

Philosophy of Science circa 1950–2000: Some Things we (should have) Learned

Diogenes
2017, Vol. 61(2) 45–58
Copyright © ICPHS 2016
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/0392192116640719
dio.sagepub.com


Harold I. Brown

Northern Illinois University, USA

During the first half of the twentieth century logical empiricism dominated philosophy of science; it began to lose its hold during the 1950s. Two factors were largely responsible for this change. First, there were persistent failures by logical empiricists to solve problems generated by their own framework, especially providing a formal account of confirmation and an analysis of theoretical terms that meets empiricist structures. Second, the emergence of a rich body of research in the history of science made it clear that the development of even the most successful sciences was more complex and less certain than had previously been assumed. Some of the emerging issues were addressed by Quine (1951) and Wilfrid Sellars (1948, 1953a, 1953b, 1954), but a general recognition that something was seriously wrong came only as the decade waned. Then it came with stunning speed. We can note six works with overlapping themes that appeared in a four-year period from several different intellectual backgrounds: Hanson (1958), Polanyi (1958), Toulmin (1961), Feyerabend (1962), Kuhn (1962), and Putnam (1962). This led to a new body of research and the quest for a new philosophical framework that could replace logical empiricism as a guide to the problems and range of acceptable solutions in philosophy of science. As Gutting has noted (2009: 151) it is now clear that this quest failed and several issues that were recently at the focus of discussion have largely disappeared from the active literature. Yet it would be unfortunate if this work faded completely from the memory of working philosophers of science because there are some important lessons about science and about of philosophy of science that we should have learned. I am going to describe these lessons from my own perspective as someone who lived and worked through this period. No doubt this attempt will be somewhat idiosyncratic; others will draw different lessons—or no lessons at all—from these endeavors. But, I will argue, the lessons I discuss here are important and should be incorporated into ongoing work. I will begin by focusing on the problem of theory choice—especially on the view that theory evaluation should be determined solely by logic and the evidence. This will lead us to several other issues.

Methodology

Formal logic—in particular the powerful tools provided by the new mathematical logic—provided a central focus of logical empiricist research. The attempt to see how much could be accomplished

Corresponding author:

Harold I. Brown, 541 N 7th Street, DeKalb, IL 60115, USA.

Email: hibrown@niu.edu

using these tools was surely a worthwhile project that had at least one major outcome: recognition that while logic provides a central part of any normative account of theory choice, logic and evidence alone are not sufficient to dictate theory choice. In order to focus on logic in this section I will temporarily assume that we have a non-problematic body of relevant evidence when evaluating theories.

One issue arises at once: the new logic is deductive logic but theory choice requires that we go beyond deduction since interesting generalizations and theories go beyond just a statement of the evidence. In the usual terminology, deduction is non-ampliative, but theory-choice requires induction which is ampliative. This difference raises initial doubts about the goal of providing a purely formal account of induction, doubts that were enhanced by the pursuit of this goal. Two key developments underlined the problems. Each generated an enormous literature.

The first development gave us the so-called “paradoxes of confirmation” (Hempel, 1945). The most obvious and straightforward attempt to specify what *counts as* confirmatory evidence for a generalization led to a surprising result. Briefly, Hempel considered the proposal that a universal generalization of the form “All A are B” is supported by items that are both A and B—that is, by items that match both the subject and predicate of the generalization.¹ But this generalization is logically equivalent to its contrapositive “All non-B are non-A” which would, according to the proposed criterion, be supported by any item that is both non-B and non-A. In the standard example, “All ravens are black” would get empirical support from every observation of a non-black non-raven, such as my brown desk. Many found this unacceptable although some, including Hempel, argued that this result is correct when properly understood. Note that this issue arose before any question of degree or strength of a confirmation came into play.

The second development was Goodman’s (1955) “New Riddle of Induction” which, in effect, introduced conceptual change into discussions of confirmation. It is important to remember that Goodman’s aim was to undermine the thesis that we can give a *complete* account of induction in syntactical terms. By moving from “All emeralds are green” to “All emeralds are grue” Goodman introduced an alternative generalization that had a different predicate concept than the original, and gave different predictions for what would be observed in some future test, but was equally confirmed by the available evidence given a purely syntactic criterion. Two points require special emphasis. First, Goodman’s own conclusion was that the syntactic criterion had to be supplemented by an additional consideration—in his view, consideration of the previous history of the use of these predicates. Second, Goodman’s challenge can now be seen as a first step toward a much wider challenge from consideration of actual conceptual change in the history of science. I will return to this topic. Goodman’s proposed solution is also a step towards an approach that I will discuss shortly in the present section. Before doing so, I want to introduce another problem with theory choice.

The problems mentioned so far arise when we focus on positive support of claims that go beyond the available evidence. But, as Popper stressed, falsification of hypotheses requires only deductive logic. Popper agreed that, given the evidence, theory evaluation should depend only on logic but insisted that logic is non-ampliative. He thus rejected the very idea of positive support for a generalization or theory. Yet we have grounds for rejecting a thesis when we deduce a testable result from that thesis and the evidence shows that result to be false. Moreover, the argument from the falsity of a conclusion to the falsity of a premise is itself a deductive argument. But this focus on falsification brought out a further limitation of the role of logic in an account of theory choice. In the most interesting and important cases, derivation of a testable conclusion requires multiple premises. The discovery that a conclusion derived from these premises is false guarantees that something is wrong among our premises but tells us nothing about how many premises are in error, or which they might be.

