect continua exist shook up the rationale for doing philosophy of science in the way the logical positivists practised it. Quine challenged the fact/framework distinction and argued that no a priori principles were necessary for science. Sellars refuted a foundationalist strand that was present in the thought of some logical positivists. And Kuhn restored the place of history in philosophical thinking about science. By the 1960s, philosophy of science had seen the advent of psychologism, naturalism, and history. The historical turn showed that attempted rational reconstructions of science were paper tigers. Yet the historicists' grand models of science were no less of paper tigers, if only because the individual sciences are not uniform enough to be lumped under grand macro-models. The 1980s saw the mushrooming of interest in the individual sciences. The renaissance of scientific realism in the 1960s resulted in an epistemic optimism with regard to science's claim to truth, though new forms of empiricism emerged in the 1980s. At the same time causation came out of the empiricist closet. In the last two decades of the century, philosophers of science had started taking seriously a number of traditional metaphysical issues that were considered meaningless in the 1930s.¹

Pressures on Kantianism

Much of the philosophical thinking before the beginning of the twentieth century had been shaped by Kant's philosophy. Immanuel Kant (1724–1804) had rejected empiricism (which denied the active role of the mind in understanding and representing the world of experience) and uncritical rationalism (which did acknowledge the active role of the mind but gave it an almost unlimited power to arrive at substantive knowledge of the world on the basis only of the lights of Reason). He famously claimed that although all knowledge starts with experience it does not arise from it: it is actively shaped by the categories of the understanding and the forms of pure intuition (space and time). The mind, as it were, imposes some conceptual structure onto the world, without which no experience could be possible. There was a notorious drawback, however. Kant thought there could be no knowledge of things as they were in themselves (the *noumena*) and only knowledge of things as they appeared to us (*phenomena*). This odd combination, Kant thought, might well be an inevitable price one has to pay in order to defeat skepticism and forgo traditional idealism. Be that as it may, his master thought was that some synthetic a priori principles should be in place for experience to be possible. These synthetic a priori principles (e.g. that space is Euclidean, that every event has a cause, that nature is law-governed, that substance is conserved, the laws of arithmetic) were necessary for the very possibility of science and of Newtonian mechanics in particular. If we were to sum up Kant's conception of a priori knowledge, this would be helpful: it is knowledge

- 1 universal, necessary and certain;
- 2 whose content is formal: it establishes conceptual connections (if analytic); it captures the form of pure intuition and formal features that phenomena have

because they are partly constituted by the pure concepts of the understanding (if synthetic);

3 constitutive of the form of experience;

4 disconnected from the content of experience; hence, unrevisable.

The new physics

A century after Kant's death, a major crisis swept across the reigning Newtonian physics. Classical mechanics crumbled. Two kinds of pressure were exerted on it.

The first came from Einstein's Special Theory of Relativity in 1905. Drawing on important considerations of symmetry, Albert Einstein (1879–1955) suggested that understanding the electrodynamics of moving bodies required a radical departure from classical mechanics. Where classical mechanics, and its extension to the electrodynamics of moving bodies by Hendrik Antoon Lorentz (1853–1928), had relied on the existence of absolute space and time, Einstein showed that no such commitment was necessary. Indeed, by taking the concept of simultaneity to be relative to a frame of reference, on the basis of the postulate that the speed of light is constant and identical in all reference frames, he showed that there was no such thing as *the* time in which an event happened. And, by postulating that laws of nature must remain invariant in all inertial frames, he showed that there was no need to posit an absolute frame of reference, typically associated with the absolute space (and the aether). In the able hands of the mathematician Hermann Minkowski (1864–1909), space and time were united in a four-dimensional spacetime framework. In 1915 Einstein advanced his General Theory of Relativity, according to which the laws of nature are the same in *all* frames of reference.

The second kind of pressure came from Max Planck's quantum of action in 1900. He showed that the explanation of a number of phenomena which were within the purview of classical mechanics required admitting a radical discontinuity: energy comes in fundamental quanta. In 1905, Einstein used Planck's idea to explain the photoelectric effect, suggesting that light radiation comes in quantized photons, while many other physicists employed it to develop an alternative to classical mechanics, known as (old) quantum theory. Its culmination came in 1913, where Neils Bohr (1865–1962) explained the structure and the stability of atoms, positing that electrons orbit the nucleus in discrete orbits. By the 1920s, Newtonian mechanics had given its place to Quantum Theory and the General Theory of Relativity.

Geometry and arithmetic

It had been already known, from the work of Nikolai Ivanovich Lobachevsky (1792–1856), János Bolyai (1802–60), and Bernhard Riemann (1826–66), that there could be consistent *geometrical* systems which represented non-Euclidean geometries. Euclid's fifth postulate, that from a point outside a line exactly one line parallel to this can be drawn, can be denied in two ways. Lobachevsky and Bolyai developed a consistent (hyperbolic) geometry which assumed that an infinite number of parallel

lines could be drawn. Riemann developed a consistent (spherical) geometry which assumed that no parallel lines could be drawn. These non-Euclidean geometries were originally admitted as interesting mathematical systems. The Kantian thought that the geometry of physical space had to be Euclidean was taken as unassailable. Yet Einstein's General Theory suggested that this Kantian thought was an illusion. Far from being flat, as Euclidean geometry required, space – that is, *physical* space – is curved. Actually, it's a space with variable curvature, the latter depending on the distribution of mass in the universe. This was a far-reaching result. All three geometries (Euclidean, Lobachevskyan, and Riemannian) posited spaces of constant curvature: zero, negative, and positive respectively. They all presupposed the Helmholtz–Lie axiom of free mobility, which, in effect, assumes that space is homogeneous. Einstein's General Theory called this axiom into question: objects in spacetime move along geodesics whose curvature is variable.

Another Kantian thought that came under fire was that space and time were the a priori forms of pure intuition. In *The Foundations of Arithmetic* (1884), Gottlob Frege (1848–1925) challenged the thought that *arithmetic* was synthetic a priori knowledge. He suggested that arithmetical truths are analytic: they are provable on the basis of general logical laws plus definitions. No intuition was necessary for establishing, and getting to know, arithmetical truths: rigorous deductive proof was enough. Arithmetic, to be sure, was still taken to embody a corpus of a priori truths. But being, in effect, logical truths, they make no empirical claims whatsoever. Frege, however, agreed with Kant that geometrical truths were synthetic a priori.

It was David Hilbert (1862–1943) in *The Foundations of Geometry* (1899) who excised all intuition from geometry. He advanced an axiomatization of geometry and showed that the proof of geometrical theorems could be achieved by strictly logical-deductive means, without *any* appeal to intuition. Hilbert's result was far-reaching because, among other things, it made available the idea of an *implicit definition*. A set of axioms implicitly defines its basic concepts in the sense that it specifies their interpretation collectively: any system of entities that satisfies the axioms is characterized by these axioms. There is no need to have an “independent” or intuitive grasp of the meanings of terms such as “point,” “line,” or “plane”: their meaning is fully specified collectively by the relevant axioms.

These developments in physics and mathematics, mixed up with the dominant Kantian tradition, created an explosive philosophical brew. The instability of the Kantian synthetic a priori principles became evident. Principles that were neatly classified in light of the three Kantian categories (analytic a priori; synthetic a priori; synthetic a posteriori) started to move about. In particular, the category of synthetic a priori truths was being drained away. The crisis in the sciences made possible the claim that the basic principles of classical mechanics and of Euclidean geometry were not a priori since they were *revisable* (they were revised). The crisis in arithmetic suggested that arithmetical truths, alongside the logical ones, were analytic. It wouldn't be an exaggeration to claim that much of philosophy of science in the first half of the twentieth century was an attempt to come to terms with the apparent collapse of the Kantian synthetic a priori.

Convention and the relativized a priori

The Kantian conception did find some support in the work of the neo-Kantian school of Marburg. In *Substance and Function* (1910) Ernst Cassirer (1874–1945) argued that, although mathematical structures were necessary for experience, in that the phenomena could be identified, organized, and structured *only* if they were embedded in such structures, these structures need not be fixed for all time and immutable. He thought that mathematical structures, though a priori (since they are required for objective experience), are revisable yet convergent: newer structures accommodate within themselves old ones. But it was Hans Reichenbach (1891–1953) in *The Theory of Relativity and A Priori Knowledge* (1921) who set the agenda for what was to come by unpacking two elements in the Kantian conception of the a priori: on the one hand, a priori truths are meant to be necessary; on the other hand, they are meant to be constitutive of the object of knowledge. Reichenbach rejected the first element of a priori knowledge, but insisted that the second was inescapable. Knowledge of the physical world, he thought, requires principles of coordination, namely, principles that connect the basic concepts of the theory with reality. These principles were taken to be constitutive of experience. Mathematics, Reichenbach thought, was indispensable precisely because it provided a framework of general rules according to which the coordination between scientific concepts and reality takes place. Once this framework is in place, a theory is presented as an axiomatic system, whose basic axioms – the “axioms of connection” – are broadly empirical in that they specify the relationships among certain physical-state variables. The axioms of coordination, Reichenbach thought, logically precede the most general axioms of connection: they are a priori, and yet revisable. Against Kant, Reichenbach argued that a priori principles, though constitutive of the object of knowledge, could be rationally revised in response to experience.

How is this possible? In *Science and Hypothesis* (1902), Henri Poincaré (1854–1912) had shown that a theoretical system *T* that comprised Euclidean geometry plus some strange physics (which allowed the length of all bodies to be distorted when they were in motion) was empirically indistinguishable from another theoretical system *T'* which comprised non-Euclidean geometry plus normal physics (which admitted rigid motion). Poincaré’s own suggestion was that the choice between such systems (and hence the adoption of a certain geometry as *the* physical geometry) was, by and large, a matter of convention. Indeed, Poincaré extended his geometrical conventionalism further by arguing that the principles of mechanics too were conventions (or definitions in disguise). His starting point was that the principles of mechanics were not a priori truths, since they could not be known independently of experience. Nor were they generalizations of experimental facts. The idealized systems to which these principles apply are not to be found in nature. Besides, no experience can conclusively confirm, or falsify, a principle of mechanics.

Poincaréan conventions are general principles which are held true, but whose truth can neither be the product of a priori reasoning nor be established by a posteriori investigation. But calling them “conventions” did not imply, for Poincaré, that

their adoption (or choice) was arbitrary. He repeatedly stressed that some principles were more convenient than others. He thought that considerations of simplicity and unity could and should “guide” the relevant choice. Indeed, he envisaged a certain hierarchy of the sciences, according to which the very possibility of empirical and testable physical science requires that there are in place (as, in the end, freely chosen conventions) the axioms of Euclidean geometry and the principles of Newtonian mechanics.

Reichenbach was influenced by this Poincaréan view, but gave it a double twist. On the one hand, he took it to imply that seemingly unrevisable principles could be revised. On the other hand, he thought that Poincaré’s own particular suggestion regarding Euclidean geometry was wrong. For whereas going for T (thereby admitting the existence of universal forces that affect indiscriminately and uniformly all moving bodies) does not allow for a unique coordination between theory and reality, the alternative framework T' (which revises the geometry of physical space) does allow for unique coordination. Einstein’s General Theory of Relativity, Reichenbach thought, falsified the claim that the geometry of physical space was Euclidean (thereby leading to a deep revision of a seemingly unassailable a priori principle). His prime thought was that *within* Einstein’s General Theory, geometry and physics become one. Hence, there is no space to play with the one at the expense of the other, as Poincaré thought.

Drawing on the thought of Poincaré and Pierre Duhem (1861–1916), Moritz Schlick (1918) argued that no theory could be tested in isolation from others. Accordingly, when geometry and physics are tested together, if there is a conflict with experience, there is space for choice about which component is to be revised, the choice being guided, by and large, by considerations of simplicity. On Schlick’s early view, there is no principled difference between constitutive principles (e.g. the principles of geometry) and other principles (e.g. the laws of physics). Reichenbach’s thought, however, was importantly different. Although he too accepted that both kinds of principle were revisable, he argued that there was a *logical* difference between them. The constitutive principles marked off the a priori, whereas the axioms of connection marked off the empirical. Reichenbach was naturally led to the conclusion that the only workable notion of a priori should be relativized: each and every theoretical framework needs some constitutive (and hence, a priori) principles, but these principles are revisable, and revised when the framework fails to come to terms with experience. Unfortunately, Reichenbach abandoned this distinction. Following Schlick in the way his thought developed, Reichenbach (1928) was led to adopt a more full-blown conventionalism, according to which even in light of Einstein’s theory, the choice of the geometry of physical space was conventional.

