

Underdetermination of Scientific Theory

First published Wed Aug 12, 2009; substantive revision Mon Sep 16, 2013

At the heart of the underdetermination of scientific theory by evidence is the simple idea that the evidence available to us at a given time may be insufficient to determine what beliefs we should hold in response to it. In a textbook example, if all I know is that you spent \$10 on apples and oranges and that apples cost \$1 while oranges cost \$2, then I know that you did not buy six oranges, but I do not know whether you bought one orange and eight apples, two oranges and six apples, and so on. A simple scientific example can be found in the rationale behind the sensible methodological adage that “correlation does not imply causation”. If watching lots of cartoons causes children to be more violent in their playground behavior, then we should (barring complications) expect to find a correlation between levels of cartoon viewing and violent playground behavior. But that is *also* what we would expect to find if children who are prone to violence tend to enjoy and seek out cartoons more than other children, or if propensities to violence and increased cartoon viewing are both caused by some third factor (like general parental neglect or excessive consumption of Twinkies). So a high correlation between cartoon viewing and violent playground behavior is evidence that (by itself) simply *underdetermines* what we should believe about the causal relationship between the two. But it turns out that this simple and familiar predicament only scratches the surface of the various ways in which problems of underdetermination can arise in the course of scientific investigation.

- [1. A First Look: Duhem, Quine, and the Problems of Underdetermination](#)
- [2. Holist Underdetermination and Challenges to Scientific Rationality](#)
 - [2.1 Holist Underdetermination: The Very Idea](#)
 - [2.2 Challenging the Rationality of Science](#)
- [3. Contrastive Underdetermination, Empirical Equivalents, and Unconceived Alternatives](#)
 - [3.1 Contrastive Underdetermination: Back to Duhem](#)
 - [3.2 Empirically Equivalent Theories](#)
 - [3.3 Unconceived Alternatives and A New Induction](#)
- [Bibliography](#)
- [Academic Tools](#)
- [Other Internet Resources](#)
- [Related Entries](#)

1. A First Look: Duhem, Quine, and the Problems of Underdetermination

The scope of the epistemic challenge arising from underdetermination is not limited only to scientific contexts, as is perhaps most readily seen in classical skeptical attacks on our knowledge more generally. René Descartes ([1640] 1996) famously sought to doubt any and all of his beliefs which could possibly be doubted by supposing that there might be an all-powerful Evil Demon who sought only to deceive him. Descartes' challenge essentially appeals to a form of underdetermination: he notes that all our sensory experiences would be just the same if they were caused by this Evil Demon rather than an external world of tables and chairs. Likewise, Nelson Goodman's (1955) “New Riddle of Induction” turns on the idea that the evidence we now have could equally well be taken to support inductive generalizations quite different from those we usually take them to support, with radically different consequences for the course of future events. ^[1] Nonetheless, underdetermination has been thought to arise in scientific contexts in a variety of distinctive and important ways that do not simply recreate such radically skeptical possibilities.

John Stuart Mill articulated a distinctively scientific version of the concern with impressive clarity in *A System of Logic*, where he writes:

Most thinkers of any degree of sobriety allow, that an hypothesis...is not to be received as probably true because it accounts for all the known phenomena, since this is a condition sometimes fulfilled tolerably well by two conflicting hypotheses...while there are probably a thousand more which are equally possible, but which, for want of anything analogous in our experience, our minds are unfitted to conceive. ([1867] 1900, 328)

However, the traditional locus classicus for underdetermination in science is the work of Pierre Duhem, a French physicist as well as historian and philosopher of science who lived at the turn of the 20th Century. In *The Aim and Structure of Physical Theory*, Duhem formulated various problems of scientific underdetermination in an especially perspicuous and compelling way, although he himself argued that these problems posed serious challenges only to our efforts to confirm theories in physics. In the middle of the 20th Century, W. V. O. Quine suggested that such challenges applied not only to the confirmation of all types of scientific theories, but to all knowledge claims whatsoever, and his incorporation and further development of these problems as part of a general account of human knowledge was one of the most significant developments of 20th Century epistemology. But neither Duhem nor Quine was careful to systematically distinguish a number of fundamentally distinct lines of thinking about underdetermination that may be discerned in their works. Perhaps the most important division is between what we might call holist and contrastive forms of underdetermination. Holist underdetermination (Section 2 below) arises whenever our inability to test hypotheses in isolation leaves us underdetermined in our *response* to a failed prediction or some other piece of disconfirming evidence. That is, because hypotheses have empirical implications or consequences only when *conjoined* with other hypotheses and/or background beliefs about the world, a failed prediction or falsified empirical consequence typically leaves open to us the possibility of blaming and abandoning one of these background beliefs and/or ‘auxiliary’ hypotheses rather than the hypothesis we set out to test in the first place. But contrastive underdetermination (Section 3 below) involves the quite different possibility that for any body of evidence confirming a theory, there might well be *other theories* that are also well confirmed by that very same body of evidence. Moreover, claims of underdetermination of either of these two fundamental varieties can vary in strength and character in any number of further ways: one might, for example, suggest that the choice between two theories or two ways of revising our beliefs is *transiently* underdetermined simply by the evidence we happen to have *at present*, or instead *permanently* underdetermined by *all possible* evidence. Indeed, the variety of forms of underdetermination that have been suggested to confront scientific inquiry, and the causes and consequences claimed for these different varieties, are sufficiently heterogeneous that attempts to address “the” problem of underdetermination for scientific theories have often engendered considerable confusion and argumentation at cross-purposes.^[2]

Moreover, such differences in the character and strength of various claims of underdetermination turn out to be crucial for resolving the significance of the issue. For example, in some recently influential discussions of science it has become commonplace for scholars in a wide variety of academic disciplines to make casual appeal to claims of underdetermination (especially of the holist variety) to support the idea that something *besides* evidence must step in to do the further work of determining beliefs and/or changes of belief in scientific contexts: perhaps most prominent among these are adherents of the sociology of scientific knowledge (SSK) movement and some feminist science critics who have argued that it is typically the sociopolitical interests and/or pursuit of power and influence by scientists themselves which play a crucial and even decisive role in determining which beliefs are actually abandoned or retained in response to conflicting evidence. As we will see in Section 2.2, however, Larry Laudan has argued that such claims depend upon simple equivocation between the comparatively weak or trivial forms of underdetermination that their partisans have managed to establish and the far stronger forms from which they draw radical conclusions about the limited reach of evidence and rationality in science. In the sections that follow we will seek to clearly characterize and distinguish the various forms of both holist and contrastive underdetermination that have been suggested to arise in scientific contexts (noting some important connections between them along the way), assess the strength and significance of the heterogeneous argumentative considerations offered in support of and against them, and consider just which forms of underdetermination pose genuinely consequential challenges for scientific inquiry.