A way of dealing with the limitations of logic as the basis for theory choice had been foreshadowed by Goodman. More detailed variations on this approach were proposed, in different terminology, by Sellars, Toulmin, Kuhn, and Putnam; Kant provided the grandmother of the approach. Kant was also concerned with the failure of logic as pointed out by Hume. In effect, Hume emphasized that induction is not deduction. Thus attempts to project past conjunctions of properties into the future are not logically necessary. No matter how many instances we have of A associated with B—whether these occur together or in sequence—the claim that they will not co-occur in new cases is logically consistent. Recognizing no other form of logic, Hume concluded that we have no grounds based in reason for maintaining that A and B will continue to be associated. Kant's response was to introduce a new form of logic—transcendental logic. For our purposes we can leave this claim to have extended logic aside and just focus on one outcome: synthetic a priori propositions. These propositions make substantive claims about the world we experience and thus have consistent negations; instances that contradict these claims are conceivable. But these claims are established a priori and thus are not subject to empirical refutation. As a result, when we encounter a case that seems to challenge a synthetic a priori proposition, such as an event for which we cannot find a cause, we conclude that the fault lies with the researcher rather than the proposition. For Kant, synthetic a priori propositions play a double role in scientific research: they make substantive claims about the world we experience (the realm, for Kant, of scientific research), and they provide a part of the methodology of science. They tell us what kinds of questions we should ask and what kinds of answers we should accept as appropriate.

Logical empiricists rejected the very notion of a synthetic a priori proposition. It was taken for granted that all propositions are either analytic a priori and thus have inconsistent negations, or are synthetic a posteriori and subject to refutation by experience. Sellars challenged this view, with full awareness of its historical antecedents. Presenting his approach in the context of a theory of meaning and of conceptual frameworks, Sellars (1953b) maintained that frameworks include propositions that are true *ex vi terminorum* but are not analytic. Unlike Kant, these are not taken to be proven results, but propositions to which we accord a special status when we adopt a framework. This status can be withdrawn if we find reasons for rejecting that framework. But while such a claim is in place, it plays the same methodological roles as Kantian synthetic a priori propositions.

We find a close parallel to Sellars' proposal in Putnam's thesis that there is an unnamed third class of propositions that do not fit into the standard empiricist dichotomy. Putnam was explicitly responding to Quine's rejection of analytic propositions. Since Quine began from the standard empiricist dichotomy, he concluded that there are only synthetic propositions although we do not treat them all the same way when faced with a refutation. Rather, Quine's holism implies that we have a choice about which propositions to reject; we protect some propositions in our corpus on pragmatic grounds such as the degree of disruption that a rejection will generate in our epistemic web. Yet it remains possible that any proposition in the web will be rejected as we continue to adjust our beliefs to our experience. Putnam argued for a different option. He accepted the existence of both analytic and synthetic propositions but maintained that our epistemic practice requires a third class of propositions that are not subject to simple refutation and that guide our research, but that remain open to reconsideration as evidence accumulates.

For Sellars, Putnam, and Quine propositions come to play this special role as a result of decisions scientists make, decisions that can be revised. Much work in epistemology takes it for granted that such decisions are epistemically dubious. With the exception of analytic propositions, only propositions that are established on the basis of evidence and logic possess genuine epistemic legitimacy. Kant fits squarely into this tradition. Yet a central theme that emerged in the second half of the twentieth century is that such choices are unavoidable. Without such choices productive scientific research becomes impossible. As we will see, some take this as a challenge to the epistemic value of science, but an alternative is to take it as challenge to that older view of epistemic acceptability.

Toulmin's "ideals of natural order" provide another version of this central theme. The same holds for at least part of Kuhn's notion of a paradigm. One problem with this notion is that Kuhn included too many items in its scope, but one of these items is the recognition that effective research requires scientists to accept, for a time, propositions that are subject to empirical challenge but are protected against such challenge and that organize research.

We also encounter here one reason why the general approach I am examining has faded from current discussions. The epistemological tradition leads us to seek an algorithm that will dictate when a proposition should be accorded this special status and when rejection is in order. It was a common theme among those who pursued philosophy of science in the new mode that this demand is inappropriate and that communal decisions by the members of the relevant community are all we should expect or require. But proponents of this view did not develop this response in a way that was found to be acceptable by the bulk of the philosophy-of-science community. In particular, its proponents did not develop an acceptable normative account of the grounds for adopting or rejecting and replacing protected propositions in the absence of an algorithm.

Evidence

While logical empiricists engaged in an extensive debate about the exact nature of observation reports, they took it for granted that the evidence we acquire by means of our senses is independent of the theories that we adopt. More recent work has taught us that the process of acquiring evidence is both richer and more complex, but less secure, than had been assumed. There is a central theme that has emerged from these discussions: the body of evidence we acquire is affected by our theories in several ways.

Let us begin with an obvious case: when we are explicitly testing a theory, that theory guides our choice of what evidence to pursue. This may lead researchers to ignore sources of evidence that will be important for other projects, but it serves to focus attention in a way that promotes effective research. To see why suppose that I want to describe the room in which I am working. The task would be overwhelming and not completable in a single lifetime without some reasons for attending to specific items in the room; which items I attend to will depend on my interests at the moment. Now this also applies to evidence collection in science where the pursuit of evidence is directed by specific views of what is worth examining—views that are, in several respects, a function of the theories currently in play, including theories that we are not currently testing. For example, scientists seek more powerful particle accelerators and study details of the cosmic microwave radiation because theories now in play indicate that these are likely to be sources of valuable evidence. Without the appropriate theories, we would have no reason for seeking evidence in these ways, no means of designing the relevant instruments, and no basis for interpreting the results. Nor would we have means of assessing the significance of any items that happened to catch our attention.