Logical positivism

Armed with the notion of convention, Schlick (1882–1936) and his Vienna Circle tried to show that there could be no synthetic a priori at all. They extended conventionalism to logic and mathematics, arguing that the only distinction there is is

between empirical (synthetic a posteriori) principles and conventional (analytic a priori) ones. The logical positivists thought they were completing Frege's agenda. Like him, they took logic and arithmetic to embody analytic truths. Unlike him, and informed by Hilbert's arithmetization of geometry, they took (pure) geometry too to embody analytic truths. So they fully rejected synthetic a priori knowledge. All a priori knowledge was analytic and, thanks to their conventionalist account of analyticity, grasping a priori (= analytic) truths required no special faculty of intuition.

The logical positivists took analytic truths to be constitutive of a language. They are true by virtue of definitions or stipulations that determine the meanings of the terms involved in them. Hence, they issue in linguistic prescriptions about how to speak about things. Attendant on this doctrine was the so-called linguistic doctrine of necessity: all and only analytic truths are necessary. This doctrine, which, following Hume, excised all necessity from nature, had already played a key role in Wittgenstein's *Tractatus Logico-Philosophicus*.

If analytic statements acquire their meaning by linguistic stipulation, how do empirical (synthetic) statements acquire theirs? Being empiricists, the logical positivists took to heart Russell's *Principle of Acquaintance*: that all meaning (of descriptive concepts) should originate in experience. They elevated this principle into a criterion of meaningfulness, known as the verifiability criterion. Non-analytic statements are meaningful (cognitively significant) if and only if their truth can be verified in experience. In slogan form: meaning is the method of verification. The logical positivists utilized this criterion to show that statements of traditional metaphysics were meaningless, since their truth (or falsity) made no difference in experience. This consequentialist criterion was akin to that employed by the American pragmatists. But unlike the pragmatists, the logical positivists did not adopt a pragmatist theory of truth. (For a discussion of pragmatism, see "American philosophy in the twentieth century," Chapter 5.)

In the course of the 1930s, the concept of verifiability moved from a strict sense of provability on the basis of experience to the much more liberal sense of confirmability. The chief problem was that the strong criterion of cognitive significance failed to deliver the goods. Apart from metaphysical statements, many ordinary scientific assertions, those that express universal laws of nature, would be meaningless, precisely because they are not, strictly speaking, verifiable. In response to this, diehard advocates of verificationism, following Schlick, took the view that law-statements were inference-tickets: their only role was to serve as major premises in valid deductive arguments, whose minor premise and conclusion were singular observational statements. The immediate challenge to this view was that it was hard to see how a statement of the form "All *F*s are *G*s" could serve as a premise in a valid deductive argument without having a truth-value.

According to the logical positivists, if put together, the Hilbert approach to geometry and the Duhem–Poincaré hypothetico-deductive account of scientific theories, give a powerful and systematic way to represent scientific theories. The basic principles of the theory are taken to be the axioms. But the terms and predicates of the theory are stripped of their interpretation/meaning. Hence, the axiomatic system

itself is entirely formal. The advantage of the axiomatic approach is that it lays bare the logical structure of the theory. What is more, the axiomatic approach unambiguously identifies the content of the theory: it is the set of logical consequences of the axioms. However, treated as a formal system, the theory lacks any empirical content. In order for the theory to get some such content, its terms and predicates should be suitably interpreted. It was a central thought of the logical positivists that a scientific theory need not be completely interpreted to be meaningful and applicable. They claimed that it is enough that only *some*, the so-called “observational,” terms and predicates be interpreted. The rest of the terms and predicates of the theory, in particular those which, taken at face value, purport to refer to unobservable entities, were deemed “theoretical,” and were taken to be only partially interpreted by means of correspondence rules, that is *mixed* sentences that linked theoretical terms with observational ones. However, it was soon realized that the correspondence rules muddle the distinction between the analytic (meaning-related) and synthetic (fact-stating) part of a scientific theory, which was central in the thought of logical positivists. For, on the one hand, they specify (even partly) the meaning of theoretical terms and, on the other, they contribute to the factual content of the theory.

Hence, it became pressing for the logical positivists to restore a way in which a theory is divided into two distinct parts, one that fixes the meaning of its theoretical terms (and which cannot be falsified if the theory is in conflict with experience) and another part of the theory which specifies its empirical content (and which can be rejected in case of conflict with experience).

Objectivity and structure

It's been a popular thought, stressed by Quine (1969), that the logical positivists, and Rudolf Carnap (1891–1970) in particular, were epistemological foundation-ists. The *locus classicus* of this view is supposed to be found in Carnap's (1928) *The Logical Structure of the World*. (For further discussion of Carnap see “The birth of analytic philosophy,” Chapter 1 and “The development of analytic philosophy: Wittgenstein and after,” Chapter 2.) The received view, in outline, comes to this. Being an empiricist and, in particular a foundationalist, Carnap aimed to materialize Russell's “supreme maxim of philosophizing,” that is, the claim that “Whenever possible logical constructions are to be substituted for inferred entities.” By showing how all concepts that can be meaningfully spoken of can be constructed out of logico-mathematical concepts and concepts which refer to immediate experiences, Carnap's *Structure* would (a) secure a certain foundation for knowledge; (b) secure empiricism about meaning and concept formation; and (c) show how the content of talk about the world could be fully captured by an unproblematic-for-empiricists talk about the “given.” Although there is something in the received view, it seems that Carnap, far from being an unqualified foundationalist empiricist, had a broadly neo-Kantian axe to grind. Carnap's main aim was to show how the epistemological enterprise should be conducted for an empiricist if the objectivity of the world of science is to be secured. That is, it was *objectivity* and not foundations that Carnap was after.

Carnap's two kinds of objectivity

There are two ways to think of objectivity. The first is akin to inter-subjectivity, understood as the “common factor” point of view: the point of view common to all subjects (see Carnap 1928: §66). Looked at in that way, Carnap’s *Structure* aimed to show how the physical world could emerge from within his constructional system as the “common factor” of the individual subjective points of view (see §§148–9). Given the gap between the world of experience and the world of science, Carnap’s attempt was to specify the assumptions that should be in place (or the structure that needs to be added to the world of experience) in order to make possible the connection between the world of experience and the world of science (see §§125–9). The Kantian connection should be obvious here. Carnap stressed that, because his construction was based on the new Frege–Russell logic, he had no place for the Kantian synthetic a priori. Carnap’s first way to explicate objectivity came to grief: there is no way to draw the distinction between subjective and objective *within* Carnap’s constructional system, since the subjective is simply equated with whatever is *outside* the system of unified science.

But objectivity may well be conceived in more radical terms: as being totally subject-independent. This is what Carnap called “desubjectivized” science (see §16). Looked at in this light, Carnap’s *Structure* amounts to an endorsement of structuralism, where structure is understood as logical structure in the sense of Russell and Whitehead’s *Principia Mathematica*. Carnap tied content (the material) to subjective experience and made formal structure the locus of objectivity. His task then was to characterize all concepts that may legitimately figure in his system of unified science by means of “purely structural definite descriptions” (see §§11–15).

Russell's structuralism

Structuralist approaches to objectivity and knowledge had been advanced before Carnap. But it was Bertrand Russell (1872–1970) who made structuralism popular by claiming, in his *The Analysis of Matter* (1927), that the abstract character of modern physics, and of the knowledge of the world that this offers, could be reconciled with the fact that all evidence for its truth came from experience. To this end, he advanced a *structuralist* account of our knowledge of the world. According to this, only the structure, i.e. the totality of formal, logico-mathematical properties, of the external world can be known, while all of its intrinsic (qualitative) properties are inherently unknown. This logico-mathematical structure, he argued, can be legitimately *inferred* from the structure of the perceived phenomena. Russell explicitly noted that this view was an attempt to show how *some* knowledge of the Kantian *noumena* (namely, their structure) can be had. Relying on the causal theory of perception (which gave him the assumption that there are external physical objects which cause perceptions), and on other substantive principles (which connect differences in the percepts with differences in their causes, and conversely), he argued that there is an isomorphism between the structure of the percepts and the structure of their causes (stimuli). But Russell’s

structuralism was met with a fatal objection, due to the Cambridge mathematician M. H. A. Newman (1928); this was that the structuralist claim is *trivial* in the sense that it merely recapitulates Russell's assumption that there is a set of physical objects which cause perceptions. Given that this set of objects has the right cardinality, Newman showed that Russell's supposed further substantive conclusion, namely that of *this* set it is also known that it has certain structure W , is wholly insubstantial. Russell conceded defeat to Newman, but it might be ironic that the very same objection applies to Carnap's structuralism. Carnap insisted on building his account of structural objectivity out of a single relation of "recollection of similarity" among elementary experiences, which (relation) he aimed to characterize in purely structural terms. But there is no guarantee that such a purely structural description can characterize the relevant relation uniquely. Besides, that *there is* a relation with the required structure is a trivial claim, as Newman pointed out.

Structural convergence

Writing at the beginning of the twentieth century, Poincaré and Duhem too favored a structuralist account of scientific knowledge. Although Poincaré took the basic axioms of geometry and mechanics to be conventions, he thought that scientific hypotheses proper, even high-level ones such as Maxwell's laws, were empirical. Faced with the problem of discontinuity in theory-change (the fact that some basic scientific hypotheses and law-statements have been abandoned in the transition from one theory to another), he argued that there is, nonetheless, some substantial continuity at the level of the mathematical equations that represent empirical as well as theoretical relations. From this, he concluded that these retained mathematical equations – together with the retained empirical content – fully capture the objective content of scientific theories. By and large, he thought, the theoretical content of scientific theories is structural: if successful, a theory represents correctly the *structure* of the world.

As in many other cases, Poincaré's structuralism had a Kantian origin. He took it that science could never offer knowledge of things as they were in themselves. But he did add to this that their relations could nonetheless be revealed by structurally-convergent scientific theories. Duhem advanced this thought further by arguing that science aims at a natural classification of the phenomena, where a classification (that is, the organization of the phenomena within a mathematical system) is *natural* if the relations it establishes among the experimental phenomena "correspond to real relations among things" (1906: 26–7). He went on to draw a principled distinction between, as he put it, the explanatory and the representative part of a scientific theory. The former advances hypotheses about the underlying causes of the phenomena, whereas the latter comprises the empirical laws as well as the mathematical equations that are used to systematize and organize these laws. His novel thought was that in theory-change, the explanatory part of the theory gets abandoned while its representative (structural) part gets retained in the successor theory. Perhaps for the first time in the history of the philosophy of science, someone was trying to substantiate his claim

by looking in detail into the history of science, and especially the history of optical theories and of mechanics. One may wonder, however, whether the explanatory/representative distinction on which Duhem's argument rests is sound. One might say that explanatory hypotheses are representative in that they too are cast in mathematical form and, normally, entail predictions which can be tested. And conversely, mathematical equations that represent laws of nature are explanatory in that they can be used as premises for the derivation of other low-level laws.

Structural realism

Poincaré and Duhem found in structuralism an account of the objectivity (and the limits) of scientific knowledge. But they worked with a notion of *mathematical* structure as this is exemplified in the mathematical equations of theories. In this sense, their structuralism was different from Russell's and Carnap's, who worked with a sharper notion of *logical* structure (see Carnap's relevant remarks in 1928: §16). But their account of objectivity as structure-invariance reappeared in the 1960s under the guise of structural realism, in the writings of Grover Maxwell (1918–81) and, in the 1980s, of John Worrall and Elie Zahar. The twist they gave to structuralism was based on the idea of Ramsey-sentences.

This idea goes back to *Theories* (1929), by Frank Ramsey (1903–30) (see also "Philosophical logic," Chapter 8). Ramsey noted that the excess content that a theory has over and above its observational consequences is best seen when the theory is formulated as expressing an *existential judgment*: there are entities which satisfy the theory. This is captured by the Ramsey-sentence of the theory, which replaces the theoretical terms with variables and binds them with existential quantifiers. Ramsey thought that reference to theoretical entities does not require names; the existentially bound variables are enough. The Ramsey-sentence of a theory, which has exactly the same observational consequences as the original theory and the very same deductive structure, is truth-evaluable and committed, if true, to the existence of entities over and above observable ones.

Structural realists offered a structuralist reading of Ramsey-sentences: the empirical content of a theory is captured by its Ramsey-sentence and the excess content that the Ramsey-sentence has over its empirical consequences is purely *structural*. Given that the Ramsey-sentence captures the logico-mathematical form of the original theory, the structuralist thought is that, if true, the Ramsey-sentence also captures the structure of reality: the logico-mathematical form of an empirically adequate Ramsey-sentence mirrors the structure of reality.