2. Holist Underdetermination and Challenges to Scientific Rationality

2.1 Holist Underdetermination: The Very Idea

Duhem's original case for holist underdetermination is, perhaps unsurprisingly, intimately bound up with his arguments for confirmational holism: the claim that theories or hypotheses can only be subjected to empirical testing in groups or collections, never in isolation. The idea here is that a single scientific hypothesis does not by itself carry any implications about what we should expect to observe in nature; rather, we can derive empirical consequences from an hypothesis only when it is conjoined with many other beliefs and hypotheses, including background assumptions about the world, beliefs about how measuring instruments operate, further hypotheses about the interactions between objects in the original hypothesis' field of study and the surrounding environment, etc. For this reason, Duhem argues, when an empirical prediction turns out to be falsified, we do not know whether the fault lies with the hypothesis we originally sought to test or with one of the many other beliefs and hypotheses that were also needed and used to generate the failed prediction:

A physicist decides to demonstrate the inaccuracy of a proposition; in order to deduce from this proposition the prediction of a phenomenon and institute the experiment which is to show whether this phenomenon is or is not produced, in order to interpret the results of this experiment and establish that the predicted phenomenon is not produced, he does not confine himself to making use of the proposition in question; he makes use also of a whole group of theories accepted by him as beyond dispute. The prediction of the phenomenon, whose nonproduction is to cut off debate, does not derive from the proposition challenged if taken by itself, but from the proposition at issue joined to that whole group of theories; if the predicted phenomenon is not produced, the only thing the experiment teaches us is that among the propositions used to predict the phenomenon and to establish whether it would be produced, there is at least one error; but where this error lies is just what it does not tell us. ([1914] 1954, 185)

Duhem supports this claim with examples from physical theory, including one designed to illustrate a celebrated further consequence he draws from it. Holist underdetermination ensures, Duhem argues, that there cannot be any such thing as a “crucial experiment”: a single experiment whose outcome is predicted differently by two competing theories and which therefore serves to definitively confirm one and refute the other. For example, in a famous scientific episode intended to resolve the ongoing heated battle between partisans of the theory that light consists of a stream of particles moving at extremely high speed (the particle or “emission” theory of light) and defenders of the view that light consists instead of waves propagated through a mechanical medium (the wave theory), the physicist Foucault designed an apparatus to test the two theories' competing claims about the speed of transmission of light in different media: the particle theory implied that light would travel faster in water than in air, while the wave theory implied that the reverse was true. Although the outcome of the experiment was taken to show that light travels faster in air than in water, ^[3] Duhem argues that this is far from a refutation of the hypothesis of emission:

in fact, what the experiment declares stained with error is the whole group of propositions accepted by Newton, and after him by Laplace and Biot, that is, the whole theory from which we deduce the relation between the index of refraction and the velocity of light in various media. But in condemning this system as a whole by declaring it stained with error, the experiment does not tell us where the error lies. Is it in the fundamental hypothesis that light consists in projectiles thrown out with great speed by luminous bodies? Is it in some other assumption concerning the actions experienced by light corpuscles due to the media in which they move? We know nothing about that. It would be rash to believe, as Arago seems to have thought, that Foucault's experiment condemns once and for all the very hypothesis of emission, i.e., the assimilation of a ray of light to a swarm of projectiles. If physicists had attached some value to this task, they would undoubtedly have succeeded in founding on this assumption a system of optics that would agree with Foucault's experiment. ([1914] 1954, p. 187)

From this and similar examples, Duhem drew the quite general conclusion that our response to the experimental or observational falsification of a theory is always underdetermined in this way. When the world does not live up to our theory-grounded expectations, we must give up *something*, but because no hypothesis is ever tested in isolation, no experiment ever tells us precisely which belief it is that we must revise or give up as mistaken:

In sum, the physicist can never subject an isolated hypothesis to experimental test, but only a whole group of hypotheses; when the experiment is in disagreement with his predictions, what he learns is that at least one of the hypotheses constituting this group is unacceptable and ought to be modified; but the experiment does not designate which one should be changed. ([1914] 1954, 187)

The predicament Duhem here identifies is no rainy day puzzle for philosophers of science, but a methodological challenge that constantly arises in the course of scientific practice itself. It is simply not true that for practical purposes and in concrete contexts a single revision of our beliefs in response to disconfirming evidence is always obviously correct, or the most promising, or the only or even most sensible avenue to pursue. To cite a classic example, when Newton's celestial mechanics failed to correctly predict the orbit of Uranus, scientists at the time did not simply abandon the theory but protected it from refutation by instead challenging the background assumption that the solar system contained only seven planets. This strategy bore fruit, notwithstanding the falsity of Newton's theory: by calculating the location of a hypothetical eighth planet influencing the orbit of Uranus, the astronomers Adams and Leverrier were eventually led to discover Neptune in 1846. But the very same strategy failed when used to try to explain the advance of the perihelion in Mercury's orbit by postulating the existence of “Vulcan”, an additional planet located between Mercury and the sun, and this phenomenon would resist satisfactory explanation until the arrival of Einstein's theory of general relativity. So it seems that Duhem was right to suggest not only that hypotheses must be tested as a group or a collection, but also that it is by no means a foregone conclusion which member of such a collection should be abandoned or revised in response to a failed empirical test or false implication. Indeed, this very example illustrates why Duhem's own rather hopeful appeal to the ‘good sense’ of scientists themselves in deciding when a given hypothesis ought to be abandoned promises very little if any relief from the general predicament of holist underdetermination.

As noted above, Duhem thought that the sort of underdetermination he had described presented a challenge only for theoretical physics, but subsequent thinking in the philosophy of science has tended to the opinion that the predicament Duhem described applies to theoretical testing in all fields of scientific inquiry. We cannot, for example, test an hypothesis about the phenotypic effects of a particular gene without presupposing a host of further beliefs about what genes are, how they work, how we can identify them, what other genes are doing, and so on. And in the middle of the 20th Century, W. V. O. Quine would incorporate confirmational holism and its associated concerns about underdetermination into an extraordinarily influential account of knowledge in general. As part of his famous (1951) critique of the widely accepted distinction between truths that are analytic (true by definition, or as a matter of logic or language alone) and those that are synthetic (true in virtue of some contingent fact about the way the world is), Quine argued instead that *all* of the beliefs we hold at any given time are linked in an interconnected web, which encounters our sensory experience only at its periphery:

The totality of our so-called knowledge or beliefs, from the most casual matters of geography and history to the profoundest laws of atomic physics or even of pure mathematics and logic, is a man-made fabric which impinges on experience only along the edges. Or, to change the figure, total science is like a field of force whose boundary conditions are experience. A conflict with experience at the periphery occasions readjustments in the interior of the field. But the total field is so underdetermined by its boundary conditions, experience, that there is much latitude of choice as to what statements to reevaluate in the light of any single contrary experience. No particular experiences are linked with any particular statements in the interior of the field, except indirectly through considerations of equilibrium affecting the field as a whole. (1951, 42–3)