These examples underline a central point about science. Scientists are continually engaged in a search for new and more precise evidence about aspects of nature. Currently available theories guide this process. In addition to the examples just mentioned, theories provide the reasons for creating telescopes that gather information from the vast portion of the electromagnetic spectrum that we cannot detect by unaided vision. More recently, scientists have introduced neutrino telescopes that take us beyond the electromagnetic spectrum. We build electron microscopes, scanning microscopes, and other types of microscopes that our ancestors did not imagine. All of these instruments enrich the empirical constraints on theories. As a result, present theories in many fields have faced much more stringent tests than theories in the past. Theory-guided empirical research generates, then, a complex situation: it yields results that are less certain than many have desired and (we

will see) complicates theory-evaluation, but it also enhances the ability to collect evidence that is fundamental to advancing science.

These considerations have led to skepticism among some about the significance of such evidence but this is the wrong response. Theory-dependence does undermine a somewhat naïve view of the nature of empirical evidence, but it does not eliminate the epistemic significance of that evidence. One reason for this becomes clear when we note that the theories guiding evidence collection do not determine the outcome of the procedure. Instruments are designed to interact with aspects of a natural world that exist apart from our theories and that may (and often do) yield outcomes that challenge existing theories. The history of solar-neutrino experiments is a striking recent example (see Bahcall, 1989, Franklin, 2001 for detailed accounts). The first of these experiments was explicitly designed to test the accepted theory of stellar-energy production by measuring the rate at which high-energy electron neutrinos from the sun arrive at the earth. It was proposed by Raymond Davis just a few years after physicists became confident that they could detect neutrinos; the accepted theory of the nature and behavior of neutrinos was central for designing the instrumental complex and interpreting the outcomes. Beginning with the first run of the experiment, the results were consistently much lower than expected. This led to the design of new detectors that were sensitive to a larger energy range, could detect other types of neutrinos, and could gather information that the original detector could not collect. In the late 1990s, after some thirty years of work, the physics community concluded that the problem did not lie with the theory initially under test, but with the theory of neutrinos that was assumed in designing the instrument. Outcomes of this sort are a permanent possibility once we set up an interaction with some part of nature.

Sometimes an anomaly is discovered when scientists engaged in focused research are not explicitly testing a theory. The clashes discovered during the nineteenth century between calculated and observed orbits of Uranus and Mercury are classic examples that were resolved in different ways. Yet the orbits of these planets were anomalous only in light of expectations provided by Newtonian theory. Without this theoretical background the observed orbits could just be entered into a database with no further consideration. This is but one instance in which theories turn relatively mundane observations into important evidence. Given appropriate background, even the failure to detect anything in a particular situation can serve as evidence. If someone removed a chair from my study I would immediately notice this on entering the room. Someone who was unfamiliar with the usual contents of this room might not notice anything of interest. In neutrino physics, since neutrinos are uncharged, the passage of a neutrino through a detector is sometimes recognized because there is no image at a particular place on a photograph.

These reflections indicate that evaluating the import of an empirical result is more difficult than logical empiricists had assumed. The most salient case occurs when an empirical outcome differs from a predicted outcome. As we have already noted, such results indicate that something is wrong somewhere in the set of premises that led to the prediction. It is now clear that we must include our understanding of our instruments in this set. The Michelson-Morely experiment provides an important variation on this theme. As originally conceived, its aim was to compare two different views of how the earth moves through the ether. The outcome of the experiment supported one of these views—the one that Michelson had hoped to refute. That outcome has now been reinterpreted as showing that there is no ether, but the design and interpretation of Michelson's new instrument assumed a wave theory of light which, at the time, made no sense without an ether. Thus our current understanding of the result of the experiment was not contemplated when the experiment was done. The experiment was set up to choose between two theories; both have now been rejected.

This pattern is not unique to this example. When Newton discussed the system of the world in Book III of *Principia* he set up a competition between the Copernican and Brahean views and concluded that both are wrong.² The items constituting our solar system move around the center

of mass of the solar system (which Newton took to be the center of the universe). Still, Newton argued, the Copernican view is a much better approximation than the Brahean view.

We have seen that, even in cases of an empirical outcome that clearly contradicts expectations, there is flexibility in deciding which of the hypotheses used to predict that outcome must be reconsidered. Some, including Quine, have held that this result allows us to maintain any selected hypothesis no matter what the evidence. Others have gone further arguing that this flexibility undermines the epistemic force of empirical evidence. But there are two points that must be noted about both theses. First, it acknowledges that an empirical outcome at variance with a prediction requires that something be changed somewhere in our corpus. As long as we are doing science, the outcome cannot be ignored. Second, as Greenwood (1990) noted, any change we make is liable to have further empirical consequences, subject to further tests which may not turn out as predicted. In the solar neutrino example the outcome has led to a new understanding of neutrinos with new predictions that are currently being tested. While we may never reach a situation in which evidence and methodology necessitate a specific choice, we do not have complete freedom to maintain any hypothesis we wish to. It is, moreover, worth repeating that while the pervasive role of theories in evidence collection makes empirical results less certain than was once believed, these theories have expanded our ability to interact with nature in new ways and thus enriched the constraints on our theories. The results of interactions with nature continue to provide constraints on future research even when interpretations of these results are changed.