It is a historical irony that the Newman-problem which plagued the Russell–Carnap structuralism is equally devastating against the Maxwell–Worrall–Zahar structuralism. For unless some non-structural restrictions are imposed on the kinds of things the existence of which the Ramsey-sentence asserts, that is, unless structuralism gives up on the claim that *only* structure can be known, it turns out that an empirically adequate Ramsey-sentence is bound to be true: truth collapses to empirical adequacy.

The analytic/synthetic distinction

In the 1930s, Carnap had to re-shape his agenda. Epistemology gave way to the “logic of science,” where the latter was taken to be a purely formal study of the language of science.

Logic of science

The thought, developed in Carnap’s *The Logical Syntax of Language* (1934), was that the development of a general theory of the logical syntax of the logico-mathematical language of science would provide a neutral framework in which: (a) scientific theories are cast and studied; (b) scientific concepts (e.g. explanation, confirmation, laws, etc.) are explicated; and (c) traditional metaphysical disputes are overcome. The whole project required a sharp analytic/synthetic distinction: philosophical statements would be analytic (about the language of science) whereas scientific ones would be synthetic (about the world). Carnap conceived of (artificially constructed) languages as systems of two kinds of rule: (a) formation rules, which specify how the sentences of the language are constructed and (b) transformation rules, which express inference rules. He distinguished between two kinds of transformation rule, namely (b₁) L-rules (the logical rules of syntax) and (b₂) P-rules (physical rules, such as the empirical laws of science). A logical language, then, consists of only L-rules, whereas a physical language (the language of a scientific theory) comprises P-rules too. The object of the general theory of logical syntax was such logical and physical languages. Carnap then equated analyticity with (suitably understood) provability *within* a language system: those sentences are analytic that are provable by means of the L-rules of the language. Hence, analyticity requires a sharp distinction between L-rules and P-rules and an equally sharp distinction between logical and descriptive symbols. But Carnap’s project came to grief. This was the result of many factors, but prominent among them were Tarski’s work on truth (which suggested that truth was an irreducibly semantic notion) and Kurt Gödel’s incompleteness theorem. Although Carnap was fully aware of Gödel’s limitative results, his own attempt to provide a neutral and minimal meta-theoretical framework (the framework of General Syntax) in which the concept of analyticity was defined fell prey to Gödel’s proof that there are mathematical (i.e. analytic) truths which are not provable within such a system.

The web of belief

Carnap’s enduring thought was that the analytic/synthetic distinction was necessary for any attempt to have a scientific understanding of the world, since the latter requires the presence of a linguistic and conceptual framework within which empirical facts are grafted. This view might plausibly be thought of as the last remnant of Kantianism in Carnap’s philosophy. This last remnant (that is, some notion of analytic a priori truths) came under heavy attack by W. V.

O. Quine (1908–2000). In “Truth by convention” (1936), Quine issued a deep challenge to the view that logic was a matter of convention. Logical truths, he argued, cannot be true by convention simply because if logic were to be derived from conventions, logic would be needed to *infer* logic from conventions. In “Two dogmas of empiricism” (1951), he went on to argue that the notion of analyticity is deeply problematic, since it requires a notion of cognitive synonymy and there is no independent criterion of cognitive synonymy. Quine’s master argument against the analytic/synthetic distinction rested on the view that “analytic” was taken to mean “unrevisable.” If analytic statements have no empirical content, experience cannot possibly have any bearing on their truth-value. So analytic statements can suffer no truth-value revision. But, Quine argued, nothing (not even logical truths) is unrevisable. Hence, there cannot be any analytic truths. Here Quine took a leaf from Duhem’s (and indeed Carnap’s) book. Confirmation and refutation are holistic: they accrue to systems (theories) as a whole and not to their constituent statements, taken individually. If a theory is confirmed, then everything it asserts is confirmed; and conversely, if a theory is refuted, then any part of it can be revised (abandoned) in order for harmony with experience to be restored. Accordingly, anything (i.e. any statement, be it logical or mathematical or empirical) can be confirmed (hence, everything has empirical content). And anything (i.e. any statement, be it logical or mathematical or empirical) can be refuted (hence, everything can be such that experience can refute it). *Ergo*, there are no analytic (= unrevisable) statements.

The image of science that emerges has no place for truths of a special status: all truths are on a par. This leads to a blurring of the distinction between the factual and the conventional. What matters, for Quine, is that a theory acquires its empirical content as a whole, by issuing in observational statements and by being confronted with experience. The whole idea, however, rests on the claim that observational statements are privileged. They are not, to be sure, privileged in the way foundationalists thought they were: they are not the certain and incorrigible foundations of all knowledge. But they are privileged in the sense that all evidence for or against a theory comes from them and also in the sense that they can command intersubjective agreement. Quine did not deny that some principles (those of logic, or mathematics, or some basic physical principles) do *seem* to have some special status. He argued, however, that this is not enough to confer on them the status of a priori truths. Rather, they are very central to our web of belief. When our theory of the world clashes with experience, they are the last to be abandoned. But they *can* be abandoned if their rejection results in a simpler web of belief. In contemplating what kind of changes we should make, Quine argued, we should be guided by principles such as simplicity, or the principle of minimal mutilation. Yet the status of these methodological principles remained unclear. Why for instance couldn’t *they* be seen as a priori principles that govern belief-revision? And if they are not, what is the argument which shows that they are also empirical principles? Surely, one cannot read them off experience in any straightforward way.

Carnap meets Quine

The cogency of Quine's attack on the a priori rests on the cogency of the following equation: a priori = unrevisable. We have already seen a whole strand in post-Kantian thinking that denied this equation, while holding on to the view that some principles are constitutive of experience. It might not be surprising then that Carnap was not particularly moved by Quine's criticism. For he *too* denied the foregoing equation (see Carnap 1934: 318). He took the analytic/synthetic distinction to be *internal* to a language and claimed that analyticity (a priority) is not invariant under language-change. In radical theory-change, the analytic/synthetic distinction has to be re-drawn within the successor theory. "Being held true, come what may" holds for both the analytic and the synthetic statements, when minor changes are required. Similarly, both analytic statements and basic synthetic postulates may be revised, when radical changes are in order.

Quine did have a point. For Carnap, analytic statements are such that (a) it is rational to accept them within a linguistic framework; (b) it is rational to reject them, when the framework changes; and (c) there is some extra characteristic which all and only analytic statements share, in distinction to synthetic ones. Even if Quine's criticisms are impotent *vis-à-vis* (a) and (b), they are quite powerful against (c). The point was simply that the dual role of correspondence rules (and the concomitant Hilbert-style implicit definition of theoretical terms) made the drawing of this distinction impossible, even *within* a theory. It took a great deal of Carnap's time and energy to find a cogent explanation of (c). In the end, he had to re-invent the Ramsey-sentences to find a plausible way to draw the line between the analytic and the synthetic (see Psillos 2000).

Confirmation and induction

The logical empiricists thought that, by devising a formal and algorithmic account of confirmation, they could kill two birds with one stone. They could bypass the problem of induction and solve the problem of the objectivity of scientific method. The former, it was thought, could be bypassed by transforming it into a logico-syntactic theory of confirmation, according to which, as it rolls in, the evidence increases the probability of a hypothesis. The latter, it was thought, could be solved by showing how the scientific method (and induction, in particular) was based on purely logical rules.

Confirmation: syntax and semantics

Hempel's (1945) logico-syntactic theory of confirmation was based on the idea that given any hypothesis of the form "All As are Bs," and given any positive instance of this hypothesis (that is an object *a* which is both A and B), this positive instance confirms the hypothesis that "All As are Bs." This view came under attack when Hempel himself noted that it leads to paradox, known as the Ravens paradox.

However, it was Nelson Goodman (1906–98) who dealt a fatal blow to Hempel's logico-syntactic approach by introducing a new predicate, "grue," which is defined as

follows: observed before 2010 and found green, or not observed before 2010 and be blue. (For further discussion of Goodman, see “American philosophy in the twentieth century,” Chapter 5.) Clearly, all *observed* emeralds are green. But they are also “grue.” Why then should we claim that the observation of a green emerald confirms the hypothesis that *All emeralds are green* instead of *All emeralds are grue*? These two generalizations are consistent with all current evidence, but they disagree on their predictions about emeralds that are observed after 2010. Hempel’s logico-syntactic theory is blind to this difference: the two competing hypotheses have exactly the same syntactic form. Goodman (1954) argued that only the first statement (“All emeralds are green”) is capable of expressing a law of nature because only this is confirmed by the observation of green emeralds. He disqualified the generalization “All emeralds are grue” on the grounds that the predicate “is grue,” unlike the predicate “is green,” does not pick out a natural kind. But the very idea of having suitable (natural-kind) predicates featuring in confirmable generalizations speaks against a purely syntactic account of confirmation.

Inductive logic

Carnap, following John Maynard Keynes (1883–1946), took pains to advance an inductive *logic* (which, in certain respects, would mimic the content-insensitive structure of deductive logic). In his *Logical Foundations of Probability* (1950) he devised a quantitative confirmation function that expressed the *degree of partial entailment* of a hypothesis by the observational evidence. The evidence (cast in observational statements) is supposed to confirm a hypothesis to the extent in which it partially entails this hypothesis, where partial entailment was a formal relation between the range of the evidence and the range of the hypothesis.

In his system of inductive logic, Carnap claimed, true sentences expressing relations of partial entailment are analytic. And if inductive logic is analytic, it is also a priori. He went as far as to claim that the contentious principle of uniformity of nature, a principle that Hume criticized not as false but as of illegitimate use in any attempt to justify induction, was also analytic within his system of inductive logic: it is statement of logical probability asserting that “on the basis of the available evidence it is very *probable* that the degree of uniformity of the world is high” (Carnap 1950: 180–1). Carnap was certainly right in noting that, expressed as above, the principle of uniformity of nature could neither be proved nor refuted by experience. But for this principle to be acceptable, we need to assume another principle, namely what Keynes called “the Principle of Limited Variety” (1921: 287). Suppose that although *C* has been invariably associated with *E* in the past, there is an unlimited variety of properties E_1, \dots, E_n such that it is possible that future occurrences of *C* will be accompanied by any of the E_i ’s, instead of *E*. Then, and if we let *n* (the variety index) tend to infinity, we cannot even start to say how likely it is that *E* will occur, given *C* and the past association of *C*s with *E*s. The Principle of Limited Variety excludes the possibility just envisaged. It is necessary for making inductive inferences of any sort. This, no doubt, is a substantive, synthetic, principle *about* the world. Carnap cannot

simply relegate it to yet another analytically true principle. Indeed, as was pointed out by Keynes long before Carnap, the very possibility of inductive inference requires that some hypotheses are given non-zero prior probability, for otherwise fresh evidence cannot raise their probability.

If “inductive logic” had to rely on substantive synthetic principles, it was no longer *logic*. Carnap’s attempt to use the quasi-logical Principle of Indifference, which dictated that all equally possible outcomes should be given equal prior probabilities of happening, led to inconsistent results. In the end, when Carnap devised the continuum of inductive methods, he drew the conclusion that there can be a variety of actual inductive methods whose results and effectiveness vary in accordance to how one picks out the value of a certain parameter, where this parameter depends on formal features of the language used. But obviously, there is no a priori reason to select a particular value of the relevant parameter; hence there is no explication of inductive inference in a unique way. Carnap suggested that it is left to the scientists to choose among different inductive methods, in view of their specific purposes. Where an a priori justification of induction was sought, the end product was based on a pragmatic decision.

Naturalism: the break with neo-Kantianism

The emergence of naturalism was a real turning point in the philosophy of science. It amounted to an ultimate break with neo-Kantianism in all of its forms (see also “Naturalism,” Chapter 6).

Against a priorism

The challenge to the very possibility of a priori knowledge was a central input in the *naturalist turn* in the philosophy of science in the 1960s. The old a prioristic way of doing philosophy of science was nicely summarized by Reichenbach: “We see such a way in the application of the *method of logical analysis* to epistemology. The results discovered by the positive sciences in continuous contact with experience presuppose principles the detection of which by means of logical analysis is the task of philosophy” (1921: 74).

Quine’s “Epistemology naturalized” (1969) claimed that, once the search for secure foundations of knowledge is shown futile, philosophy loses its presumed status as the privileged framework (equipped with a privileged source of knowledge: a priori reflection and logical analysis) aiming to validate science. Philosophy becomes continuous with the sciences in that there is no privileged philosophical method, distinct from the scientific method, and that the findings of empirical sciences are central to understanding philosophical issues and disputes. Quine went as far as to suggest a replacement thesis: epistemology, as traditionally understood, should give way to psychology. The aim of this psychologized epistemology should be the understanding of the link between observations (the meager input of scientific method) and theory (its torrential output), an issue that, strictly speaking, belongs to psychology.