One consequence of this general picture of human knowledge is that any and all of our beliefs are tested against experience only as a corporate body—or as Quine sometimes puts it, “The unit of empirical significance is the whole of science” (1951, p. 42).^[4] A mismatch between what the web as a whole leads us to expect and the sensory experiences we actually receive will occasion *some* revision in our beliefs, but which revision we should make to bring the web as a whole back into conformity with our experiences is radically underdetermined by those experiences themselves. If we find our belief that there are brick houses on Elm Street to be in conflict with our immediate sense experience, we might revise our beliefs about the houses on Elm Street, but we might equally well modify instead our beliefs about the appearance of brick, or about our present location, or innumerable other beliefs constituting the interconnected web—in a pinch we might even decide that our present sensory experiences are simply hallucinations! Quine's point was not that

any of these are particularly likely responses to recalcitrant experiences (indeed, an important part of his account is the explanation of why they are not), but instead that they would serve equally well to bring the web of belief as a whole in line with our experience. And if the belief that there are brick houses on Elm Street were sufficiently important to us, Quine insisted, it would be possible for us to preserve it “come what may” (in the way of empirical evidence), by making sufficiently radical adjustments elsewhere in the web of belief. It is in principle open to us, Quine argued, to revise even beliefs about logic, mathematics, or the meanings of our terms in response to recalcitrant experience; it might seem a tempting solution to certain persistent difficulties in quantum mechanics, for example, to reject classical logic’s law of the excluded middle (allowing physical particles to both have and not have some determinate classical physical property like position or momentum at a given time). The only test of a belief, Quine argued, is whether it fits into a web of connected beliefs that accords well with our experience *on the whole*. And because this leaves any and all beliefs in that web at least potentially subject to revision on the basis of our ongoing sense experience or empirical evidence, he insisted, there simply are no beliefs that are analytic in the originally supposed sense of immune to revision in light of experience or true no matter what the world is like.

Quine recognized, of course, that many of the logically possible ways of revising our beliefs in response to recalcitrant experiences that remain open to us strike us as ad hoc, perfectly ridiculous, or worse. He argues (1955) that our actual revisions of the web of belief seek to maximize the theoretical “virtues” of simplicity, familiarity, scope, and fecundity, along with conformity to experience, and elsewhere suggests that we typically seek to resolve conflicts between the web of our beliefs and our sensory experiences in accordance with a principle of “conservatism”, that is, by making the smallest possible number of changes to the least central beliefs we can that will suffice to reconcile the web with experience. That is, Quine recognized that when we encounter recalcitrant experience we are not usually at a loss to decide which of our beliefs to revise in response to it, but he claimed that this is simply because we are strongly disposed as a matter of fundamental psychology to prefer whatever revision requires the most minimal mutilation of the existing web of beliefs and/or maximizes virtues that he explicitly characterizes as pragmatic. Indeed, it would seem that on Quine’s view the very notion of a belief being more central or peripheral or in lesser or greater “proximity” to sense experience should be cashed out simply as a measure of our willingness to revise it in response to recalcitrant experience. That is, it would seem that what it *means* for one belief to be located “closer” to the sensory periphery of the web than another is simply that we are more likely to revise the first than the second if doing so would enable us to bring the web as a whole into conformity with otherwise recalcitrant sense experience. Thus, Quine saw the traditional distinction between analytic and synthetic beliefs as simply registering the endpoints of a psychological continuum ordering our beliefs according to the ease and likelihood with which we are prepared to revise them in order to reconcile the web as a whole with our sense experience.

2.2 Challenging the Rationality of Science

Perhaps it is unsurprising that such holist underdetermination has often been taken to pose a threat to the fundamental rationality of the scientific enterprise. The claim that the empirical evidence alone underdetermines our response to failed predictions or recalcitrant experience might even seem to fairly *invite* the suggestion that what systematically steps into the breach to do the further work of singling out just one or a few candidate responses to disconfirming evidence is (even if “pragmatic”) something irrational or at least arational in character. Imre Lakatos and Paul Feyerabend each suggested that because of underdetermination, the difference between empirically successful and unsuccessful theories or research programs is largely a function of the differences in talent, creativity, resolve, and resources of those who advocate them. And at least since the influential work of Thomas Kuhn, one important line of thinking about science has held that it is ultimately the social and political interests (in a suitably broad sense) of scientists themselves which serve to determine their responses to disconfirming evidence and therefore the further empirical, methodological, and other commitments of any given scientist or scientific community. Mary Hesse suggests that Quinean underdetermination showed why certain “non-logical” and “extra-empirical” considerations must play a role in theory choice, and claims that “it is only a short step from this philosophy of science to the suggestion that adoption of such criteria, that can be seen to be different for different groups and at different periods, should be explicable by social rather than logical factors” (1980, 33). And perhaps the most prominent modern day inheritors of this line of thinking are those scholars in the sociology of scientific knowledge (SSK) movement and in feminist science studies who argue that it is typically the career interests, political affiliations, intellectual allegiances, gender biases, and/or pursuit of power and influence by scientists

themselves which play a crucial or even decisive role in determining precisely which beliefs are abandoned or retained in response to conflicting evidence. The shared argumentative schema here is one on which holist underdetermination ensures that the evidence alone cannot do the work of picking out a single response to such conflicting evidence, thus something else must step in to do the job, and sociologists of scientific knowledge, feminist critics of science, and other interest-driven theorists of science each have their favored suggestions close to hand.

In a justly celebrated discussion, Larry Laudan (1990) argues that the significance of such underdetermination has been greatly exaggerated. Underdetermination actually comes in a wide variety of strengths, he insists, depending on precisely what is being asserted about the character, the availability, and (most importantly) the *rational defensibility* of the various competing hypotheses or ways of revising our beliefs that the evidence supposedly leaves us free to accept. Laudan usefully distinguishes a number of different dimensions along which claims of underdetermination vary in strength, and he goes on to insist that those who attribute dramatic significance to the thesis that our scientific theories are underdetermined by the evidence invariably defend only the weaker versions of that thesis, while they go on to draw dire consequences and shocking morals regarding the character and status of the scientific enterprise from much stronger versions. He suggests, for instance, that Quine's famous claim that any hypothesis can be preserved "come what may" can perhaps be defended simply as a description of what it is *psychologically* possible for human beings to do, but Laudan insists that in this form the thesis is simply bereft of interesting or important consequences for epistemology—the study of *knowledge*. The strong version of the thesis along this dimension instead asserts that it is always normatively or rationally *defensible* to retain any hypothesis in the light of any evidence whatsoever, but this latter, stronger version of the claim, Laudan suggests, is one for which no convincing evidence or argument has ever been offered. More generally, he insists, arguments for underdetermination turn on implausibly treating all logically possible responses to the evidence as equally justified or rationally defensible. For example, Laudan suggests that we might reasonably hold the resources of *deductive logic* to be insufficient to single out just one acceptable response to disconfirming evidence, but not that deductive logic *plus the sorts of ampliative principles of good reasoning typically deployed in scientific contexts* are insufficient to do so. Similarly, defenders of underdetermination might assert the *nonuniqueness* claim that for any given theory or web of beliefs there is at least one alternative that can also be reconciled with the available evidence, or the stronger *egalitarian* claim that *all* of the contraries of any given theory can be reconciled with the available evidence equally well. And the claim of such "reconciliation" itself disguises a wide range of further alternative possibilities: that our theories can be made *logically compatible* with any amount of disconfirming evidence (perhaps by the simple expedient of removing any claim(s) with which the evidence is in conflict), that any theory may be reformulated or revised so as to *entail* any piece of previously disconfirming evidence, or so as to *explain* previously disconfirming evidence, or that any theory can be made to be *as well supported empirically* by any collection of evidence as any other theory. And in *all* of these respects, Laudan claims, partisans have defended only the weaker forms of underdetermination while founding their further claims about and conceptions of the scientific enterprise on versions much stronger than those they have managed or even attempted to defend.