Now let's consider an additional way in which theory-guidance supports research. Since William Herschel discovered infra-red solar radiation in 1800 we have learned that much of the physical world cannot be detected by our senses alone. Yet we have also found that items outside the range of our senses—items such as radioactivity and genes—play a major role in determining how nature behaves. As a result, recent science has made more progress on several dimensions (which I discuss below) than it did in the previous millennia of studying items that we can easily detect. But theory-guided research provides our only means of access to this central research realm.

There is another, more subtle, respect in which evidence is dependent on theories. For an empirical result to be relevant to evaluating a particular theory it must be described using the concepts of that theory. Sometimes competing theories yield different descriptions of a body of evidence. The significance of this point will be clearer after we have discussed conceptual change.

Conceptual change

Conceptual change was a central topic throughout the period under review. Discussion of this topic raised some pseudo-problems that I will discuss, but also taught us important lessons about how science develops. The key issue concerns the introduction of new concepts that cannot be completely reduced to older concepts. This issue could not arise for the early logical empiricists who held that a set of basic concepts derived directly from experience provides the empirical content for all scientific concepts. Any new concepts would be new combinations of these basic concepts and thus could be translated without loss of content into basic concepts, eliminating any apparent conceptual disparities between different theories. This thesis failed because it could not account for the way that scientific concepts embody more than just summaries of the evidence used to recognize instances of a concept. The logical empiricist response was a series of retreats in which the relation between experiential concepts and theoretical concepts became progressively more tenuous and indirect. Some versions of the later approach acknowledged that part of the content of theoretical concepts derives from systemic relations among the concepts in a theory. This, in turn, provided an entering wedge for maintaining that concepts derive their content solely from these systemic relations and that theory imposes meaning on experience (see Brown, 1979: ch 3 and Brown, 2007 for details). I will approach the issues that arise by first noting three forms of conceptual change.

One form is fairly straightforward: elimination of concepts that were once part of science along with the terminology associated with those concepts. Well-known examples include phlogiston, caloric, and the Aristotelian notion of natural place, but the extent of the phenomenon is somewhat hidden because many rejected concepts are familiar only to historians of particular periods. I will add two lesser-known examples. One of these is *telegony*: the presumed effect of the father of a female's first child on all her subsequent children. In nineteenth-century England, practical animal breeders and more theoretical biologists took this to be a well-documented phenomenon. Darwin, for example, offered the ability to explain telegony as one virtue of his pangenesis theory of heredity.

The second example had a shorter life-span. In 1899 the Curies discovered cases in which a non-radioactive material placed near a radioactive material acquires a temporary radioactivity. They interpreted this as a case of *radioactive induction* by analogy with electromagnetic induction. This interpretation was widely accepted until 1903 when Rutherford and Soddy, working together, discovered that the non-radioactive material did not become radioactive, but was contaminated by a radioactive substance. While the concept of radioactive induction had no long-lasting impact, it illustrates an important feature of research at the boundaries of current science: researchers try out new concepts, many of which are soon abandoned. Historians can, no doubt, provide further examples in their fields of expertise.

A second, straightforward form of conceptual change occurs when we introduce new concepts that were not contemplated by our ancestors, along with new terminology. These arise because of genuinely new discoveries. Clear examples include entropy, fermion, and reverse transcriptase.

The trickiest kind of conceptual innovation occurs when the content of a new concept overlaps that of an older concept while the older terminology is retained. Consider *planet*. Before Copernicus there were two parts to the content of this concept. First, five celestial objects—Mercury, Venus, Mars, Jupiter, and Saturn—were designated planets; sometimes the Moon and Sun were included. Second, planets were picked out because they appeared to move around the stationary earth in non-circular annual paths. The Earth was, by definition, not a planet. Once the Copernican revolution had been consolidated planets were defined as items that (to a good approximation) move around the Sun; the Earth was just another planet. Both the defining characteristics of the planets and the status of a key item in the earlier conceptualization of planets changed. Moreover, the Sun and Moon were now clearly excluded from the class of planets. Yet there is continuity with the older framework since the five major planets retained their status.

Perhaps the most discussed example is *mass* in the transition from classical mechanics to relativity. Classically mass was considered invariant with respect to velocity. Newton introduced mass in creating a new contrast with *weight*; he did not contemplate the possibility that mass might depend on velocity, as it does in relativity theory. There is continuity with the older view captured in the new invariant *rest mass*, the fact that mass essentially reduces to rest mass in many low-velocity situations, and in the fact that mass retains its function as a measure of resistance to acceleration. But a distinction between mass and rest mass makes no sense in classical mechanics; neither does the relation between mass and energy embodied in the new equation $E=mc^2$. This combination of conceptual continuity and conceptual innovation has kept this example at the focus of discussion.

The speed of light (in *vacuo*), c , provides a revealing example that has not been much discussed. In classical mechanics this speed has no special significance; in relativity this speed takes on a variety of new and fundamental roles. The most important is that c is invariant across observational frames while all other speeds are frame relative. This requires revision of the formula for compounding velocities, modification of the rules for transforming laws between uniformly moving frameworks, a new formula for the Doppler effect, and other changes. Indeed, c appears in just about every major formula of the new theory. There is, again, overlap with the older framework: the magnitude of c is not changed and its special role often has practical significance only at high

velocities; otherwise the classical formulas remain satisfactory. Conceptual innovations, especially of this third type, generated some widely discussed issues.