Why, then, call it epistemology? Because, Quine says, we are still concerned with a central question of traditional epistemology: how do theories relate to *evidence*? Quine made capital on the vivid metaphor of Neurath's boat: "Neurath has linked science to a boat which, if we are to rebuild it, we must rebuild plank by plank while staying afloat in it" (1960: 3). Accordingly, philosophy has no special status; nor are there parts of our conceptual scheme (the findings of science in particular) that cannot be relied upon when revision elsewhere in it are necessary.

Against anti-psychologism

Naturalism's other input was the rejection of anti-psychologism that had dominated philosophy of science in the first half of the twentieth century. Broadly speaking, anti-psychologism is the view that the content of the key concepts of epistemology and philosophy of science (knowledge, evidence, justification, method, etc.) should be analyzed independently of the psychological processes that implement them to cognitive subjects. More specifically, anti-psychologism holds that logical questions should be sharply distinguished from psychological questions. This distinction led to a sharp separation between the *context of discovery* and the *context of justification* (which was introduced by Reichenbach). In *The Logic of Scientific Discovery* (1959), Karl Popper (1902–94) argued that, being psychological, the process of coming up with theories is not subject to rigorous logical analysis, nor is it interesting philosophically (on Popper, see also "The development of analytic philosophy: Wittgenstein and after," Chapter 2). He lumped it under the rubric of devising conjectures and argued that the important philosophical issue was how hypotheses are tested after they have been suggested. He took it that the process of testing is a formal process of trying to find an *inconsistency* between the predictions of a theory and the relevant observations.

Popper's anti-psychologism was based on his stance on Hume's problem of induction. But whereas Hume, who was a proto-naturalist, thought that induction was an inevitable natural-psychological process of forming beliefs, Popper went as far as to argue that induction was a myth: real science (as well as ordinary thinking) doesn't (and shouldn't) employ induction. In its place he tried to put the idea that scientific method needs no more than deduction: even if scientific theories cannot be proved (or confirmed) on the basis of experience, they can still be falsified by the evidence. This thought, however, is simplistic. For one thing, as Duhem and Quine have pointed out, no theory is, strictly speaking, refutable. Typically, in mature sciences, the theory being tested cannot generate empirical predictions without the use of several auxiliary assumptions. Hence, if the prediction is not fulfilled, the only thing we can *logically* infer is that either the auxiliaries are, or the theory is, false. That is, logic alone cannot distribute blame to the premises of an argument with a false conclusion. For another, the key Popperian thought, namely that deductive logic is the only logic that we have or need and that the only valid arguments are deductively valid arguments, is deeply problematic. This last view, known as deductivism, is flawed both as a descriptive and as a normative thesis.

The historical turn

Although the logical positivists favored inductivism and Popper deductivism, an assumption shared by both was that scientific method should be justified a priori. Popper, to be sure, was more inclined to the view that scientific methods are conventions, but he did subscribe to the positivist idea that a theory of scientific method is a *rational reconstruction* of the scientific game: a reconstruction that presents science as a rule-governed activity. These rules are meant to constitute the very essence of the scientific game very much like the rules of chess are constitutive of a game of chess. This idea gives rise to an important question: if there are many competing rational reconstructions of the scientific game, which one is to be preferred? As was fully acknowledged by the history-oriented methodologies of science that succeeded Popper and the positivists, this question could be answered only if the very game of rational reconstruction gave way to accounts informed by the actual history and practice of science.

The historical turn of the 1960s was a central component of the naturalist turn. Thomas S. Kuhn's (1922–96) theory of science should be seen as the outcome of two inputs: (a) a reflection on the actual scientific practice as well as the actual historical development and succession of scientific theories, and (b) a reaction to what was perceived to be the dominant logical empiricist and Popperian images of scientific growth: a progressive and cumulative process governed by specific rules on how evidence relates to theory. A coarse-grained explication of Kuhn's *The Structure of Scientific Revolutions* (1962) goes as follows. The emergence of a scientific discipline is characterized by the adoption by a community of a paradigm, that is, as of a shared body of theories, models, exemplars, values, and methods. A long period of "normal science" emerges, in which scientists attempt to apply, develop, and explore the consequences of the paradigm. This activity is akin to puzzle-solving, in that scientists follow the rules (or, the concrete exemplars) laid out by the paradigm in order (a) to determine which problems are solvable and (b) to solve them. This rule-bound (or, better, exemplar-bound) activity that characterizes normal science goes on until an anomaly appears. An anomaly is a problem that falls under the scope of the paradigm, is supposed to be solvable, but persistently resists solution. The emergence of anomalies signifies a serious decline in the puzzle-solving efficacy of the paradigm. The community enters a stage of crisis that is ultimately resolved by a revolutionary transition (paradigm shift) from the old paradigm to a new one. The new paradigm employs a different conceptual framework, and sets new problems as well as rules (exemplars) for their solutions. A new period of normal science emerges. Crucially, for Kuhn the change of paradigm is not rule-governed. It's nothing to do with degrees of confirmation or conclusive refutations. Nor does it amount to a slow transition from one paradigm to the other. Rather, it's an abrupt change in which the new paradigm completely replaces the old one. The new paradigm gets accepted not because of superior arguments or massive evidence in its favor, but rather because of its proponents' powers of rhetoric and persuasion.

The Kuhnian approach had a mixed reception. Many philosophers of science thought that it opened up the way to relativist and irrationalist views of science.

Imre Lakatos's (1922–74) lasting contribution to this debate was his attempt to combine Popper's and Kuhn's images of science in one model of theory-change that preserves progress and rationality while avoiding Popper's naive falsificationism *and* doing justice to the actual history of radical theory-change in science. To this end, Lakatos (1970) developed the "Methodology of Scientific Research Programmes." A research program is a sequence of theories and consists of a hard core, the negative heuristic, and the positive heuristic. The *hard core* comprises all those theoretical hypotheses that any theory which belongs to the research program must share. The advocates of the research program hold these hypotheses immune to revision. This methodological decision to protect the hard core constitutes the negative heuristic. The process of articulating directives about how the research program will be developed, either in the face of anomalies or in an attempt to cover new phenomena, constitutes the positive heuristic. It creates a *protective belt* around the hard core, which absorbs all potential blows from anomalies. A research program is progressive as long as it issues in novel predictions, some of which are corroborated. It is degenerating when it offers only *post hoc* accommodations of facts, either discovered by chance or predicted by a rival research program. Progress in science occurs when a progressive research program supersedes a degenerating one. But although this Lakatosian program is highly plausible and suggestive, it suffers from a severe drawback: as an account of scientific methodology, it is retroactive. It provides no way to tell which of two *currently* competing research programs is and will continue to be progressive. For even if one of them seems stagnated, it may stage an impressive comeback in the future.

Methodological naturalism

The real bite of the naturalist turn is that it made available a totally different view of how scientific methods are justified. The core of methodological naturalism is the view that methodology is an empirical discipline and that, as such, it is part and parcel of natural science. This approach suggests the following. (1) Normative claims are *instrumental*: methodological rules link up aims with methods which will bring them about, and recommend what action is more likely to achieve one's favored aim. (2) The soundness of methodological rules depends on whether they lead to successful action, and their justification is a function of their effectiveness in bringing about their aims. Where traditional approaches to method were seen as issuing in *categorical imperatives*, the naturalist suggestion was that methodological claims should be seen as issuing in *hypothetical imperatives*. These imperatives were seen as resting on (supervening upon) factual claims which capture correlations, causal links or statistical laws between "doing X" and "achieving Y." Hence, these (hypothetical) imperatives depend on contingent features of the world. This view implied a big break with the traditional a prioristic conception of method. After all, we want our methods to be effective in this world, i.e. to guide scientists to correct decisions and correct strategies for extracting information from nature. In this sense, the methods scientists adopt must be amenable to substantive information about the actual world.

Methodological naturalism carried a number of problems in its train. The first is the problem of circularity: were methods to be justified empirically, wouldn't some methods (the very same ones?) be presupposed in coming to legitimately accept the empirical findings that are relevant to their vindication? There are many things one could say in reply to this charge. But the short answer is that we are all passengers on *HMS Neurath*. There is no dry-dock in which we can place our conceptual scheme as a whole and examine it bit by bit. Rather, we are engaged in a process of mutual adjustment of its pieces while keeping it afloat: we fix our methods in light of empirical facts and ascertain facts while keeping our methods fixed. If the aim of philosophy of science is no longer to validate science and its method, the ensuing threat of circularity is moot.

The second problem with methodological naturalism is that it seems to lead to epistemic relativism. This concern forces naturalism to adopt an axiology, that is, a general theory of the constraints that govern rational choice of aims and goals. More specifically, naturalism has to accept truth as the *basic cognitive virtue*. This move blunts the threat of relativism. In fact, naturalism should take a leaf from the book of reliabilism in epistemology. Its central point is that justification is issued by reliable processes and methods, where a process or method is reliable if it tends to generate true beliefs. Accordingly, methodological and reasoning strategies should be evaluated by their success in producing and maintaining true beliefs, or tending to do so, where judging this success, or establishing the tendency, is open to empirical findings and investigation. Methodological naturalism should appeal to reliabilism in order to meet the third challenge it faces, namely that it leaves no room for normative judgment. Reliabilism supplements methodological naturalism with a normative meta-perspective of truth-linked judgments.

Semantic holism and the theory/observation distinction

A new approach to meaning – semantic holism – started to become popular in the 1960s. Instead of trying to specify the meaning of each and every term in isolation from the others, semantic holism claimed that terms (and, in particular, so-called “theoretical terms”) get their meaning from the theories and network of nomological statements in which they are embedded. Although the logical empiricists subscribed to semantic atomism, they advocated confirmational holism, that is, the view that theories are confirmed as wholes. Already in *Syntax*, Carnap stressed: “Thus *the test applies, at bottom, not to a single hypothesis but to the whole system of physics as a system of hypotheses*” (Duhem, Poincaré) (1934: 318). But confirmational holism, conjoined with the denial of the analytic/synthetic distinction and the thought that confirmable theories are meaningful, leads to semantic holism. Later empiricists, notably Carl Hempel (1907–97), accepted semantic holism. But Carnap (1956) resisted this conclusion to very end.

Theory-ladenness and incommensurability

An increasingly popular view in the 1960s was that there was no way to draw a sharp line between theoretical statements and observational ones. In its *strong* version, it

claimed that, strictly speaking, there can be *no* observational terms at all. This claim, inspired by the view that all observation is theory-laden, led to the thought that there cannot *possibly* be a theory-neutral observation language. This view was supported by Norwood Russell Hanson (1924–67), Kuhn, and Paul Feyerabend (1924–94). But it was advanced at the beginning of the century by Duhem (1906). In the 1960s, when this thesis resurfaced, it drew on a mass of empirical evidence coming from psychology to the effect that perceptual experience is theoretically interpreted. Hanson, Kuhn, and Feyerabend pushed the theory-ladenness-of-observation thesis to its extremes, by arguing that each theory (or paradigm) determines the meaning of *all* terms that occur in it and that there is *no* neutral language which can be used to assess different theories (or paradigms). If the meanings of terms are determined by the theory as a whole, it could be claimed that every time the theory changes, the meanings of *all* terms change, too. We have then a thesis of radical meaning variance. If, on top of that, it is accepted that meaning determines reference (as the traditional descriptive theories of meaning have it), an even more radical thesis follows, namely, reference variance.

Kuhn argued that a paradigm defines a world. Accordingly, there is a sense in which when a new paradigm is adopted, the world changes: Priestley, with his phlogiston-based paradigm lived in a different world from Lavoisier, with his oxygen-based theory. Remarks such as this have led to a lot of interpretative claims. But, I think, the best way to understand Kuhn's general philosophical perspective is to say that Kuhn's underlying philosophy is relativized neo-Kantianism. It's neo-Kantianism because it implies a distinction between the world-in-itself, which is epistemically inaccessible to cognizers, and the phenomenal world, which is constituted by the cognizers' concepts and categories. But Kuhn's neo-Kantianism is *relativized* because he thought there was a plurality of phenomenal worlds, each being dependent on, or constituted by, some community's paradigm. The paradigm imposes, so to speak, a structure on the world of appearances: it carves up this world into "natural kinds." This is how a phenomenal world is "created." But different paradigms carve up the world of appearances in different networks of natural kinds. Incommensurability then follows (even if it is local rather than global) since it is claimed that there are no ways to match up the natural-kind structure of one paradigm with that of another.