Laudan is certainly right to distinguish these various versions of holist underdetermination, and he is equally right to suggest that many of the thinkers he confronts have derived grand morals concerning the scientific enterprise from much stronger versions of underdetermination than they are able to defend, but the underlying situation is somewhat more complex than he suggests. Laudan's overarching claim is that champions of holist underdetermination show only that a wide variety of responses to disconfirming evidence are logically possible (or even just psychologically possible), rather than that these are all rationally defensible or equally well-supported by the evidence. But his straightforward appeal to further epistemic resources like ampliative principles of belief revision that are supposed to help narrow the merely logical possibilities down to those which are reasonable or rationally defensible is itself problematic, at least as part of any attempt to respond to Quine. This is because on Quine's holist picture of knowledge such further ampliative principles governing legitimate belief revision are, of course, *themselves* simply part of the web of our beliefs, and are therefore open to revision in response to recalcitrant experience as well—indeed, this is true even for the principles of deductive logic and the (consequent) demand for particular forms of logical consistency between parts of the web itself! So while it is true that the ampliative principles we currently embrace do not leave all logically or even psychologically possible responses to the evidence open to us (or leave us free to preserve any hypothesis "come what may"), our continued adherence to *these very principles*, rather than being willing to revise the web of belief so as to abandon them, is part of the phenomenon to which Quine is using underdetermination to draw our attention and cannot be taken for

granted without begging the question. Put another way, Quine does not simply ignore the further principles that function to ensure that we revise the web of belief in one way rather than others, but it follows from his account that such principles are themselves part of the web and therefore candidates for revision in our efforts to bring the web of beliefs into conformity (by the *resulting* web's own lights) with sensory experience. This recognition makes clear why it will be extremely difficult to say how the shift to an alternative web of belief (with alternative ampliative or even deductive principles of belief revision) should or even can be evaluated for its rational defensibility—each proposed revision will be maximally rational by the lights of the principles it itself sanctions.^[5] Of course we can rightly say that many candidate revisions would violate our *presently accepted* ampliative principles of rational belief revision, but the preference we have for those rather than the alternatives is itself a matter of their position in the existing web of belief we have inherited and the role that they themselves play in guiding the revisions we are inclined to make to that web in light of ongoing experience.

Thus, if we accept Quine's general picture of knowledge, it becomes quite difficult to disentangle normative from descriptive issues, or questions about the psychology of human belief revision from questions about the justifiability or rational defensibility of such revisions. It is in part for this reason that Quine famously suggests (1969, 82; see also p 75–76) that epistemology itself “falls into place as a chapter of psychology and hence of natural science”: the point is not that epistemology should simply be abandoned in favor of psychology, but instead that there is ultimately no way to draw a meaningful distinction between the two. (James Woodward, in comments on an earlier draft of this entry, pointed out that this makes it all the harder to assess the significance of Quinean underdetermination in light of Laudan's complaint or even know the rules for doing so, but in an important way this difficulty was Quine's point all along!) Quine's claim is that “[e]ach man is given a scientific heritage plus a continuing barrage of sensory stimulation; and the considerations which guide him in warping his scientific heritage to fit his continuing sensory promptings are, where rational, pragmatic” (1951, 46), but the role of these “pragmatic” considerations or principles in selecting just one of the many possible revisions of the web of belief in response to recalcitrant experience is not to be contrasted with those same principles having a rational or epistemic justification. Far from conflicting with or even being orthogonal to the search for truth and our efforts to render our beliefs maximally responsive to the evidence, Quine insists, revising our beliefs in accordance with such pragmatic principles “at bottom, is what evidence is” (1955, 251). Whether or not this strongly naturalistic conception of epistemology can ultimately be defended, it is misleading for Laudan to suggest that the thesis of underdetermination becomes trivial or obviously insupportable the moment we inquire into the rational defensibility rather than the mere logical or psychological possibility of alternative revisions to the holist's web of belief.

In fact, there is an important connection between this lacuna in Laudan's famous discussion and the further uses made of the thesis of underdetermination by sociologists of scientific knowledge, feminist epistemologists, and other vocal champions of holist underdetermination. When faced with the invocation of further ampliative standards or principles that supposedly rule out some responses to disconfirmation as irrational or unreasonable, these thinkers typically respond by insisting that the embrace of such further standards or principles (or perhaps their application to particular cases) is *itself* underdetermined, historically contingent, and/or subject to ongoing social negotiation. For this reason, they suggest, such appeals (and their success or failure in convincing the members of a given community) should be explained by reference to the same broadly social and political interests that they claim are at the root of theory choice and belief change in science more generally (see, e.g., Shapin and Schaffer, 1982). On both accounts, then, our response to recalcitrant evidence or a failed prediction is constrained in important ways by preexisting features of the existing web of beliefs, but for Quine the continuing force of these constraints is ultimately imposed by the fundamental principles of human psychology (such as our preference for minimal mutilation of the web, or the pragmatic virtues of simplicity, fecundity, etc.), while for interest-driven theorists of science the continuing force of any such constraints is limited only by the ongoing negotiated agreement of the communities of scientists who respect them.

As this last contrast makes clear, however, recognizing the limitations of Laudan's critique of Quine and the fact that we cannot dismiss holist underdetermination with any straightforward appeal to ampliative principles of good reasoning by itself does nothing to establish the further *positive* claims about belief revision advanced by interest-driven theorists of science. Even simply conceding that theory choice or belief revision in science is indeed underdetermined by the evidence in just the ways that Duhem and/or Quine suggested leaves entirely open whether it is instead the social or political interests of scientists themselves

that do the further work of singling out the particular beliefs or responses to falsifying evidence that any particular scientist or scientific community will actually adopt or find compelling. Even many of those philosophers of science who are most strongly convinced of the general significance of various forms of underdetermination itself remain deeply skeptical of this latter thesis and thoroughly unconvinced by the empirical evidence that has been offered in support of it (usually in the form of case studies of particular historical episodes in science).

3. Contrastive Underdetermination, Empirical Equivalents, and Unconceived Alternatives

3.1 Contrastive Underdetermination: Back to Duhem

Although it is also a form of underdetermination, what we described in Section 1 above as contrastive underdetermination raises fundamentally different issues from the holist variety considered in Section 2 (a recent book-length treatment of many of these issues is Bonk 2008). This is clearly evident in Duhem's original writings concerning so-called crucial experiments, where he seeks to show that even when we explicitly suspend any concerns about holist underdetermination, the contrastive variety remains an obstacle to our discovery of truth in theoretical science:

But let us admit for a moment that in each of these systems [concerning the nature of light] everything is compelled to be necessary by strict logic, except a single hypothesis; consequently, let us admit that the facts, in condemning one of the two systems, condemn once and for all the single doubtful assumption it contains. Does it follow that we can find in the 'crucial experiment' an irrefutable procedure for transforming one of the two hypotheses before us into a demonstrated truth? Between two contradictory theorems of geometry there is no room for a third judgment; if one is false, the other is necessarily true. Do two hypotheses in physics ever constitute such a strict dilemma? Shall we ever dare to assert that no other hypothesis is imaginable? Light may be a swarm of projectiles, or it may be a vibratory motion whose waves are propagated in a medium; is it forbidden to be anything else at all? ([1914] 1954, 189)