First, there seems to be a distinction between new discoveries that fit an existing conceptual framework and discoveries that yield framework changes. For example, reclassifying the Earth as a planet was a fundamental change, the discovery of new planets does not disrupt the post-Newtonian framework; it need not have disrupted the older framework. In a familiar way of thinking, this situation calls for a rule that will allow us to distinguish the two cases; no such rule has been forthcoming. While some view the failure to find an appropriate rule as a major failing of those who discussed conceptual innovation, there may be a different, more fundamental, lesson that we should learn: that the search for a rule is the wrong way to proceed. Rather, we have an array of cases in which more and less drastic changes occur in our way of thinking about some topic. There are many clear cases of fundamental change, and of new discoveries that do not alter the framework, but no sharp line can be drawn between these. Many will respond that without a principled division between the two kinds of cases we are left with arbitrary choices and intellectual chaos. I will challenge this view as we proceed and argue that rule-driven choices and arbitrary decisions are not the only options. But first I want to introduce some further issues raised by the focus on conceptual innovation.

Traditional empiricists thought about concepts individually. Each of the basic concepts is independent of any other concept while the more complex concepts that we introduce can be individually reduced to basic concepts. I have already noted that, in later discussions, many logical empiricists moved away from this view, acknowledging some role for systemic relations in determining conceptual content. This view is also central to Sellars' work and is embodied in Putnam's notion of a law-cluster concept. Feyerabend and Kuhn, at times, took this view to an extreme by arguing that *only* systemic relations are involved in determining conceptual content. But the introduction of systemic relations into conceptual change generated two further problems: the demand for criteria that determine what constitutes a single conceptual system and prevents our lapsing into an all-consuming holism, and the problem of incommensurability. I will consider the latter next.

Incommensurability was the most widely debated issue raised by the thesis that scientific revolutions involve deep conceptual change. The status of empirical evidence now became a central topic. As noted above, for a body of evidence to be relevant to the evaluation of a theory, that evidence must be described using the concepts of that theory. But if conceptual content is completely determined by internal relations among the concepts of a theory, then different theories will embody different conceptual systems. As a result, it seems, there will not be any body of data that is relevant to the evaluation of both theories. Genuine comparison becomes impossible and even the sense in which two theories are in competition becomes unclear.

Two steps are needed to get beyond this worry. Note, first, that part of the problem arises from extension of a useful metaphor—the idea that a scientific theory is a language—beyond the range in which it is helpful. Scientific theories embody concepts and a technical vocabulary that, to some degree, constrain the way practitioners think about their subject. Part of the process of learning a theory is to learn its language. But scientific theories exist in a wider culture and language community that is typically shared by advocates of competing theories; this provides resources on which disputants can draw. Often these resources allow creative thinkers to describe relevant situations in ways that abstract from the concepts of their preferred theory. An experiment that Galileo proposed is a good example. The behavior of an object dropped from a high tower provided an important argument against the earth's daily rotation. From the perspective of Aristotelian physics a dropped object falls straight down towards the center of the earth. Thus, it was argued, if the earth is turning as the object falls, it will not land at the foot of the tower but a considerable distance away.³ Since all agreed that the object falls at the foot of the tower, this seemed to refute the rotating-earth thesis. Galileo recognized that, for this and other reasons, a moving earth is not compatible with

Aristotelian physics, and sought an alternative. On his account the falling object shares the motion of the earth and thus would indeed land at the foot of the tower. Since the tower experiment cannot provide a reason for preferring one of the theories, Galileo proposed a related experiment: drop an object from the top of the mast of a moving ship. On the Aristotelian account the object would land towards the rear of the ship; on Galileo's account it would land at the foot of the mast. In both cases, where the object lands is described in language that is independent of the competing theories. On the ship, the assessment of where the object lands can be made by an independent referee who has no knowledge of either theory. Moreover, confirmation of Galileo's prediction could provide a motivation for trying to understand how he arrived at it. We encountered a similar situation when discussing interpretations of the Michelson-Morley experiment where there was a visible outcome on which all agreed.

Sometimes the "theory as language" metaphor was enhanced by the odd view that people are epistemically trapped by their language and unable to think beyond its boundaries. This was held even in the face of familiar evidence of bilingual people as well as examples of scientists who have mastered multiple theories teaching, say, classical mechanics but doing research in quantum theory.⁴ Major theoretical innovators were typically masters of the theories they sought to replace. Galileo deployed such knowledge when he proposed the ship experiment.

We come, then, to the second step mentioned above. Logical empiricists focused on relations between explicitly formulated theories and observation reports; any consideration of the cognitive process by which researchers produce these was considered, as a matter of principle, irrelevant. But researchers are people who come into research situations with a rich body of cognitive skills; without these skills no research would be possible. These skills provide the key to dealing with many of the problems that bedeviled both logical empiricism and the proposed successors. Kuhn eventually took this step in the case of incommensurability, acknowledging that people can learn an unfamiliar framework, although this may require significant effort (2000: 220).