Social constructivism

Kuhn's views inspired social constructivists. Social constructivism is an agglomeration of views with varying degrees of radicalness and plausibility. Here is a sketchy list of them. The truth of a belief has nothing to do with its acceptability; beliefs are determined by social, political, and ideological forces, which constitute their causes. Scientific facts are constructed out of social interactions and negotiations. Scientific objects are created in the laboratory. The acceptability of scientific theories is largely, if not solely, a matter of social negotiation and a function of the prevailing social and political values. Science is only one of any number of possible "discourses," no one of which is fundamentally truer than any other. What unites this cluster of

views are (vague) slogans such as “scientific truth is a matter of social authority” or “nature plays (little or) no role” in how science works. It might be useful to draw a distinction between a weak form of social constructivism and a strong one. The weak view has it that some categories (or entities) are socially constructed: they exist because we brought them into existence and persist as long as we make them do so. Money, the Red Cross, and football games are cases in point. But this view, though not problem-free, is almost harmless. On the strong view, all reality (including the physical world) is socially constructed: it is a mere projection of our socially inculcated conceptualizations. I am not sure anyone can really *believe* this. But there is a rather strong argument against the (more sensible) Kuhnian relativized neo-Kantianism. Roughly put, the argument is that the world enters the picture as that which offers *resistance* to our attempts to conceptualize it. Unless we think there is a world with an objective natural-kind structure, we cannot explain the appearance, persistence, and stubbornness of anomalies in our theories.

Causal theories of reference

The advent of the causal theories of reference in the 1970s challenged the alleged inevitability of incommensurability. Hilary Putnam (1973) extended Saul Kripke’s (1972) causal theory of proper names to the reference of natural-kind terms and to physical-magnitude terms. (For further discussion of Kripke, see “American philosophy in the twentieth century,” Chapter 5 and “Philosophy of language,” Chapter 9.) The thrust of the causal theory is that the relation between a word and an object is *direct* – a direct causal link – unmediated by a concept. Generally, when confronted with some observable phenomena, it is reasonable to assume that there is a physical magnitude, or entity, which causes them. Then we dub this magnitude with a term and associate this magnitude with the production of these phenomena. The descriptions that will, typically, surround the introduction of the term might be incomplete, misguided, or even mistaken. Yet, on the causal theory of reference, one has nonetheless introduced *existentially* a referent: an entity causally responsible for certain effects to which the introduced term refers.

It is easily seen that the causal theory disposes of semantic incommensurability. Besides, it lends credence to the claim that even though past scientists had partially or fully incorrect beliefs about the properties of a causal agent, their investigations were continuous with the investigations of subsequent scientists, since they were aiming to identify the nature of the same causal agent. Yet there is a sense in which the causal theory of reference makes reference-continuity in theory-change all too easy. If the reference of theoretical terms is fixed purely existentially, then insofar as *there is* a causal agent behind the relevant phenomena, the term is bound to end up referring to it. Hence, there can be no referential failure, even in cases where it is counterintuitive to expect successful reference. This, and other problems, have led many theorists to combine elements of both the causal and the descriptive theories of reference into a causal-descriptivist account. On this mixed account, the burden of reference of theoretical terms lies with *some* descriptions which specify the kind-constitutive properties in virtue of which the referent, if it exists, plays its causal role.

The role and structure of scientific theories

In spite of the empiricists' best efforts, the meaning of a theoretical term cannot be exhausted by a set of empirical applications. It should also include the *entity* to which this term putatively refers. Consequently, terms such as "electron" or "magnetic flux" are no less meaningful than terms such as "table" or "is red." The latter refer to observable entities, whereas the former can be seen as putatively referring to unobservable entities.

Semantic realism and instrumentalism

This breath of fresh air came to be known as "semantic realism." It was championed in the 1950s by Herbert Feigl (1902–88). He claimed that the empiricist program had been a hostage to verificationism for too long. Verificationism runs together two separate issues: the evidential basis for the truth of the assertion and the semantic relation of designation (i.e. reference). It thereby conflates the issue of what constitutes evidence for the truth of an assertion with the issue of what make this assertion true. If evidence-conditions and truth-conditions are separated, verificationism loses its bite. An observational statement as well as a theoretical statement are true if and only if their truth-conditions obtain. Accordingly, theoretical terms as well as (and no less than) observational terms have putative factual reference. If theoretical statements cannot be given truth-conditions in an ontology that dispenses with theoretical entities, then a full and just explication of scientific theories simply requires commitment to irreducible unobservable entities, no less than it requires commitment to observable entities.

A strong pressure on semantic realism came from the application to philosophy of science of Craig's theorem: for any scientific theory T , T is replaceable by another (axiomatizable) theory $\text{Craig}(T)$, consisting of all and only the theorems of T which are formulated in terms of the observational vocabulary (Craig 1956). The gist of Craig's theorem is that a theory is a conservative extension of the deductive systematization of its observational consequences. This point was readily seized upon by instrumentalists. Instrumentalism claims that theories should be seen as (useful) instruments for the organization, classification, and prediction of observable phenomena. So, the "cash value" of scientific theories is fully captured by what theories say about the observable world. Instrumentalists made use of Craig's theorem in order to argue that theoretical commitments in science were dispensable: theoretical terms could be eliminated *en bloc*, without loss in the deductive connections between the observable consequences of the theory.

This debate led Hempel (1958) to formulate what he called "the theoretician's dilemma." If the theoretical terms and the theoretical principles of a theory do not serve their purpose of a deductive systematization of the empirical consequences of a theory, they are dispensable. But, even if they do serve their purpose, given Craig's theorem, they can be dispensed with. Theoretical terms and principles of a theory either serve their purpose or they do not. Hence, the theoretical terms and principles

of any theory are dispensable. But is this dilemma compelling? As Hempel himself stressed, it is implausible to think of theories as establishing solely a deductive systematization of observable phenomena. Theories also offer inductive systematizations in the sense that they are used to establish inductive connections among them: they function as premises in *inductive* arguments whose other premises concern observable phenomena and whose conclusions refer to observable phenomena. But then, the Craig-theorem-based instrumentalist argument is blocked: theories, seen as aiming to establish inductive connections among observables, are *indispensable*.

The semantic view of theories

In the 1970s, the dominant formal-syntactic view of theories came under heavy attack by a number of philosophers of science. What emerged was a cluster of views known as the *semantic view of theories*. The core of this view was that theories represent via models and hence that the characterization of theories, as well as the understanding of how they represent the world, should rely on the notion of model. Where the logical empiricists favored a formal axiomatization of theories in terms of first-order logic, thinking that models can play only an illustrative role, the advocates of the semantic view opted for a looser account of theories, based on mathematics rather than meta-mathematics. As an answer to the question "What is a scientific theory?," the semantic view claims that a scientific theory should be thought of as something *extra-linguistic*: a certain structure (or class of structures) that admits of different linguistic garbs. This thought goes back to Patrick Suppes and was further pursued by Fred Suppe and Bas van Fraassen. Accordingly, to present a theory amounts to presenting a family of models. A key argument for the semantic view was that its own characterization of theories tallies better with the actual scientific conception of theories.

An immediate challenge to this view was that it was unclear how theories can represent anything empirical and hence how they can have empirical content. This challenge has been met in several ways, but two ways became prominent. The first way was that the representational relation was, ultimately, some mathematical *morphism*. The theory represents the world by having one of its models *isomorphic* to the world or by having the empirical phenomena *embedded* in a model of the theory. However, mathematical morphisms preserve only structure and hence it is not clear how a theory acquires any specific empirical content, and in particular how it can be judged as true. The second way, which was advanced to meet this last challenge, was that theories are mixed entities: they consist of mathematical models plus theoretical hypotheses. The latter are linguistic constructions which claim that a certain model of the theory represents (say by being similar to) a certain worldly system. On this view, advanced by Ronald Giere (1988), theoretical hypotheses provide the *link* between the model and the world. But in light of this mixed account of theories, it may well be that a proper understanding of theories requires the applications of both linguistic and extra-linguistic resources and that, in the end, there is not a single medium (the mathematical model) via which theories represent the world: theories may well be complexes of representational media.

Empiricism versus realism

Empiricists thought that a theory-free observational language was necessary for the existence of a neutral ground on which competing theories could be confirmed and compared. However, at least some of them also thought that such a language could capture “the given” in experience, which acted as the certain foundation of all knowledge. Yet these two functions do not come to the same thing. Even if observation is theory-laden, it can be laden by theories (or beliefs) that are accepted by all sides of a theoretical dispute. Observations can be fallible and corrigible and yet they can confirm theories. That observations can be the epistemic foundation of all knowledge, in the sense of being an autonomous stratum of knowledge, is a dogma of empiricism that was refuted by Wilfrid Sellars (1912–89) in his famous attack on the “myth of the given” (1956). (For further discussion of Sellars, see “American philosophy in the twentieth century,” Chapter 5.)

Realism and explanation

Sellars took a further step. He claimed that there was another popular myth that empiricism capitalized on: *the myth of the levels* (1963). This myth rested on the following image. There is the bottom level of observable entities. Then there is the intermediate level of the observational framework, which consists of empirical generalizations about observable entities. And finally, there is yet another (higher) level: the theoretical framework of scientific theories which posits unobservable entities and laws about them. It is part of this image that while the observational framework is explanatory of observable entities, the theoretical framework explains the inductively established generalizations of the observational framework. But then, Sellars says, the empiricist will rightly protest that the higher level is dispensable. For all the explanatory work vis-à-vis the bottom level is done by the observational framework and its inductive generalizations. Why then posit a higher level in the first place? Sellars’s reply was that the unobservables posited by a theory explain *directly* why (the individual) observable entities behave the way they do and obey the empirical laws they do (to the extent that they do obey such laws). He, therefore, offered an indispensability argument for the existence of unobservable entities: they are indispensable elements of scientific explanation.

By the 1960s, the tide had started to move the scientific realists’ way. Jack (J. J. C.) Smart (1963) put the scientific realism issue in proper perspective by arguing that it rests on an abductive argument, *aka inference to the best explanation*. Smart attacked the instrumentalists’ position by saying that they must believe in cosmic coincidence. On the instrumentalist view of theories, a vast number of ontologically disconnected observable phenomena are “connected” only by virtue of a purely instrumental theory: they just *happen* to be, and just *happen* to be related to one another in the way suggested by the theory. Scientific realism, on the other hand, leaves no space for a cosmic-scale coincidence: it is *because* theories are true and *because* the unobservable entities they posit exist that the phenomena are, and are related to one another, the way they are.

Smart's point was that scientific realism (and its concomitant view of science) should be accepted because it offers the best explanation of why the observable phenomena are as they are predicted by scientific theories. Hilary Putnam (1975) and Richard Boyd (1973) argued that inference to the best explanation is the very method scientists use to form and justify their beliefs in unobservable entities and that realism should be seen as an overarching empirical hypothesis which gets support from the fact that it offers the best explanation of the success of science. The Putnam–Boyd argument came to be known as “the no-miracles argument” since, in Putnam’s words, “The positive argument for realism is that it is the only philosophy that does not make the success of science a miracle” (1975: 73).

Explanation and metaphysics

Oddly enough, a similar debate had been conducted at the beginning of the twentieth century. But then it was mostly shaped by an emergent scientific theory, atomism, and some philosophical reactions to it. Atomism posited the existence of unobservable entities – the atoms – to account for a host of observable phenomena (from chemical bonding to Brownian motion). Although many scientists adopted atomism right away, there was some strong resistance to it by some eminent scientists. Ernst Mach resisted atomism on the basis of the concept-empiricist claim that the very concept of “atom” was problematic because it was radically different from ordinary empirical concepts. Poincaré was initially very skeptical of the atomic hypothesis, but near the end of his life he came to accept it as a result of the extremely accurate calculation of Avogadro’s number by the French physicist Jean Perrin (1870–1942), who employed thirteen distinct ways to specify the precise value of this number. This spectacular development suggested to Poincaré an irresistible argument in favor of the reality of atoms: we can calculate how many they are, therefore they exist. Yet resistance still persisted and was best exemplified in the writings of Duhem. He built his resistance to atoms on a sharp distinction between science and metaphysics. He claimed that explanation by postulation – that is, explanation in terms of unobservable entities and mechanisms – belonged to metaphysics and not to science.

Duhem’s theory of science rested on a very restricted understanding of scientific method, which can be captured by the slogan: *scientific method = experience + logic*. On this view, whatever cannot be proved from experience with the help of logic is irredeemably suspect. The irony is that Duhem himself offered some of the best arguments against an instrumentalist conception of theories. The central one comes from the possibility of novel predictions: if a theory were just a “rack filled with tools,” it would be hard to understand how it can be “a prophet for us” (1906: 27). Duhem’s point was that the fact that some theories generate *novel* predictions could not be accounted for on a purely instrumentalist understanding of scientific theories. However, he thought that explanatory arguments were outside the scope of scientific method. Yet this thought is deeply problematic. Scientists employ explanatory arguments in order to accept theories. So, contrary to Duhem, it is *part and parcel* of science and its method to rely on explanatory considerations in order to form and defend rational

belief. In fact, this point was forcefully made by the American pragmatist Charles Saunders Peirce (1839–1914) in his spirited defense of abduction.