Contrastive underdetermination is so-called because it questions the ability of the evidence to confirm any given hypothesis *against alternatives*, and the central focus of discussion in this connection (equally often regarded as "the" problem of underdetermination) concerns the character of the supposed alternatives. Of course the two problems are not entirely disconnected, because it is open to us to consider alternative possible modifications of the web of beliefs as alternative theories or theoretical "systems" between which the empirical evidence alone is powerless to decide. But we have already seen that one *need* not think of the alternative responses to recalcitrant experience as competing theoretical alternatives to appreciate the character of the holist's challenge, and we will see that one need not embrace any version of holism about confirmation to appreciate the quite distinct problem that the available evidence might support more than one theoretical alternative. It is perhaps most useful here to think of holist underdetermination as starting from a particular theory or body of beliefs and claiming that our revision of those beliefs in response to new evidence may be underdetermined, while contrastive underdetermination instead starts from a given body of evidence and claims that more than one theory may be well-supported by that very evidence. Part of what has contributed to the conflation of these two problems is the holist presuppositions of those who originally made them famous. After all, on Quine's view we simply revise the web of belief in response to recalcitrant experience, and so the suggestion that there are multiple possible revisions of the web available in response to any particular evidential finding just *is* the claim that there are in fact many different "theories" (i.e. candidate webs of belief) that are equally well-supported by any given body of data.^[6] But if we give up such extreme holist views of evidence, meaning, and/or confirmation, the two problems take on very different identities, with very different considerations in favor of taking them seriously, very different consequences, and very different candidate solutions. Notice, for instance, that even if we somehow knew that no other hypothesis on a given subject was well-confirmed by a given body of data, that would not tell us where to place the blame or which of our beliefs to give up if the remaining hypothesis in conjunction with others subsequently resulted in a failed empirical prediction. And as Duhem suggests above, even if we supposed that we somehow knew exactly which of our hypotheses to blame in response to a failed empirical

prediction, this would not help us to decide whether or not there are other hypotheses available that are equally well-confirmed by the data we actually have.

One way to see why not is to consider an analogy that champions of contrastive underdetermination have sometimes used to support their case. If we consider any finite group of data points, an elementary proof reveals that there are an infinite number of distinct mathematical functions describing different curves that will pass through all of them. As we add further data to our initial set we will definitively eliminate functions describing curves which no longer capture all of the data points in the new, larger set, but no matter how much data we accumulate, the proof guarantees that there will always be an infinite number of functions *remaining* that define curves including all the data points in the new set and which would therefore seem to be equally well supported by the empirical evidence. *No* finite amount of data will *ever* be able to narrow the possibilities down to just a single function or indeed, any finite number of candidate functions, from which the distribution of data points we have might have been generated. Each new data point we gather eliminates an infinite number of curves that *previously* fit all the data (so the problem here is not the holist's challenge that we do not know which beliefs to give up in response to failed predictions or disconfirming evidence), but also leaves an infinite number still in contention.

3.2 Empirically Equivalent Theories

Of course, generating and testing fundamental scientific hypotheses is rarely if ever a matter of finding curves that fit collections of data points, so nothing follows directly from this mathematical analogy for the significance of contrastive underdetermination in most scientific contexts. But Bas van Fraassen has offered an extremely influential line of argument intended to show that such contrastive underdetermination is a serious concern for scientific theorizing more generally. In *The Scientific Image* (1980), van Fraassen uses a now-classic example to illustrate the possibility that even our best scientific theories might have *empirical equivalents*: that is, alternative theories making the very same empirical predictions, and which therefore cannot be better or worse supported by any *possible* body of evidence. Consider Newton's cosmology, with its laws of motion and gravitational attraction. As Newton himself realized, van Fraassen points out, exactly the same predictions are made by the theory whether we assume that the entire universe is at rest or assume instead that it is moving with some constant velocity in any given direction: from our position within it, we have no way to detect constant, absolute motion by the universe as a whole. Thus, van Fraassen argues, we are here faced with empirically equivalent scientific theories: Newtonian mechanics and gravitation conjoined either with the fundamental assumption that the universe is at absolute rest (as Newton himself believed), or with any one of an infinite variety of alternative assumptions about the constant velocity with which the universe is moving in some particular direction. All of these theories make all and only the same empirical predictions, so no evidence will ever permit us to decide between them on empirical grounds.^[7]

Van Fraassen is widely regarded as holding that the prospect of contrastive underdetermination grounded in such empirical equivalents should lead us to restrict our epistemic ambitions for the scientific enterprise. His constructive empiricism holds that the aim of science is not to find true theories, but only theories that are empirically adequate: that is, theories whose claims about *observable phenomena* are all true. Since the empirical adequacy of a theory is not threatened by the existence of another that is empirically equivalent to it, fulfilling this aim has nothing to fear from the possibility of such empirical equivalents. In reply, many critics have suggested that van Fraassen gives no reasons for restricting belief to empirical adequacy that could not also be used to argue for suspending our belief in the *future* empirical adequacy of our best present theories: of course there *could* be empirical equivalents to our best theories, but there could also be theories equally well-supported by all the evidence up to the present which diverge in their predictions about observables in future cases not yet tested. This challenge seems to miss the point of Van Fraassen's epistemic voluntarism: his claim is that we should believe no more *but also no less* than we need to make sense of and take full advantage of our scientific theories, and a commitment to the empirical adequacy of our theories, he suggests, is the least we can get away with in this regard. Of course it is true that we are running some epistemic risk in believing in even the full empirical adequacy of our present theories, but the risk is considerably less than what we assume in believing in their truth, it is the minimum we need to take full advantage of the fruits of our scientific labors, and, he suggests, “it is not an epistemic principle that one might as well hang for a sheep as a lamb” (1980, 72).

In an influential discussion, Larry Laudan and Jarrett Leplin (1991) argue that philosophers of science have invested even the bare possibility that our theories might have empirical equivalents with far too much epistemic significance. Notwithstanding the popularity of the presumption that there are empirically equivalent rivals to every theory, they argue, the conjunction of several familiar and relatively uncontroversial epistemological theses is sufficient to defeat it. Because the boundaries of what is observable change as we develop new experimental methods and instruments, because auxiliary assumptions are always needed to derive empirical consequences from a theory (cf. confirmational holism, above), and because these auxiliary assumptions are themselves subject to change over time, Laudan and Leplin conclude that there simply is no guarantee that any two theories judged to be empirically equivalent at a given time will remain so as the state of our knowledge advances. Accordingly, any judgment of empirical equivalence is both defeasible and relativized to a particular state of science. So even if two theories are empirically equivalent at a given time this is no guarantee that they will *remain* so, and thus there is no foundation for a general pessimism about our ability to distinguish theories that are empirically equivalent to each other on empirical grounds. Although they concede that we could have good reason to think that particular theories have empirically equivalent rivals, this must be established case-by-case rather than by any general argument or presumption.