We must, then, take human cognitive abilities into account if we are to understand how science works and arrive at a proper evaluation of its epistemic import. While this is as great a departure from the perspective of logical empiricism as we are likely to find, some reflection just on deduction will underline the limits of the logical empiricist approach. For the fact that a given set of premises entails a conclusion does not make that relation part of the body of human knowledge. This can occur only if some of us *recognize* that relation. Put differently, whether an entailment relation holds between propositions is an objective fact, but only objective facts that we are aware of become a part of human knowledge. In mathematics, where deduction reigns supreme, we require proofs. Proofs do not create logical relations, they demonstrate to human beings that such relations obtain. Moreover, proofs must be tailored to human cognitive capacities. We begin with some ability to recognize simple logical relations and we learn about more complex relations through a process that links premises and conclusions by means of simpler steps. Often a proof is difficult and an entailment relation may resist the best efforts of mathematicians for generations or even centuries. The recent proof of Fermat's Last Theorem some three hundred years after the theorem was announced, demonstrated (to the few who could follow it) that a particular set of premises entails that theorem. What we cannot demand is that human beings should just recognize an entailment that obtains. This is beyond our capability.

Once this point of principle has been made, it becomes clear that the actual process of arriving at and evaluating scientific proposals depends on our abilities at every stage. Indeed, science has been improved and enriched as we have learned more about our abilities—and limitations. Consider two examples. First, double-blind testing has a history (Kaptchuk, 1998); it only became part of the methodology of certain fields as researchers became aware of human limitations that were not always apparent. Second, while our senses allow us to detect only a limited range of items in the physical world, this has not stopped us from conceiving of the existence of items we cannot sense

and learning to interact with and thereby learn about them. Different beings, with different abilities, might not be able to do this.

Reflecting on our cognitive abilities puts us in a position to move beyond the view that human psychology provides an impediment on the path to knowledge and should be excluded as rigorously as possible. Our cognitive abilities are the source of our epistemic strengths. To be sure, as cognitive psychologists have been pointing out since the 1970s, there are many normative failings in our cognitive behavior. But this very project is possible only because some among us have discovered the appropriate norms against which we can evaluate the behaviors in question. It is a serious error to view human psychology only as a source of epistemic failings.

With these considerations in mind, let us return to language and incommensurability. Obviously the ability to use language is among the basic human capacities. Children learn their local language without effort; children raised in a multilingual community learn multiple languages. People who encounter others with a different language rapidly produce a pidgin that will, if the interaction is maintained, develop into a full-blown language (a creole) within a generation or two. Even adults can learn another language although—as is the case with all human achievements—some are better at this than others.

We also have the ability to invent and learn new concepts that cannot be reduced to older concepts, and this is crucial for the advancement of science. A significant amount of scientific research consists of inventing and trying out new concepts because we discover theoretical failings in older frameworks, and because our interactions with nature bring us face-to-face with phenomena not previously imagined. While the need to introduce new concepts complicates the development of science, our ability to do this is a vital resource without which science could not exist; it is not a bar to the coherent development of science.

The same considerations allow us to deal with the threat of rampant holism. At a given time, the connections that are implicit in a conceptual system do not overwhelm us because researchers focus their attention more narrowly, and do not draw out these implications—although the discovery of connections between subjects that were not previously viewed as related is sometimes important for the development of science. Another classic problem is amenable to the same approach. There have been theories, such as Bohr's theory of the atom, that were inconsistent, and recognized as inconsistent. A basic result in deductive logic tells us that every proposition is implied by a contradiction, but working scientists did not use the theory in this way. Rather, they limited the scope of the conclusions that they derived from the theory, even while recognizing that inconsistency is a defect and seeking a replacement.

While I have been concerned here with conceptual incommensurability, the literature includes a second form of incommensurability: competing theories may embody different criteria for evaluating theories. This is to be expected given the aspects of methodology already discussed. The status of causality as a central norm for science provides a striking example. A long tradition considers the search for causes as the central function of science and thus views quantum theory—on the dominant interpretation—as unacceptable. Others respond by rejecting this norm. This is not a unique situation. Taken in its historical context, it is no more extreme than the elimination of teleology from physics and then from biology, or the introduction of the systematic study of the part of the world that we cannot sense. Reevaluations of basic methodology are part of the ongoing development of science and are well within the range of our cognitive capacities.

Let me summarize the key outcomes of this section. The first is to acknowledge that conceptual innovation (along with methodological innovation) is a central feature of the process of scientific discovery. We have learned things about nature (including ourselves) that our ancestors did not imagine. Our psychology—our cognitive abilities—provide the key resources that enable this kind of science. While there are pitfalls that come along with our psychology, we have also (through scientific research) been learning about these pitfalls and ways to overcome them. Ignoring the role

of human cognition in the pursuit of science only blocks the road to an adequate understanding of science and to the goal of learning how to better pursue scientific discovery.⁵

Social resources

It should be obvious that science is a social endeavor. The development of science takes place over multiple generations; researchers at each stage build on work of their predecessors, even when this results in correcting earlier work. At a given time, information and skills are distributed across multiple individuals and productive research often requires marshaling and organizing that information and those skills. In some fields, such as high-energy physics, there is a fairly sharp distinction between theoreticians and experimentalists that is a direct result of the need for large amounts of specialized knowledge. Some contemporary experimentation requires very large teams. In some cases the need for large teams is generated by the sheer quantity of work involved. The first experimental paper reporting the lifetime of charm particles took some 280 person-years and had 99 authors (Hardwig, 1991). One of the first papers reporting evidence for the existence of top quarks had 500 authors; these included people of diverse skills. We draw on social resources when we look up values in a handbook or use off-the-shelf hardware or software. In addition, individuals are often too enamored of their own ideas to notice problems that will be apparent to others.

