Refusing the Devil's Bargain: What Kind of Underdetermination Should We Take Seriously?

Author(s): P. Kyle Stanford

Source: *Philosophy of Science*, Vol. 68, No. 3, Supplement: Proceedings of the 2000 Biennial Meeting of the Philosophy of Science Association. Part I: Contributed Papers (Sep., 2001), pp. S1-S12

Published by: The University of Chicago Press on behalf of the Philosophy of Science Association

Stable URL: http://www.jstor.org/stable/3080930

Accessed: 26-09-2016 19:20 UTC

Refusing the Devil's Bargain: What Kind of Underdetermination Should We Take Seriously?

P. Kyle Stanford†‡
University of California, Irvine

Advocates have sought to *prove* that underdetermination obtains because all theories have *empirical equivalents*. But algorithms for generating empirical equivalents simply exchange underdetermination for familiar philosophical chestnuts, while the few convincing examples of empirical equivalents will not support the desired sweeping conclusions. Nonetheless, underdetermination does not depend on empirical equivalents: our warrant for current theories is equally undermined by presently unconceived alternatives as well-confirmed merely by the *existing* evidence, so long as this transient predicament recurs for each theory and body of evidence we consider. The historical record supports the claim that this recurrent, transient underdetermination predicament is our own.

1. Introduction. Nearly a century ago, Pierre Duhem wondered whether there might not be alternatives to even our best scientific theories that remained unconceived by us despite being supported by the available evidence:

> 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

†Send requests for reprints to the author, Program in Logic and Philosophy of Science, University of California, Irvine, CA 92697–5100; email: stanford@uci.edu.

‡My thanks to David Malament, Jeff Barrett, Philip Kitcher, Pen Maddy, Baron Reed, the members of my Winter 1999 Realism seminar, and several anonymous reviewers from the National Science Foundation for helpful comments and discussion of the material in this paper, as well as to members of audiences at the Claremont Colleges and at meetings of the Philosophy of Science Association and the British Society for the Philosophy of Science. Thanks also to the University of California at Irvine for supporting this work with a Career Development Award.

Philosophy of Science, 68 (Proceedings) pp. S1–S12. 0031-8248/2001/68supp-0001$0.00

> whose waves are propagated in a medium; is it forbidden to be anything else at all? ([1914] 1954, 189–190)[1]

The startling suggestion that even our best efforts at scientific theorizing might fail to exhaust the set of well-confirmed possibilities has enjoyed considerable influence, but its plausibility is hard to assess: critics have consistently wondered why we should either just *assume* that there are always alternative theories equally well-confirmed by the evidence or let the bare possibility that there *might* be prevent us from believing the best-confirmed theories we do have.

Challenged to show that underdetermination is anything more than a speculative possibility, advocates have sought to demonstrate that we can (ideally in an algorithmic fashion) always produce *empirical equivalents* to a given theory, i.e., competitors that make identical empirical predictions and therefore cannot be distinguished from it by any possible evidence.[2] Here I will argue that the case from empirical equivalents earns a Pyrrhic victory, for it succeeds only where it gives up any distinctive problem of scientific underdetermination altogether, but that the genuine threat of such underdetermination is nonetheless supported by a New Induction over the History of Science.

2. Empirical Equivalents and Underdetermination. Algorithms for generating empirical equivalents fall roughly but reliably into *global* and *local* varieties. Global algorithms are designed to produce empirical equivalents from absolutely any theory and are perhaps best exemplified by Kukla's (1996; see also 1993) appeals to such all-purpose alternatives to any theory T 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), the hypothesis of the Makers (the debatably coherent fantasy that we and our apparently T-governed world are part of an elaborate computer simulation), and the hypothesis of the Manipulators (that our experience is manipulated by powerful beings in such a way as to make it appear that T is true). Kukla devotes his efforts to defending such proposals from the accusation

1. Notice that this worry depends in no way on Duhem's equally famous commitment to holism, although holism grounds the distinct worry that theory choice is underdetermined because any theory may be retained in light of any evidence if we make suitable changes to it or the set of auxiliary assumptions we accept (of course, such changes might also be presently unconceived).

2. For Earman (1993), the crucial sense of empirical equivalence (his EI_3) obtains between two hypotheses just in case two worlds in which those two hypotheses are respectively true need not be distinguished by some piece of empirical evidence. The differences between these formulations will not matter for our purposes.

(see Laudan and Leplin 1991; Hoefer and Rosenberg 1994) that they are not "real theories" at all.

But this is beside the point, I suggest, for *whether or not* such farfetched scenarios are real theories they amount to no more than a salient presentation of the possibility of radical or Cartesian skepticism.[3] While many contemporary philosophers are inclined to grant the irrefutability of such skepticism, underdetermination was supposed to represent a distinct and important problem, arising perspicuously in the context of scientific theorizing about inaccessible domains of nature and troubling *even those who never hoped to defend their scientific beliefs to the truly radical skeptic.* Thus, if Cartesian fantasies are the only reasons we can give for taking underdetermination seriously, then there simply *is no* distinctive problem of scientific underdetermination to worry about, for the worry *just is* the familiar specter of radical skepticism.