A fairly standard reply to this line of argument is to suggest that what Laudan and Leplin really show is that the notion of empirical equivalence must be applied to larger collections of beliefs than those traditionally identified as scientific theories—at least large enough to encompass the auxiliary assumptions needed to derive empirical predictions from them. At the extreme, perhaps this means that the notion of empirical equivalents (or at least timeless empirical equivalents) cannot be applied to anything less than “systems of the world” (i.e. total Quinean webs of belief), but even that is not fatal: what the champion of contrastive underdetermination asserts is that there are empirically equivalent systems of the world that *incorporate* different theories of the nature of light, or spacetime, or whatever. On the other hand, it might seem that quick examples like van Fraassen's variants of Newtonian cosmology do not serve to make *this* thesis as plausible as the more limited claim of empirical equivalence for individual theories. It seems equally natural, however, to respond to Laudan and Leplin simply by conceding the variability in empirical equivalence but insisting that this is not enough to undermine the problem. Empirical equivalents create a serious obstacle to belief in a theory so long as there is *some* empirical equivalent to that theory at any given time, but it need not be the same one at each time. On this line of thinking, cases like van Fraassen's Newtonian example illustrate how easy it is for theories to admit of empirical equivalents at any given time, and thus constitute a reason for thinking that there probably are or will be empirical equivalents to any given theory at any particular time we consider it, assuring that whenever the question of belief in a given theory arises, the challenge posed to it by contrastive underdetermination arises as well.

Laudan and Leplin also suggest, however, that even if the universal existence of empirical equivalents were conceded, this would do much less to establish the significance of underdetermination than its champions have supposed, because “theories with exactly the same empirical consequences may admit of differing degrees of evidential support” (1991, 465). A theory may be better supported than an empirical equivalent, for instance, because the former but not the latter is derivable from a more general theory whose consequences include a third, well supported, hypothesis. More generally, the belief-worthiness of an hypothesis depends crucially on how it is connected or related to other things we believe and the evidential support we have for those other beliefs.^[8] Laudan and Leplin suggest that we have invited the specter of rampant underdetermination only by failing to keep this familiar home truth in mind and instead implausibly identifying the evidence bearing on a theory exclusively with the theory's own entailments or empirical consequences (but cf. Tulodziecki 2012). This impoverished view of evidential support, they argue, is in turn the legacy of a failed foundationalist and positivistic approach to the philosophy of science which mistakenly assimilates epistemic questions about how to decide whether or not to believe a theory to semantic questions about how to establish a theory's meaning or truth-conditions.

John Earman (1993) has argued that this dismissive diagnosis does not do justice to the threat posed by underdetermination. He argues that worries about underdetermination are an aspect of the more general question of the reliability of our inductive methods for determining beliefs, and notes that we cannot decide how serious a problem underdetermination poses without specifying (as Laudan and Leplin do not) the inductive methods we are considering. Earman regards some version of Bayesianism as our most promising form of inductive methodology, and he proceeds to show that challenges to the long-run reliability of our Bayesian methods can be motivated by considerations of the empirical indistinguishability (in several

different and precisely specified senses) of hypotheses stated in any language richer than that of the evidence itself that do not amount simply to general skepticism about those inductive methods. In other words, he shows that there are more reasons to worry about underdetermination concerning inferences to hypotheses about unobservables than to, say, inferences about unobserved observables. He also goes on to argue that at least two genuine cosmological theories have serious, nonskeptical, and nonparasitic empirical equivalents: the first essentially replaces the gravitational field in Newtonian mechanics with curvature in spacetime itself, [9] while the second recognizes that Einstein's General Theory of Relativity permits cosmological models exhibiting different global topological features which cannot be distinguished by any evidence inside the light cones of even idealized observers who live forever. [10] And he suggests that “the production of a few concrete examples is enough to generate the worry that only a lack of imagination on our part prevents us from seeing comparable examples of underdetermination all over the map” (1993, 31) even as he concedes that his case leaves open just how far the threat of underdetermination extends (1993, 36).

Most philosophers of science, however, have not embraced the idea that it is only lack of imagination which prevents us from finding empirical equivalents to our scientific theories generally. They note that the convincing examples of empirical equivalents we do have are all drawn from a single domain of highly mathematized scientific theorizing in which the background constraints on serious theoretical alternatives are far from clear, and suggest that it is therefore reasonable to ask whether even a small handful of such examples should make us believe that there are probably empirical equivalents to most of our scientific theories most of the time. They concede that it is always *possible* that there are empirical equivalents to even our best scientific theories concerning any domain of nature, but insist that we should not be willing to suspend belief in any *particular* theory until some convincing alternative to it can actually be produced: as Philip Kitcher puts it, “give us a rival explanation, and we'll consider whether it is sufficiently serious to threaten our confidence” (1993, 154; see also Leplin 1997, Achinstein 2002). That is, these thinkers insist that until we are able to *actually construct* an empirically equivalent alternative to a given theory, the bare possibility that such equivalents exist is insufficient to justify suspending belief in the best theories we do have. And for this same reason most philosophers of science are unwilling to follow van Fraassen into what they regard as constructive empiricism's unwarranted epistemic modesty. Even if van Fraassen is right about the most minimal beliefs we must hold in order to take full advantage of our scientific theories, most thinkers do not see why we should believe the least we can get away with rather than believing the most we are entitled to by the evidence we have.

Philosophers of science have responded in a variety of ways to the suggestion that a few or even a small handful of serious examples of empirical equivalents does not suffice to establish that there are probably such equivalents to most scientific theories in most domains of inquiry. One such reaction has been to invite more careful attention to the details of particular examples of putative underdetermination: considerable work has been devoted to assessing the threat of underdetermination in the case of particular scientific theories (for a recent example, see Butterfield (2012)). Another reaction has been to investigate whether particular kinds of theories or domains of science (e.g. 'historical' vs. 'experimental' sciences) are more vulnerable to problems of underdetermination than others and, if so, why (see Cleland (2002), Carman (2005), Turner (2005, 2007), Stanford (2010), Forber and Griffith (2011)). But champions of contrastive underdetermination have most frequently responded by seeking to argue that *all* theories have empirical equivalents, typically by proposing something like an algorithmic procedure for generating such equivalents from any theory whatsoever. Stanford (2001, 2006) suggests that these efforts to *prove* that all our theories must have empirical equivalents fall roughly but reliably into global and local varieties, and that neither makes a convincing case for a distinctive scientific problem of contrastive underdetermination. Global algorithms are well-represented by Andre Kukla's (1996) suggestion that from any theory T we can immediately generate such empirical equivalents as T' (the claim that T 's observable consequences are true, but T itself is false), T'' (the claim that the world behaves according to T when observed, but some specific incompatible alternative otherwise), and the hypothesis that our experience is being manipulated by powerful beings in such a way as to make it appear that T is true. But such possibilities, Stanford argues, amount to nothing more than the sort of Evil Deceiver to which Descartes appealed in order to doubt any of his beliefs that could possibly be doubted (see Section 1, above). Such radically skeptical scenarios pose an equally powerful (or powerless) challenge to any knowledge claim whatsoever, no matter how it is arrived at or justified, and thus pose no *special* problem or challenge for beliefs offered to us by theoretical science. If global algorithms like Kukla's are the only reasons we can give for taking underdetermination seriously in a

scientific context, then there is no distinctive problem of the underdetermination of scientific theories by data, only a salient reminder of the irrefutability of classically Cartesian or radical skepticism.^[11]