In spite of the clear social basis for science, epistemology has been dominated by an individualist view that considers the pursuit of knowledge to be solely a matter of what individual thinkers can establish within the privacy of their own minds. When Kuhn included a social element in his original characterization of a paradigm (1962: 10) many philosophers responded that this was mere sociology—even advocacy of mob rule—not epistemology. For logical empiricists, in particular, any consideration of social factors is irrelevant to philosophy of science; intrusion of social factors into research is an impediment to the proper development of science.

A substantial number of sociologists agreed with the last conclusion and took the social basis of science as a reason for challenging the epistemic credentials of much science—not including their own. But two key items were lost in the resulting “science wars.” First, the social side of science is only one feature of the process by which we pursue knowledge of the world. It must be balanced by, and integrated with, other features we have noted including logic, evidence, and skills. Overemphasis on just one of a set of elements that are all required for successful science is a theme that runs through many of the debates we have been considering.

Second, there was a failure to understand that the social basis of science provides crucial resources without which research could not proceed. Here too there are pitfalls that come with these resources, but we do not avoid such pitfalls by ignoring them. In addition to the resources noted in the opening paragraph of this section, I will add that the expanded communication that became available in recent decades provides a major enhancement of the ability to pursue science (along with important dangers). Scientists now have easy access to a much wider body of evidence and information as well as expanded sources of analysis and critique of their theories and their evidence-gathering procedures. Again we find that the pursuit of science is more complex and less certain than had previously been conceived, and that a feature of research, as carried out by human scientists, that worried some philosophers of science, actually plays a positive role in the pursuit of science.

Progress

The nature, and even the existence, of scientific progress came to special prominence in the period I am surveying. To a large extent doubts about progress resulted from undermining earlier, overly optimistic, views of the development of science. For example, given the many cases in which once-accepted theories are rejected as false and replaced by new theories that are later found to be false,

we cannot simply define progress in terms of accumulation of truths.⁶ The appropriate response is to seek a more modest and nuanced view of the accomplishments and prospects for progress while recognizing that the appropriate criteria for judging progress are different in different domains. I will explore some examples.

Technology provides a rich array of cases in which there are clear examples of improvement as measured by explicit criteria. Once a new technology has been invented we have, it seems, a powerful ability to improve it. For example, over time we have learned to produce cars that are more reliable and use less fuel than older cars; computers have become faster, smaller, and more powerful; surgery has become safer and more effective than it was in the past. Some technological improvements impinge directly on the development of theoretical science. As we discovered that the world is replete with items we cannot detect with our unaided senses, we have learned to design instruments that allow us to interact with those objects and increase the body of evidence we can use in testing theories. We have also learned to build instruments that increase the precision of our evidence collection even in the perceptible realm. Both of these abilities have been enhanced by increased computer power; in some fields modern experimentation would not be possible without this computer power. Technologies have thus improved our ability to pursue truth, even if we lack any guarantee that we are achieving or approaching truth in many fields.

We also find clear progress in our ability to predict outcomes of experiments, what we will find from observations, and the results of our practical activities. Physicists, for example, were able to predict the result of bringing together a critical mass of U235, rather than finding this out by trial and error. In theoretical science improved predictions provide one motivation for developing new means of testing theories. Developments in mathematics provide a major source of this increased predictive power. A contemporary student with a semester of calculus can easily solve problems that eluded some of the best minds in the world in the seventeenth century. Such developments have extended our ability to make and evaluate predictions even in fields that are not engaged in constructing mathematical theories.

The accomplishments noted thus far contribute to the development of effective research. This includes research that generates a challenge to the very theory that guided that research. It took a great deal of mathematical and technical development to bring us to the point at which an anomaly in the shift of the Mercury's perihelion of twenty-two seconds of arc per century could pose a serious challenge to the most successful scientific theory that had yet existed.

The hardest problem about progress is our ability to assess whether we are learning what nature is like *in some fields*. There are fields in which there is no reasonable doubt that we are achieving such knowledge. To mention just a few examples, we know more than our ancestors did about the various land masses and cultures on our planet, about human anatomy and physiology, and about the variety of items in the cosmos. The hard question arises in those areas of mathematical physics that are widely considered to be the most fundamental and the most successful in terms of predictive ability and support for new technologies. In the key case of quantum theory we have a formalism that we know how to associate with measurements and predictions, and that provided the basis for designing transistors—and thus all of the technologies that they enabled—as well as MRI machines and more. Yet when we attempt to decide what nature is like at the quantum level we find only disagreement and paradox. Here it is not clear if we know more than our ancestors did or even if we know much at all.

Is science rational?

Many debates in the latter half of the twentieth century were presented as evaluations of the rationality of science. I have avoided that terminology here because I think it generated more confusion than insight. One set of problems arises because “rational” and its cognates are used to describe too many different things. Much of the work I have been examining focused on the process of science. Yet here too we must be careful because we sometimes describe a process as rational if it has

certain characteristics even though the actors were unaware of this. One can, on this usage, stumble mindlessly into being rational.

Even if we focus more narrowly on the considerations scientists, in their historical context, invoke when making such decisions as whether to accept a theory as part of their research framework, or view it as worthy of further testing, or discard it, the language of rationality carries confusing baggage. The traditional identification of “the rational” with “the logical” is especially prominent. Kant, we noted, sought to defend the rationality of science by introducing a new form of logic. Logical empiricists did not accept this move, and they recognized only a limited scope for rationality. Carnap (1956), for example, maintained that rational decisions occur only within a framework and that framework choice is not rational. This much-admired view provides one reason why the thesis that the history of science exhibits changes of framework was taken, by many, as equivalent to claiming that science is not fully rational. It has also been widely taken for granted that “the rational” and “the social” are diametrically opposed while we have seen that social factors are of crucial importance for the advancement of science. Rather than entering into these, often verbal, debates, I have avoided this terminology in my account of things we should have learned. I submit that this terminological decision has not prevented me from saying anything of substance.