Duhem's understanding of metaphysics as bound up with any attempt to offer explanation by postulation became part of the empiricist dogma for many decades. It led many empiricists to reject scientific realism. However, the key question is what kinds of methods are compatible with empiricism. Even if we grant, as we should, that all knowledge starts with experience, its boundaries depend on the credentials of the methods employed. It is perfectly compatible with empiricism, as Reichenbach and others noted, to accept ampliative (inductive) methods and to accept the existence of unobservable entities on their basis. So there is no incompatibility between being an empiricist (which, after all, is an epistemological stance) and being a scientific realist (which, after all, is a metaphysical stance).

Constructive empiricism

The rivalry of realism and empiricism was fostered by van Fraassen's influential doctrine of constructive empiricism in the 1980s. This is a view about science according to which science aims at empirically adequate theories, and acceptance of scientific theories involves belief only in their empirical adequacy (though acceptance involves *more* than belief, namely commitment to a theory). Van Fraassen (1941–) took realism to be, by and large, an *axiological* thesis: the aim of science is true theories. He supplemented it with a *doxastic* thesis: acceptance of theories implies belief in their truth. Seen that way, realism and constructive empiricism are rivals. But of course, a lot depends on whether an empiricist ought to be a *constructive* empiricist. There is no logical obstacle for an empiricist (who, ultimately, thinks that all knowledge stems from experience) to fostering methods that warrant belief in the truth of theories in a way that goes beyond belief in their empirical adequacy, and hence, to be scientific realist. Similarly, there is no logical obstacle for an empiricist to being stricter than constructive empiricism. Constructive empiricism does set the boundaries of experience a little further than strict empiricism, and since what empiricism is *not* carved in stone, there is no logical obstacle to setting the boundaries of experience (that is, the reach of legitimate applications of scientific method) even further, as realists demand.

Van Fraassen tied empiricism to a sharp distinction between observable and unobservable *entities*. This is a step forward *vis-à-vis* the more traditional empiricist distinction in terms of observational and theoretical terms and predicates. Drawing the distinction in terms of entities allows for the description of observable entities to be fully theory-laden. Yet, van Fraassen insisted, even theoretically described, an entity does not cease to be observable if a suitably placed observer *could* perceive it with the naked eye. Long before van Fraassen, Grover Maxwell (1962) had denied this entity-based distinction, arguing that observability is a vague notion and that, in essence, *all* entities are observable under suitable circumstances. He rested on the view that “observability” should be best understood as detectability through or by means of something. If observability is understood thus, there are continuous degrees

of observability, and hence there is no natural and non-arbitrary way to draw a line between observable and unobservable entities. The rebuttal of Maxwell's argument requires that naked-eye observations (which are required to tell us what entities are *strictly* observable) form a special kind of detection that is set apart from any other way of detecting the presence of an entity (e.g. by the microscope). Be that as it may, the issue is not whether the entity-based distinction can be drawn but its epistemic relevance: why should the observable/unobservable distinction define the border between what is epistemically accessible and what is not?

What is scientific realism?

I take scientific realism to consist in three theses (or stances) (see Psillos 1999).

The metaphysical thesis: The world has a definite and mind-independent structure.

The semantic thesis: Scientific theories should be taken at face-value.

The epistemic thesis: Mature and predictively successful scientific theories are (approximately) true of the world.

In the last quarter of the twentieth century it was the third thesis of scientific realism that was primarily the focus. Briefly put, this is an epistemically optimistic thesis: science can and does deliver theoretical truth no less than it can and does deliver observational truth. The centrality of this thesis for realism was dictated by the perception of what the opposing view was: *agnostic* or *skeptical* versions of empiricism. A strong critique of realism was that the abductive-ampliative methodology of science fails to connect empirical success and truth robustly. Hence, it was argued, the claim that scientific theories are true is never warranted. Two arguments were put forward in favor of this view. The first relies on the so-called underdetermination of theories by evidence: two or more mutually incompatible theories can nonetheless be empirically congruent and hence equally empirically warranted. Given that at most one of them can be true, the semantic thesis can still stand but accompanied by a skeptical attitude towards the truth of scientific theories. The second argument was so-called pessimistic induction. As Larry Laudan (1984) pointed out, the history of science is replete with theories which were once considered to be empirically successful and fruitful, but which turned out to be false and were abandoned. If the history of science is the wasteland of aborted "best theoretical explanations" of the evidence, it might well be that current best explanatory theories will take the route to this wasteland in due course.

The argument from underdetermination rests on two questionable premises. (1): for any theory T there is at least another one incompatible theory T' which is empirically congruent with T. (2): if two theories are empirically equivalent, then they are epistemically equivalent too. Both premises have been forcefully challenged by realists. Some have challenged (1) on the grounds that the thesis it encapsulates is not proven. Others have objected to (2). There are, on the face of it, two strategies available. One (2a) is to argue that even if we take only empirical evidence to bear on the epistemic support of the theory, it does not follow that the class of the observational consequences of the theory is coextensional with the class of empirical facts that support to the theory. An obvious counterexample to the claim of coextensionality is that a theory can get *indirect*

support from evidence it does not directly entail. The other strategy (2b) is to note that theoretical virtues (and explanatory power, in particular) are epistemic in character and hence they can bear on the support of the theory. Here again, there are two options available to realists: (2b.i) to argue (rather implausibly) that some theoretical virtues are constitutive marks of truth; or (2b.ii) to argue for a broad conception of evidence which takes the theoretical virtues to be broadly empirical and contingent marks of truth. Option (2b.ii) is an attractive strategy because it challenges the strictly empiricist conception of evidence and its relation to rational belief.

When it comes to the pessimistic induction, the best defense of realism has been to try to reconcile the historical record with some form of realism. In order to do this, realists should be more selective in what they are realists about. A claim that has emerged with some force is that theory-change is not as radical and discontinuous as the opponents of realism have suggested. Realists have aimed to show that there are ways to identify the theoretical constituents of abandoned scientific theories which essentially contributed to their successes, separate them from others that were “idle” and demonstrate that those components which made essential contributions to the theory’s empirical success were retained in subsequent theories of the same domain (see Kitcher 1993; Psillos 1999). Then, the fact that our current best theories may be replaced by others does not, necessarily, undermine scientific realism. All it shows is that (a) we cannot get at the truth all at once; and (b) our judgments from empirical support to approximate truth should be more refined and cautious, in that they should only commit us to the theoretical constituents that do enjoy evidential support and contribute to the empirical successes of the theory. Realists ground their epistemic optimism on the fact that newer theories incorporate many theoretical constituents of their superseded predecessors, especially those constituents that have led to empirical successes. The substantive continuity in theory-change suggests that a rather stable network of theoretical principles and explanatory hypotheses has emerged, which has survived revolutionary changes, and has become part and parcel of our evolving scientific image of the world.

Causation and explanation

The collapse of the synthetic a priori took with it the Kantian idea that the principle of causation (“everything that happens, that is, begins to be, presupposes something upon which it follows by rule”), was synthetic a priori. The logical positivists took to heart Hume’s critique of the supposed necessary connection between cause and effect. The twist they gave to this critique was based on their verificationist criterion of meaning: positing a necessary link between two events would be tantamount to committing a kind of nonsense since all attempts to verify it would be futile.

Causal explanation

A central element of the empiricist project was to legitimize – and demystify – the concept of causation by subsuming it under the concept of lawful explanation, which

in turn, was modeled on deductive arguments. This project culminated in Hempel and Oppenheim's (1948) Deductive-Nomological (DN) model of explanation. According to this, to offer an explanation of an event e is to construct a valid deductive argument of the following form:

Antecedent/initial conditions

Statements of laws

Therefore, e (event/fact to be explained)

So, when the claim is made that event c causes event e , it should be understood as follows: there are relevant laws in virtue of which the occurrence of the antecedent condition c is nomologically sufficient for the occurrence of the event e . It has been a standard criticism of the DN-model that, insofar as it aims to offer sufficient and necessary conditions for an argument to count as a bona fide explanation, it fails. A view that became dominant in the 1960s was that the DN-model fails precisely because it ignores the role of *causation* in explanation. In other words, there is more to the concept of causation than what can be captured by DN-explanations.

The Humean view may be entitled the "regularity view of causation." But an opposite view that became prominent in the twentieth century, owing mostly to the work of Curt John Ducasse (1881–1969), is that what makes a sequence of events causal is a local tie between the cause and the effect, or an intrinsic feature of the particular sequence. Causation, non-Humeans argue, is essentially singular: a matter of *this* causing *that*. Concomitant to this view is the thought that causal explanation is also singular. It is not a matter of subsuming the explanandum under a law but rather a matter of offering causally relevant information about why an event occurred.

The platitudes of causation

No matter how one thinks about causation, there are certain platitudes that this concept should satisfy. One of them may be called the "difference platitude": causes are difference-makers, that is, things would be different if the causes of some effects were absent. This platitude is normally cast in two ways. The first is the *counterfactual* way: if the cause hadn't been, the effect wouldn't have been either. David Lewis (1941–2001) defined causation in terms of the *counterfactual dependence* of the effect on the cause: the cause is rendered counterfactually necessary for the effect (1986). The other is the *probabilistic* way: causes raise the *chances* of their effects, namely, the probability that a certain event happens is higher if we take into account its cause than if we don't. This thought has led to the development of theories of *probabilistic causation*. Some philosophers, most notably Patrick Suppes (1984) and Nancy Cartwright (1983), claimed that the existence of probabilistic causation is already a good argument against the view that causation is connected with invariable sequences or regularities.

The origin of the probabilistic understanding of causation is found in Hempel's (1965) attempts to offer a general model of statistical explanation. Hempel argued that singular events whose probability of happening is less than unity can be explained by being subsumed under a statistical (or probabilistic) law. He advanced an *inductive-statistical* model of explanation according to which to explain an event is to construct an inductive argument one premise of which states a statistical generalization. But it became immediately obvious that this model suffered from problems. For instance, it can only explain events that are likely to happen, although unlikely events do happen and they too require explanation. Besides, statistical laws require fixing reference classes and it is not obvious what principles should govern this fixing. Wesley Salmon (1925–2001) advanced an alternative model of statistical explanation (1971), known as the *statistical-relevance* model, according to which in judging whether a factor C is relevant to the explanation of an event that falls under type E , we look at how taking C into account affects the probability of E . In particular, a factor C explains the occurrence of an event E , if $\text{prob}(E/C) > \text{prob}(E)$. Note that it is *not* required that the probability $\text{prob}(E/C)$ be high. All that is required is that there is a difference, no matter how small, between the two probabilities. Although it was initially thought that this model could offer an account of causal relevance, Salmon soon accepted the view that there is more to probabilistic causation than relations of statistical dependence. Mere correlations can conform to the statistical relevance model, but they are not causal. The theories of probabilistic causation that emerged in the 1980s attempted to figure out what more should be added to relations of statistical relevance to get causation.

Another central platitude of the concept of causation may be called the “recipe platitude”: causes are recipes for producing or preventing their effects. This platitude is normally cast in terms of manipulability: causes can be manipulated to bring about certain effects. G. H. von Wright (1906–2003) developed this thought into a full-blown theory of causation. Since manipulation is a distinctively human action, he thought that the causal relation is dependent upon the concept of human action. But his views were taken to be too anthropomorphic. Yet in the last quarter of the century there were important attempts to give a more objective gloss to the idea of manipulation (mostly developed by James Woodward).

In the same period, several philosophers tried to show that there is more to causation than regular succession by positing a *physical mechanism* that links cause and effect. Salmon advanced a mechanistic approach (1984), saying roughly that an event c causes an event e if and only if there is a causal process that connects c and e . He characterized as “causal” those processes that are capable of transmitting a mark, where a mark is a modification of the structure of a process. Later on, Salmon and Phil Dowe took causation to consist in the exchange or transfer of some conserved quantity, such as energy-momentum or charge. Such accounts may be called “transference models” because they claim that causation consists in the transfer of something (some physical quantity) between the cause and its effect. But there is a drawback. Even if it is granted that these models offer neat accounts of causation at the level of physical processes, they can be generalized as accounts of causation *simpliciter* only if they are married to

strong reductionist views that all worldly phenomena (be they social or psychological or biological) are, ultimately, reducible to physical phenomena.

Despite the centrality of the concept of causation in the ways we conceptualize the world and in the analysis of a number of other concepts, there is hardly any agreement about what causation is. It might not be implausible to think that, although traditionally causation has been taken to be a single, unitary, concept, there may well be two (or more) concepts of causation.