By contrast to such global strategies for generating empirical equivalents, local algorithmic strategies instead begin with some particular scientific theory and proceed to generate alternative versions that are equally well supported by all possible evidence. This is what van Fraassen does with the example of Newtonian cosmology, showing that an infinite variety of supposed empirical equivalents can be produced by ascribing different constant absolute velocities to the universe as a whole. But Stanford suggests that empirical equivalents generated in this way are also insufficient to show that there is a distinctive and genuinely troubling form of underdetermination afflicting scientific theories, because they rely on simply saddling particular scientific theories with further claims for which those theories themselves (together with whatever background beliefs we actually hold) imply that we cannot have any evidence. Such empirical equivalents invite the natural response that they force our theories to undertake commitments that they never should have in the first place. Such claims, it seems, should simply be excised from the theories themselves, leaving over just the claims that sensible defenders would have held were all we were entitled to believe by the evidence in any case. In van Fraassen's Newtonian example, for instance, this could be done simply by undertaking no commitment concerning the absolute velocity and direction (or lack thereof) of the universe as a whole. To put the point another way, if we believe a given scientific theory when one of the empirical equivalents we could generate from it by the local algorithmic strategy is correct instead, most of what we originally believed will nonetheless turn out to be straightforwardly true.

3.3 Unconceived Alternatives and A New Induction

Stanford (2001, 2006) concludes that no convincing general case has been made for the presumption that there are empirically equivalent rivals to all or most scientific theories, or to any theories besides those for which such equivalents can actually be constructed. But he goes on to insist that empirical equivalents are no essential part of the case for a significant problem of contrastive underdetermination. Our efforts to confirm scientific theories, he suggests, are no less threatened by what Larry Sklar (1975, 1981) has called “transient” underdetermination, that is, theories which are *not* empirically equivalent but are equally (or at least reasonably) well confirmed by all the evidence we happen to have in hand at the moment, so long as this transient predicament is also “recurrent”, that is, so long as we think that there is (probably) at least one such (fundamentally distinct) alternative available—and thus the transient predicament re-arises—whenever we are faced with a decision about whether to believe a given theory at a given time. Stanford argues that a convincing case for contrastive underdetermination of this recurrent, transient variety can indeed be made, and that the evidence for it is available in the historical record of scientific inquiry itself.

Stanford concedes that present theories are not transiently underdetermined by the theoretical alternatives we have actually developed and considered to date: we think that our own scientific theories are considerably better confirmed by the evidence than any rivals we have actually produced. The central question, he argues, is whether we should believe that there are well confirmed alternatives to our best scientific theories that are *presently unconceived* by us. And the primary reason we should believe that there are, he claims, is the long history of repeated transient underdetermination by *previously* unconceived alternatives across the course of scientific inquiry. In the progression from Aristotelian to Cartesian to Newtonian to contemporary mechanical theories, for instance, the evidence available at the time each earlier theory dominated the practice of its day also offered compelling support for each of the later alternatives (unconceived at the time) that would ultimately come to displace it. Stanford's “New Induction” over the history of science claims that this situation is typical; that is, that “we have, throughout the history of scientific inquiry and in virtually every scientific field, repeatedly occupied an epistemic position in which we could conceive of only one or a few theories that were well confirmed by the available evidence, while subsequent inquiry would routinely (if not invariably) reveal further, radically distinct alternatives as well confirmed by the previously available evidence as those we were inclined to accept on the strength of that evidence” (2006, 19). In other words, Stanford claims that in the past we have repeatedly failed to exhaust the space of fundamentally distinct theoretical possibilities that were well confirmed by the existing evidence, and that we have every reason to believe that we are probably also failing to exhaust the space of such alternatives that are well confirmed by the evidence we have at present. Much of the rest of his case is taken up with discussing historical examples illustrating that earlier scientists did not simply ignore or dismiss, but instead genuinely *failed to conceive of* the serious, fundamentally distinct theoretical possibilities that would ultimately come to displace the

theories they defended, only to be displaced in turn by others that were similarly unconceived at the time. He concludes that “the history of scientific inquiry itself offers a straightforward rationale for thinking that there typically are alternatives to our best theories equally well confirmed by the evidence, even when we are unable to conceive of them at the time” (2006, 20; for reservations and criticisms concerning this line of argument, see Magnus 2006, 2010; Godfrey-Smith 2008; Chakravartty 2008; Devitt 2011; Ruhmkorff 2011). Stanford concedes, however, that the historical record can offer only fallible evidence of a distinctive, general problem of contrastive scientific underdetermination, rather than the kind of deductive proof that champions of the case from empirical equivalents have typically sought. Thus, claims and arguments about the various forms that underdetermination may take, their causes and consequences, and the further significance they hold for the scientific enterprise as a whole continue to evolve in the light of ongoing controversy, and the underdetermination of scientific theory by evidence remains very much a live and unresolved issue in the philosophy of science.

Bibliography

- Achinstein, P., 2002, “Is There A Valid Experimental Argument for Scientific Realism?”, *Journal of Philosophy*, 99: 470–495.
- Bonk, T., 2008, *Underdetermination: An Essay on Evidence and the Limits of Natural Knowledge*, Dordrecht, The Netherlands: Springer.
- Butterfield, J., 2012, “Underdetermination in Cosmology: An Invitation”, *Proceedings of the Aristotelian Society* (Supplementary Volume), 86: 1–18.
- Carman, C., 2005, “The Electrons of the Dinosaurs and the Center of the Earth”, *Studies in History and Philosophy of Science*, 36: 171–174.
- Chakravartty, A., 2008, “What You Don't Know Can't Hurt You: Realism and the Unconceived”, *Philosophical Studies*, 137: 149–158.
- Cleland, C., 2002, “Methodological and Epistemic Differences Between Historical Science and Experimental Science”, *Philosophy of Science*, 69: 474–496.
- Descartes, R., [1640] 1996, *Meditations on First Philosophy*, trans. by John Cottingham, Cambridge: Cambridge University Press.
- Devitt, M., 2011, “Are Unconceived Alternatives a Problem for Scientific Realism”, *Journal for General Philosophy of Science*, 42: 285–293.
- Duhem, P., [1914] 1954, *The Aim and Structure of Physical Theory*, trans. from 2nd ed. by P. W. Wiener; originally published as *La Théorie Physique: Son Objet et sa Structure* (Paris: Marcel Riviera & Cie.), Princeton, NJ: Princeton University Press.
- Earman, J., 1993, “Underdetermination, Realism, and Reason”, *Midwest Studies in Philosophy*, 18: 19–38.
- Feyerabend, P., 1975, *Against Method*, London: Verso.
- Forber, P. and Griffith, E., 2011, “Historical Reconstruction: Gaining Epistemic Access to the Deep Past”, *Philosophy and Theory in Biology*, 3. doi:10.3998/ptb.6959004.0003.003
- Gilles, D., 1993, “The Duhem Thesis and the Quine Thesis,” in *Philosophy of Science in the Twentieth Century*, Oxford: Blackwell Publishers, pp. 98–116.
- Glymour, C., 1970, “Theoretical Equivalence and Theoretical Realism”, *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1970: 275–288.
- , 1977, “The Epistemology of Geometry”, *Noûs*, 11: 227–251.
- , 1980, *Theory and Evidence*, Princeton, NJ.: Princeton University Press.
- , 2013, “Theoretical Equivalence and the Semantic View of Theories”, *Philosophy of Science*, 80: 286–297.
- Godfrey-Smith, P., 2008, “Recurrent, Transient Underdetermination and the Glass Half-Full”, *Philosophical Studies*, 137: 141–148.
- Goodman, N., 1955, *Fact, Fiction, and Forecast*, Indianapolis: Bobbs-Merrill.
- Halvorson, H., 2012, “What Scientific Theories Could Not Be”, *Philosophy of Science*, 79: 183–206.
- , 2013, “The Semantic View, If Plausible, Is Syntactic”, *Philosophy of Science*, 80: 475–478.
- Hesse, M., 1980, *Revolutions and Reconstructions in the Philosophy of Science*, Brighton: Harvester Press.
- Kitcher, P., 1993, *The Advancement of Science*, New York: Oxford University Press.
- Kuhn, T., [1962] 1996, *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press, 3rd edition.
- Kukla, A., 1996, “Does Every Theory Have Empirically Equivalent Rivals?”, *Erkenntnis*, 44: 137–166.