Conclusion

Given that the newer work challenged well-established views in philosophy of science, it is no surprise that one immediate response was to read it as an attack on the epistemic value of science itself. Time and reflection have led to a different evaluation. It is, again, now clear that scientific research is more complex, and its results less certain, than had been thought. But once this is absorbed, many items that were initially seen as a threat to science can be recognized as vital resources for the development of science—although these resources have pitfalls and must be used with care. These include the decision to maintain selected views as free from challenge for substantial periods of time. This produces carefully focused research—research that may itself lead to reconsideration of protected theses. They also include theory-guided research which has increased the scope and power of science; the ability to introduce and absorb new concepts and thus new ways of thinking about a subject; and the ability to deploy resources that are spread across members of the community. And all of these are mediated by the various skills that individuals bring to their research.

Finally, I want to underline a theme that is implicit in the above discussion: science is an ongoing quest that takes time, often long periods of time. As research proceeds, surprises are the norm. It is a key epistemic virtue of the sciences that they continually seek out situations that may yield surprises—even surprises that lead to changes in the aims and methods of science. Many philosophers seek timeless rules that will dictate proper scientific procedure. Perhaps such rules exist, but finding them is a task at least as difficult as finding results within the sciences that will never be subject to challenge. It would be nice to find a methodology that will allow us to avoid all error, but this is beyond human capacity. Instead, science has developed an exceptionally powerful means of recognizing and correcting errors. The ability to do this is a central feature of the genius of science.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Notes

1. That such generalizations are refuted by items that are A and not-B was recognized as non-problematic.
2. Brahe held that the earth is stationary, the planets move around the sun, and the entire complex moves

around the earth. When Newton addressed the subject, the Ptolemaic view was effectively dead.

3. Using contemporary figures for a three-hundred-foot tower at the latitude of Pisa, the distance would be a little under 9/10 mile. At one point Galileo gives a time of fall that assumes a slower rate of fall and gives a distance of a little over 9/10 mile.
4. Kuhn began stressing this point circa 1983. See Kuhn (2000: 53, 77, 238).
5. Seventeenth and eighteenth-century empiricists did not share the view that our psychology is irrelevant to understanding human knowledge, although they mainly viewed our psychology as a source of limitations on the scope of our knowledge. Systematic rejection of any role for psychology was a feature of analytic epistemology as it developed in the twentieth century. See Price (1940) for a systematic attempt to eliminate any psychological elements from Hume's epistemology.
6. It is, I hope, obvious that we cannot define progress as getting closer to the truth in a domain—a truth that we do not know, and that may require the introduction of concepts that we have not yet imagined.

References

- Bahcall J (1989) *Neutrino Astrophysics*. Cambridge: Cambridge UP.
- Brown HI (1979) *Perception, Theory, and Commitment*. Chicago: University of Chicago Press.
- Brown HI (2007) *Conceptual Systems*. Oxford/New York: Routledge.
- Carnap R (1956) Empiricism, semantics, and ontology. In Carnap, *Meaning and Necessity*. Chicago: University of Chicago Press, pp. 205–221.
- Feyerabend P (1962) Explanation, reduction, and empiricism. In H Feigl and G Maxwell (eds), *Minnesota Studies in the Philosophy of Science* III. Minneapolis: University of Minnesota Press, pp. 28–97.
- Franklin A (2001) *Are There Really Neutrinos? An Evidential History*. Cambridge, MA: Perseus.
- Goodman N (1955) *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard UP.
- Greenwood J (1990) Two dogmas of neo-empiricism. *Philosophy of Science* 57: 553–574.
- Gutting G (2009) *What Philosophers Know*. Cambridge: Cambridge UP.
- Hanson NR (1958) *Patterns of Discovery*. Cambridge: Cambridge UP.
- Hardwig J (1991) The role of trust in knowledge. *Journal of Philosophy* 88: 693–708.
- Hempel C (1945) Studies in the logic of confirmation. *Mind* 54: 1–26, 97–121.
- Kaptchuk T (1998) Intentional ignorance: a history of blind assessment and placebo controls in medicine. *Bulletin of the History of Medicine* 72: 389–433.
- Kuhn TS (1962) *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Kuhn TS (2000) *The Road Since Structure*. Chicago: University of Chicago Press.
- Polanyi M (1958) *Personal Knowledge*. Chicago: University of Chicago Press.
- Price, HH (1940) The permanent significance of Hume's philosophy. *Philosophy* 15: 7–37.
- Putnam H (1962) The analytic and the synthetic. In H Feigl and G Maxwell (eds), *Minnesota Studies in the Philosophy of Science* III. Minneapolis: University of Minnesota Press, pp. 358–397.
- Quine WV (1951) Two dogmas of empiricism. *Philosophical Review* 60: 20–43.
- Sellars W (1948) Concepts as involving laws and inconceivable without them. *Philosophy of Science* 15: 287–315.
- Sellars W (1953a) Inference and meaning. *Mind* 62: 313–338.
- Sellars W (1953b) Is there a synthetic a priori? *Philosophy of Science* 20: 121–138.
- Sellars W (1954) Some reflections on language games. *Philosophy of Science* 21: 204–228.
- Toulmin S (1961) *Foresight and Understanding*. New York: Harper & Row.