What is a law of nature?

The Deductive-Nomological model of explanation, as well as any attempt to tie causation to laws, faced a significant conceptual difficulty: the problem of how to determine what the laws of nature were. Most Humeans adopted the “regularity view of laws”: laws of nature are regularities. Yet they had a hurdle to jump: not all regularities are causal. Nor can all regularities be deemed laws of nature. So they were forced to draw a distinction between the good regularities (those that constitute the laws of nature) and the bad ones, i.e. those that are, as John Stuart Mill put it, “conjunctions in some sense accidental.” Only the former can underpin causation and play a role in explanation. The predicament that Humeans were caught in is this. Something (let’s call it the property of lawlikeness) must be added to a regularity to make it a law of nature. But what can this be?

Humean conceptions of laws

The first systematic attempt to pin down this elusive property of lawlikeness was broadly epistemic. The thought, advanced by A. J. Ayer (1901–89), Richard Braithwaite (1900–90), and Nelson Goodman among others, was that inquirers have different epistemic attitudes towards laws and accidents. Lawlikeness was taken to be the property of those generalizations that play a certain epistemic role: they are believed to be true, and they are so believed because they are confirmed by their instances and are used in proper inductive reasoning. But this purely epistemic account of lawlikeness fails to draw a robust line between laws and accidents.

A much more promising attempt to draw the line between laws and accidents is what may be called the “web of laws” view. According to this, the regularities that constitute the laws of nature are those that are expressed by the axioms and theorems of an ideal deductive system of our knowledge of the world, a system that strikes the best balance between simplicity and strength. Whatever regularity is not part of this best system is merely accidental: it fails to be a genuine law of nature. The gist of this approach, which was advocated by Mill, and in the twentieth century by Ramsey and Lewis, is that no regularity, taken in isolation, can be deemed a law of nature. The regularities that constitute laws of nature are determined in a kind of holistic fashion by being parts of a structure. Despite its many attractions, this view faces the charge that it cannot offer a fully objective account of laws of nature. For instance, it is commonly argued that how our knowledge of the world is organized into a simple and

strong deductive system is, by and large, a subjective matter. But this kind of criticism is overstated. There is nothing in the web of laws approach that makes laws mind-dependent. The regularities that are laws are fully objective, and they govern the world irrespective of our knowledge of them, and of our being able to identify them. In any case, as Ramsey pointed out, it is a fact about the world that some regularities form, objectively, a system; this is that *the world has an objective nomological structure*, in which regularities stand in certain relations to each other; these relations can be captured (or expressed) by relations of deductive entailment in an ideal deductive system of our knowledge of the world. Ramsey's suggestion grounds an objective distinction between laws and accidents in a *worldly* feature: that the world has a certain nomological structure (see Psillos 2002).

Non-Humean conceptions of laws

In the 1970s, David Armstrong (1983), Fred Dretske (1977), and Michael Tooley (1977) put forward the view that lawhood cannot be reduced to regularity. Lawhood, they claimed, is a certain necessitating relation among properties (i.e. universals). An attraction of this view is that it makes clear how laws can cause anything to happen: they do so because they embody causal relations among properties. But the central concept of nomic necessitation is still not sufficiently clear. In particular, it is not clear how the necessitating relation between the property of *F*-ness and the property of *G*-ness makes it the case that All *F*s are *G*s.

Both the Humeans and the advocates of the Armstrong–Dretske–Tooley view agreed that laws of nature are *contingent*. But a growing rival thought was that if laws did not hold with some kind of objective necessity, they could not be robust enough to support either causation or explanation. Up until the early 1970s, this talk of objective necessity in nature was considered almost unintelligible. It was Kripke's (1972) liberating views that changed the scene radically. He broke with the Kantian tradition which equated the necessary and the a priori as well as with the empiricist tradition which equated the necessary with the analytic. He argued that there are necessarily true statements which can be known a posteriori.

As a result of this, there has been a growing tendency among non-Humeans to take laws of nature to be metaphysically necessary. This amounts to a radical denial of the contingency of laws. Along with it came a resurgence of Aristotelianism in the philosophy of science. The advocates of the view that the laws are *contingent* necessitating relations among properties thought that, although an appeal to (natural) properties is indispensable for the explication of lawhood, the properties themselves are passive and freely recombinable. Accordingly, there can be possible worlds in which some properties are not related in the way they are related in the actual world. But the advocates of metaphysical necessity took the stronger line that laws of nature flow from the essences of properties. Insofar as properties have essences, and insofar as it is part of their essence to endow their bearers with a certain behavior, it follows that the bearers of properties *must* obey certain laws, those that are issued by their properties. Essentialism was treated with suspicion in most of the twentieth century,

partly because essences were taken to be discredited by the advent of modern science and partly because the admission of essences (and the concomitant distinction between essential and accidental properties) created logical difficulties. Essentialism required the existence of *de re* necessity, that is, natural necessity, since if it is of the essence of an entity to be thus-and-so, it is *necessarily* thus-and-so. But before Kripke's work, the dominant view was that all necessity was *de dicto*, that is, it applies, if at all, to propositions and not to things in the world.

The thought that laws are metaphysically necessary gained support from the (neo-Aristotelian) claim that properties are active powers. Many philosophers argue that properties are best understood as powers since the only way to identify them is via their causal role. Accordingly, two seemingly distinct properties that have exactly the same powers are, in fact, one and the same property. And similarly, one cannot ascribe different powers to a property without changing this property. It's a short step from these thoughts to the idea that properties are not freely recombinable: there cannot be worlds in which two properties are combined by a different law from the one that unites them in the actual world. Actually, on this view, it does not even make sense to say that properties are united by laws. Rather, properties – qua powers – *ground* the laws.

The emergence of a new dogma?

From the 1970s on, a new over-arching approach to many problems in the philosophy of science has emerged: Bayesianism. This is a mathematical theory based on the probability calculus that aims to provide a general framework in which key concepts such as rationality, scientific method, confirmation, evidential support, and sound inductive inference are cast and analyzed. It borrows its name from a theorem of probability calculus: Bayes's Theorem. Let H be a hypothesis and e the evidence. Bayes's theorem says:

$$\text{prob}(H/e) = \frac{\text{prob}(e/H)\text{prob}(H)}{\text{prob}(e)}, \text{ where } \text{prob}(e) = \text{prob}(e/H)\text{prob}(H) + \text{prob}(e/-H)\text{prob}(-H).$$

The unconditional $\text{prob}(H)$ is called the prior probability of the hypothesis, the conditional $\text{prob}(H/e)$ is called the posterior probability of the hypothesis *given* the evidence, and the $\text{prob}(e/H)$ is called the likelihood of the evidence given the hypothesis. (The theorem can be easily reformulated so that background knowledge is taken into account.) In its dominant version, Bayesianism is subjective or personalist because it claims that probabilities express subjective degrees of belief. It is based on the significant mathematical result – proved by Ramsey and Bruno de Finetti (1906–85) – that subjective degrees of beliefs (expressed as fair betting quotients) satisfy Kolmogorov's axioms for probability functions. The central idea, known as the Dutch-book theorem, is that unless the degrees of beliefs that an agent possesses, *at any given time*, satisfy the axioms of the probability calculus, she is subject to a Dutch-book, that is, to a set of synchronic bets such that they are all fair by her own

lights, and yet, taken together, make her suffer a net loss come what may. The thrust of the Dutch-book theorem is that there is a *structural incoherence* in a system of degrees of belief that violate the axioms of the probability calculus.

There have been attempts to extend Bayesianism to belief-revision by the technique of conditionalization. It is supposed to be a canon of rationality that agents should *update* their degrees of belief by conditionalizing on the evidence: $\text{Prob}_{\text{new}}(-) = \text{Prob}_{\text{old}}(-/e)$, where e is the total evidence. (Conditionalization can be either strict, where the probability of the learned evidence is unity, or Jeffrey (due to Richard Jeffrey, 1926–2002) where the evidence one updates on can have probability less than 1.) The penalty for not conditionalizing on the evidence is liability to a Dutch-book strategy: the agent can be offered a set of bets over time such that (a) each of them taken individually will seem fair to her at the time when it is offered; but (b) taken collectively, they lead her to suffer a net loss, come what may. But critics point out that there is no general proof of the conditionalization rule (see Earman 1992: 46–51).

The Bayesians' reliance on subjective prior probabilities has been a constant source of dissatisfaction among their critics. It is claimed that purely subjective prior probabilities fail to capture the all-important notion of rational or reasonable degrees of belief. But, in all fairness, it's been extremely difficult to articulate the notion of a *rational* degree of belief. Attempts to advance a more objectivist Bayesian theory (based on the principle of indifference, or, worse, based on Keynes's thought that "logical intuition" enables agents to "see" the logical relation between the evidence and the hypothesis) have failed. This has led Bayesians to insist on the indispensability of subjective priors in inductive reasoning. It has been argued that, in the long run, the prior probabilities wash out: even widely different prior probabilities will converge, in the limit, to the same posterior probability, if agents conditionalize on the same evidence. But this is little consolation because, apart from the fact that in the long run we are all dead, the convergence theorem holds only under limited and very well-defined circumstances that can hardly be met in ordinary scientific cases. Since the 1980s, subjective Bayesianism has become the orthodoxy in confirmation theory. The truth is that this theory of confirmation has had many successes. Old tangles, like the Ravens paradox or the grue problem, have been resolved. New tangles, like the problem of old evidence (how can a piece of evidence that is already known and entailed by the theory raise the posterior probability of the theory above its prior?), have been, to some extent, resolved. But to many, this success is not very surprising since the subjective nature of prior probabilities makes, in effect, prior probabilities free parameters that Bayesians can play with.

Can there be non-probabilistic accounts of confirmation? Clark Glymour (1980) developed his own "bootstrapping" account of confirmation that was meant to be an improvement over Hempel's theory. Glymour's key idea was that confirmation is three-place relation: the evidence confirms a hypothesis *relative* to a theory (which may be the very theory in which the hypothesis under test belongs). This account gave prominent role to explanatory considerations, but failed to show how confirmation of the hypothesis can give us reasons to believe in the hypothesis. The thought here

is that unless probabilities are introduced into a theory of confirmation, there is no connection between confirmation and reasons for belief.

Among the alternatives to Bayesian confirmation that were developed over the last two decades of the century and that relied on probabilities, two stand out. The first is Deborah Mayo's (1996) error-statistical approach, which rests on the standard Neyman–Pearson statistics and utilizes error-probabilities understood as objective frequencies. They refer to the experimental process itself and specify how reliably it can discriminate between alternative hypotheses. Mayo focuses on learning from experiments and argues that experimental learning has a life and a model of its own. Experimental learning may be messy and piecemeal, it may require a tool-kit of techniques that ensure low error probabilities, but it can be embedded in a rather neat statistical framework that makes a virtue out of our errors. Yet, in Mayo's account, it is not clear how theories are confirmed by the evidence.

The other alternative to Bayesian confirmation is by Peter Achinstein (2001). He has advanced an absolute notion of confirmation, according to which e is evidence for H only if e is not evidence for the denial of H . This is meant to capture the view that evidence should provide a good reason to believe. But Achinstein's real novelty consists in his claim that this absolute conception of evidence is not sufficient for reasonable belief. What must be added is that there is an *explanatory connection* between H and e in order for e to be evidence for H .

Indeed, to many philosophers, explanatory considerations should be a rational constraint on inference and should inform any account of scientific method which tallies with scientific practice. But it is hard to see how explanatory considerations can be accommodated within the Bayesian account of scientific method, which, in effect, accepts the following equation: *scientific method* = *experience* + *probability calculus (including logic)*. This is certainly an improvement over Duhem's equation: *scientific method* = *experience* + *logic*. The need to import explanatory considerations into the scientific method (almost) swayed Duhem away from his strict equation (though Duhem rightly noted that explanatory considerations do not admit of an algorithmic account since they are values, broadly understood). Bayesians may not be easily swayed out of their own equation, since they may argue that explanatory considerations can be reflected in a suitable distribution of subjective prior probabilities among competing theories. But this kind of answer is inadequate if only because explanatory considerations are supposed to be objective, or in any case, not merely subjective. In fact, Bayesian accounts of confirmation fail to account for the methodological truism that goodness of explanation goes hand in hand with warrant for belief. It's not hard to see that on a Bayesian account of confirmation, and given that the theory T entails the set of its observational consequences T_O , no theory T can ever be more credible than a "theory" T_O . Yet T may well explain T_O , whereas T_O does not explain itself. As Glymour (1980: 83) put it, Bayesians render theories a "gratuitous risk."