- Lakatos, I. 1970, “Falsification and the Methodology of Scientific Research Programmes”, in *Criticism and the Growth of Knowledge*, I. Lakatos and A. Musgrave (eds.), Cambridge: Cambridge University Press, pp. 91–196.
- Laudan, L., 1990, “Demystifying Underdetermination”, in *Scientific Theories*, C. Wade Savage (ed.), (Series: Minnesota Studies in the Philosophy of Science, vol. 14), Minneapolis: University of Minnesota Press, pp. 267–297.
- Laudan, L. and Leplin, J., 1991, “Empirical Equivalence and Underdetermination”, *Journal of Philosophy*, 88: 449–472.
- Leplin, J., 1997, *A Novel Defense of Scientific Realism*, New York: Oxford University Press.
- Magnus, P. 2006, “What's New About the New Induction?”, *Synthese*, 148: 295–301.
- , 2010, “Inductions, Red Herrings, and the Best Explanation for the Mixed Record of Science”, *British Journal for the Philosophy of Science*, 61: 803–819.
- Manchak, J. 2009, “Can We Know the Global Structure of Spacetime?”, *Studies in History and Philosophy of Modern Physics*, 40: 53–56.
- Mill, J. S., [1867] 1900, *A System of Logic, Ratiocinative and Inductive, Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*, New York: Longmans, Green, and Co.
- Norton, J. 2008, “Must Evidence Underdetermine Theory?”, in *The Challenge of the Social and the Pressure of Practice: Science and Values Revisited*, M. Carrier, D. Howard, and J. Kourany (eds.), Pittsburgh: University of Pittsburgh Press, 17–44.
- Quine, W. V. O., 1951, “Two Dogmas of Empiricism”, Reprinted in *From a Logical Point of View*, 2nd Ed., Cambridge, MA: Harvard University Press, pp. 20–46.
- , 1955, “Posits and Reality”, Reprinted in *The Ways of Paradox and Other Essays*, 2nd Ed., Cambridge, MA: Harvard University Press, pp. 246–254.
- , 1969, “Epistemology Naturalized”, in *Ontological Relativity and Other Essays*, New York: Columbia University Press, pp. 69–90.
- , 1975, “On Empirically Equivalent Systems of the World”, *Erkenntnis*, 9: 313–328.
- , 1990, “Three Indeterminacies”, in *Perspectives on Quine*, R. B. Barrett and R. F. Gibson, (eds.), Cambridge, MA: Blackwell, pp. 1–16.
- Ruhmkorff, S., 2011, “Some Difficulties for the Problem of Unconceived Alternatives”, *Philosophy of Science*, 78: 875–886.
- Shapin, S. and Shaffer, S., 1982, *Leviathan and the Air Pump*, Princeton: Princeton University Press.
- Sklar, L., 1975, “Methodological Conservatism”, *Philosophical Review*, 84: 384–400.
- , 1981, “Do Unborn Hypotheses Have Rights?”, *Pacific Philosophical Quarterly*, 62: 17–29.
- , 1982, “Saving the Noumena”, *Philosophical Topics*, 13: 49–72.
- Stanford, P. K., 2001, “Refusing the Devil's Bargain: What Kind of Underdetermination Should We Take Seriously?”, *Philosophy of Science*, 68 [Proceedings]: S1–S12.
- , 2006, *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*, New York: Oxford University Press.
- , 2010, “Getting Real: The Hypothesis of Organic Fossil Origins”, *The Modern Schoolman*, 87: 219–243
- Tulodziecki, D., 2012, “Epistemic Equivalence and Epistemic Incapacitation”, *British Journal for the Philosophy of Science*, 63: 313–328.
- Turner, D., 2005, “Local Underdetermination in Historical Science”, *Philosophy of Science*, 72: 209–230.
- , 2007, *Making Prehistory: Historical Science and the Scientific Realism Debate*, Cambridge: Cambridge University Press.
- Van Fraassen, B., 1980, *The Scientific Image*, Oxford: Oxford University Press.

Academic Tools



[How to cite this entry.](#)



[Preview the PDF version of this entry](#) at the [Friends of the SEP Society](#).



[Look up this entry topic](#) at the [Indiana Philosophy Ontology Project](#) (InPhO).



[Enhanced bibliography for this entry](#) at [PhilPapers](#), with links to its database.

Other Internet Resources

[Please contact the author with suggestions.]

Related Entries

[confirmation](#) | [constructive empiricism](#) | [Duhem, Pierre](#) | [epistemology: naturalism in](#) | [feminist \(interventions\): epistemology and philosophy of science](#) | [Feyerabend, Paul](#) | [induction: problem of](#) | [Quine, Willard van Orman](#) | [scientific knowledge: social dimensions of](#) | [scientific realism](#)

Acknowledgments

I have benefited from discussing both the organization and content of this article with many people including audiences and participants at the 2009 Pittsburgh Workshop on Underdetermination and the 2009 Southern California Philosophers of Science retreat, as well as the participants in graduate seminars both at UC Irvine and Pittsburgh. Special thanks are owed to John Norton, P. D. Magnus, John Manchak, Bennett Holman, Penelope Maddy, Jeff Barrett, David Malament, John Earman, and James Woodward.

[Copyright © 2013](#) by
Kyle Stanford <stanford@uci.edu>

Open access to the Encyclopedia has been made possible by a [world-wide funding initiative](#). See the list of [contributing institutions](#). If your institution is not on the list, please consider asking your librarians to contribute.

Stanford | Center for the Study of
Language and Information

The Stanford Encyclopedia of Philosophy is [copyright © 2016](#) by [The Metaphysics Research Lab](#), Center for the Study of Language and Information (CSLI), Stanford University

Library of Congress Catalog Data: ISSN 1095-5054