Bayesianism is an admirable theory but, to my mind, there is more to scientific method than subjective prior probability and Bayes's Theorem. Of course, the tough issue is to offer a systematic and well-worked-out alternative to Bayesianism. Perhaps this alternative should be based on the broad idea that the kernel of scientific method

is eliminative induction, where elimination is informed by explanatory considerations. But any account of scientific method should make room for epistemic values. Philosophy of science in the twentieth century has been animated by a legend: that theory-testing and theory-evaluation in science is based on a fully objective and general, value-free, presuppositions-free, evidence-driven, and domain-neutral scientific method. The naturalist turn did a lot to rebut the a prioristic version of the legend but it seems it's not gone all the way!

Outlook

Kant raised the question "How is metaphysics possible at all?" and went on to answer it by defending the synthetic a priori character of metaphysical propositions, which he thought would make science possible too. In the twentieth century, some of the philosophers who were deeply influenced by Kant's thought ended up considering metaphysics meaningless because it transgresses the bounds of meaningful discourse captured by mathematics and science. But as the century was drawing to a close, the genie came out of the bottle again: philosophers of science had to swim in deep metaphysical waters in order to address a number of key issues. Far from being a philosophical dirty word, "metaphysics" captures a legitimate and indispensable area of philosophical engagement with science. Understanding the basic structure of the world requires doing serious metaphysics, although now we know full well that no metaphysics can be seriously engaged with in abstention from what science tells us about the world. The big movement of the twentieth century was from issues related to language and meaning to issues related directly to the furniture of the world. It was also a movement from rational reconstructions and macro-models of science to the details of the individual sciences (especially, sciences other than physics).

Much to my disappointment, the present survey has not touched upon a number of issues. My greatest regret is that there is no discussion of feminist philosophies of science, a distinctive product of the twentieth century that has brought to light a key aspect of the issue of scientific objectivity. As the twentieth century recedes into history, what kinds of issue are dead and what are still alive?

Although in philosophy no issue really dies, it's safe to say that some positions are no longer viable (e.g. reductive understandings of the meaning of theoretical terms, strong instrumentalist accounts of scientific theories, naive regularity views of laws, algorithmic approaches to scientific method, global type-reductive accounts of inter-theoretic relations). Perhaps the whole philosophical issue of conceptual change is pretty much exhausted (with the exception of empirical studies of cognitive models of science), as is the issue of devising grand theories of science. Perhaps one should also be skeptical about the insights that can be gained by formal approaches to the nature and structure of scientific theories.

Surprisingly, the realism debate keeps coming back. The nature of rational judgment and the role of values in it are certainly fertile areas. And so are the nature of causation and the structure of causal inference. Adequate descriptions of ampliative reasoning are still in demand. And the prospects of structuralism occupy the thought

of many philosophers of science. Most importantly, more work should be done on the relations between central issues in the philosophy of science and topics and debates in other areas of philosophy (most notably, metaphysics and epistemology).

A breath of fresh air has come from the conscious effort of Michael Friedman (1999) and many others to re-evaluate and reappraise the major philosophical schools and the major philosophers of science of the twentieth century. The philosophical battlegrounds of the twentieth century saw many attacks on strawmen and a number of pyrrhic victories. The time is now ripe to start repainting the complex landscape of twentieth-century philosophy of science.

Note

- 1 This survey reflects my own prejudices and debts; my own intellectual influences and debts. Others would have done things differently. I hope that I have not been unfair to the broad tradition I have been brought up in. Many thanks to Marc Lange and to three anonymous readers for their encouragement and insightful comments.

References

- Achinstein, P. (2001) *The Book of Evidence*. New York: Oxford University Press.
- Armstrong, D. M. (1983) *What Is a Law of Nature?* Cambridge: Cambridge University Press.
- Boyd, R. (1973) "Realism, underdetermination and the causal theory of evidence." *Noûs* 7: 1–12.
- Carnap, R. (1928) *The Logical Structure of the World*. English translation by R. A. George, Berkeley: University of California Press, 1961.
- (1934) *The Logical Syntax of Language*. English translation by A. Smeaton, London: Kegan Paul, 1937.
- (1950) *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- (1956) "The methodological character of theoretical concepts." In *Minnesota Studies in the Philosophy of Science* 1, Minneapolis: University of Minnesota Press, pp. 38–76.
- Cartwright, N. (1983) *How the Laws of Physics Lie*. Oxford: Clarendon Press.
- Cassirer, E. (1910) *Substance and Function*. English translation by W. C. and M. C. Swabey, Chicago: Open Court, 1923.
- Craig, W. (1956) "Replacements of auxiliary assumptions." *Philosophical Review* 65: 38–55.
- Dretske, F. I. (1977) "Laws of nature." *Philosophy of Science* 44: 248–68.
- Duhem, P. (1906) *The Aim and Structure of Physical Theory*. English translation by P. Wiener, Princeton, NJ: Princeton University Press, 1954.
- Earman, J. (1992) *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*, Cambridge MA: MIT Press.
- Frege, G. (1884) *The Foundations of Arithmetic*. English translation by J. L. Austin, Evanston, IL: Northwestern University Press, 1980.
- Friedman, M. (1999) *Reconsidering Logical Positivism*. Cambridge: Cambridge University Press.
- Giere, R. (1988) *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press.
- Glymour, C. (1980) *Theory and Evidence*. Princeton, NJ: Princeton University Press.
- Goodman, N. (1954) *Fact, Fiction and Forecast*. Cambridge, MA: Harvard University Press.
- Hempel, C. (1945) "Studies in the logic of confirmation." *Mind* 54: 1–26.
- (1958) "The theoretician's dilemma: a study in the logic of theory construction." In *Minnesota Studies in the Philosophy of Science* 2, Minneapolis: University of Minnesota Press, pp. 37–98.
- (1965) *Aspects of Scientific Explanation*. New York: Free Press.
- Hempel, C. and P. Oppenheim (1948) "Studies in the logic of explanation." *Philosophy of Science* 15: 135–75.

- Hilbert, D. (1899) *The Foundations of Geometry*. English translation by L. Unger, La Salle, IL: Open Court, 1971.
- Keynes, J. M. (1921) *A Treatise on Probability: The Collected Works of J. M. Keynes*, vol. 7. London and Cambridge: MacMillan and Cambridge University Press.
- Kitcher, P. (1993) *The Advancement of Science*. Oxford: Oxford University Press.
- Kripke, S. (1972) "Naming and necessity." In D. Davidson and G. Harman (eds.) *Semantics of Natural Language*, Dordrecht: Reidel, pp. 253–355, 763–9.
- Kuhn, T. S. (1962) *The Structure of Scientific Revolutions*. 2nd, enlarged edn., 1970. Chicago: University of Chicago Press.
- Lakatos, I. (1970) "Falsification and the methodology of scientific research programmes." In I. Lakatos and A. Musgrave (eds.) *Criticism and the Growth of Knowledge*, Cambridge: Cambridge University Press, pp. 91–196.
- Laudan, L. (1984) *Science and Values*. Berkeley: University of California Press.
- Lewis, D. (1986) "Causation." In *Philosophical Papers*, vol. 2, Oxford: Oxford University Press, pp. 159–213.
- Maxwell, G. (1962) "The ontological status of theoretical entities." In *Minnesota Studies in the Philosophy of Science* 3, Minneapolis: University of Minnesota Press, pp. 3–27.
- Mayo, D. G. (1996) *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press.
- Newman, M. H. A. (1928) "Mr. Russell's 'Causal theory of perception'." *Mind* 37: 137–48.
- Poincaré, H. (1902) *Science and Hypothesis*. English translation, New York: Dover, 1905.
- Popper, K. (1959) *The Logic of Scientific Discovery*. London: Hutchinson.
- Psillos, S. (1999) *Scientific Realism: How Science Tracks Truth*. London and New York: Routledge.
- (2000) "An introduction to Carnap's 'Theoretical concepts in science'" (with Carnap's: "Theoretical concepts in science"). *Studies in History and Philosophy of Science* 31: 151–72.
- (2002) *Causation and Explanation*. Montreal: McGill-Queens University Press.
- Putnam, H. (1973) "Explanation and reference." In G. Pearce and P. Maynard (eds.) *Conceptual Change*, Dordrecht: Reidel, pp. 199–221.
- (1975) *Mathematics, Matter and Method: Philosophical Papers*, vol.1. Cambridge: Cambridge University Press.
- Quine, W. V. O. (1936) "Truth by convention." In O. H. Lee (ed.) *Philosophical Essays for A. N. Whitehead*, New York: Longmans, pp. 90–124. Repr. in Quine's *The Ways of Paradox and Other Essays*, Cambridge, MA: Harvard University Press, 1976, pp. 77–106.
- (1951) "Two dogmas of empiricism." *Philosophical Review* 60: 20–43.
- (1960) *Word and Object*. Cambridge, MA: MIT Press.
- (1969) "Epistemology naturalized." In *Ontological Relativity and Other Essays*, Cambridge, MA: Harvard University Press, pp. 69–90.
- Ramsey, F. (1929) "Theories." In *The Foundations of Mathematics and Other Essays*, ed. R. B. Braithwaite, London: Routledge & Kegan Paul, 1931, pp. 212–36.
- Reichenbach, H. (1921) *The Theory of Relativity and A Priori Knowledge*. English translation by M. Reichenbach, Berkeley: University of California Press, 1965.
- (1928) *Philosophy of Space and Time*. English translation by M. Reichenbach and J. Freund, New York: Dover, 1958.
- Russell, B. (1927) *The Analysis of Matter*. London: Routledge & Kegan Paul.
- Salmon, W., R. C. Jeffrey, and J. G. Greeno (eds.) (1971) *Statistical Explanation and Statistical Relevance*. Pittsburgh: University of Pittsburgh Press.
- (1984) *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.
- Schlick, M. (1918) *General Theory of Knowledge*. 2nd German edn., 1925. English translation by A. E. Blumberg, Vienna and New York: Springer-Verlag, 1974.
- Sellars, W. (1956) "Empiricism and the philosophy of mind." In *Minnesota Studies in the Philosophy of Science* 1, Minneapolis: University of Minnesota Press, pp. 253–329.
- (1963) *Science, Perception and Reality*. Atascadero, CA: Ridgeview, 1991.
- Smart J. J. C. (1963) *Philosophy and Scientific Realism*. London: Routledge & Kegan Paul.

- Suppe, Frederick. (1980). *The Semantic Conception of Theories and Scientific Realism*. Urbana: University of Illinois Press
- Suppes, P. (1984) *Probabilistic Metaphysics*. Oxford: Blackwell.
- Tooley, M. (1977) "The nature of laws." *Canadian Journal of Philosophy* 7: 667–98.
- van Fraassen, Bas C. (1980) *The Scientific Image*. Oxford: Clarendon Press.

Further reading

- Bird, A. (1998) *Philosophy of Science*. Montreal: McGill-Queens University Press. (A thorough critical introduction to most central areas and debates in the philosophy of science, written with elegance and insight).
- Hempel, C. (1966) *Philosophy of Natural Science*. Englewood Cliffs, NJ: Prentice-Hall. (An all-time classic introduction to the philosophy of science, tantalizingly modern).
- Hitchcock, C. (ed.) (2004) *Contemporary Debates in Philosophy of Science*. Oxford: Blackwell. (Well-known contemporary thinkers present opposing arguments on eight hotly debated issues in the philosophy of science).
- Ladyman, J. (2002) *Understanding Philosophy of Science*. London and New York: Routledge.
- Lange, M. (ed.) (2006) *Philosophy of Science: An Anthology*. Oxford: Blackwell. (A comprehensive collection of influential papers on a number of central topics in the philosophy of science arranged thematically and accompanied by thoughtful introductions by the editor).
- Papineau, D. (ed.) (1996) *The Philosophy of Science*, Oxford: Oxford University Press. (A short collection of the most influential papers on general methodological issues in the philosophy of science, with an excellent introduction by the editor).
- Psillos, S. (2007) *Philosophy of Science A–Z*. Edinburgh: Edinburgh University Press. (A dictionary of most central concepts, schools of thought, arguments, and thinkers; especially suitable for beginners).
- Psillos, S. and M. Curd (eds.) (2008) *The Routledge Companion to the Philosophy of Science*. London and New York: Routledge. (Fifty-five specially commissioned chapters on all major issues in the philosophy of science, written by distinguished authors and offering a balanced critical assessment of the discipline and issues of current debate).
- Salmon, M. et al. (eds.) (1992) *Introduction to the Philosophy of Science*. Englewood Cliffs, NJ: Prentice-Hall. (Written by members of the History and Philosophy of Science Department of the University of Pittsburgh, it contains thought-provoking advanced introductions to major debates about general issues (such as confirmation and explanation) as well as more special issues in the foundations of the sciences)