The same response applies to some famous non-algorithmic *examples* of empirical equivalents, like the notorious prospect of a continuously shrinking universe with compensatory changes in physical constants making this state of affairs undetectable to us (i.e., theories we describe as making unmotivated and/or wildly implausible assumptions about nature). *Some* judgments of prior plausibility are required in order to escape radical or Cartesian skepticism in the first place, and we are no less entitled to these resources in a scientific context than any other.[4] Again Cartesian fantasies simply *replace* our worry about scientific underdetermination with a quite different (perhaps insoluble, but familiar) general skeptical problem.

A similarly subtle change of subject arises with the demand that we consider the 'Craigian reduction' of a theory (i.e., a statement of that theory's observable consequences) as a competitor when trying to assess the plausible threat of underdetermination. Perhaps even Craigian reductions are "real theories," but the worry was that there might be too many different accounts of the inaccessible workings of nature well-confirmed by the evidence, *not* simply that there are (as we already knew) multiple options for beliefs about the world that the evidence leaves us free to accept. Agnosticism about all accounts of the inaccessible aspects of nature is always *possible*, but is *defensible* only if the underdetermination of theory by evidence (or some other ground for suspicion about all theories)

3. Kukla sometimes appreciates this point (1996, 158), but not how it undermines his case (see below).

4. The need for such judgments will not alone evade underdetermination, however, for the prior plausibility of electrons, phlogiston, or curved spacetime is simply not on par with that of Cartesian Evil Demons (cf. Van Fraassen 1980, 36). I suspect that this difference is what is really at issue in the (misleading) claim that some scenarios are too farfetched to constitute "real theories" at all (e.g., Leplin and Laudan 1993, 11).

is independently established, for we surely want the *strongest* set of beliefs to which we are entitled by the evidence. It is not enough that the epistemically more modest choice to believe only a theory's claims about observable phenomena is always left open by the evidence (cf. Van Fraassen 1980); for that matter, so is choosing to believe nothing at all.

In contrast to the global strategy's Cartesian fantasies, the local algorithmic strategy seeks instead to take advantage of one or more formal features of a particular theory to show that an infinite or indefinite number of serious scientific empirical equivalents to that theory can be produced by varying the feature(s) in question. Consider the now-famous example of TN(0): Newtonian mechanics and gravitational theory, including Newton's claim that the universe is at rest in absolute space. This theory supports any number of empirical equivalents of the form TN(v), where v ascribes some constant absolute velocity to the universe.

But such empirical equivalents prove too little. The sensible realist will surely insist that we are not here faced with a range of *competing* theories making identical predictions about the observable phenomena, but instead just a *single* theory being conjoined to various factual claims about the world for which *that very theory* (along with the auxiliary hypotheses we accept) implies that we cannot have any empirical evidence.[5] It is not always trivial to determine which elements of a proposed theory are otiose by its own lights, but the sensible realist will counsel realism only about those theoretical claims (*whatever* they are) that our theories themselves imply are amenable to empirical investigation. This realism should no more extend to the *conjunction* of Newtonian theory with claims about the absolute velocity of the universe than with claims about the existence of God.

Another way to see this point is to note that empirical equivalents of the TN(v) variety pose no threat to the *approximate* truth of our theories: if the realist believes TN(0) when one of the various TN(v) obtains, *most* of her theoretical beliefs about the relevant domain will be straightforwardly true. Thus, empirical equivalents of the TN(v) variety show at most that we would have been unjustified in taking any stand on the constant absolute velocity of the universe, not in accepting the other theoretical claims of Newton's theory.

5. Although this implication requires auxiliary assumptions, these will be the same assumptions needed to assert the empirical equivalence of the various TN(v). Of course, changes in the accepted auxiliary assumptions over time may defeat the claim of empirical equivalence (a central point in Laudan and Leplin's [1991] attack on underdetermination), but we are here concerned with what to make of the prospect of theories that are empirically equivalent given (or 'indexed to'; see Kukla 1993) a particular set of auxiliary assumptions, or alternatively, with empirically equivalent "global theories" or "systems of the world" (see Hoefer and Rosenberg 1994).

Our response to the local algorithmic strategy, like the global, applies equally well to some famous non-algorithmic examples. John Earman suggests (drawing on results from Clark Glymour and David Malament), for example, that underdetermination threatens because "even idealized observers who live forever may be unable to empirically distinguish hypotheses about global topological features of some of the cosmological models allowed by Einstein's field equations for gravitation" (1993, 31). But such claims about global topology—concerning, for example, the compactness of space (as determined relative to some canonical foliation of spacetime)—are simply factual claims about the world for which the General Theory of Relativity itself (again, given accepted auxiliary hypotheses) suggests that we are (or may be) unable to acquire evidence. And there is surely something pathological about the claim that hooking one or another such claim to the General Theory of Relativity produces genuinely distinct, empirically indistinguishable theories: once again, the sensible realist will surely counsel realism only about those aspects of well-confirmed theories that *those theories themselves* hold to be empirically significant.

This suggests that the local strategy (like the global) actually trades in underdetermination for another long-standing philosophical problem, this time in the theory of confirmation: if true empirical consequences of a theory are all that matter to its confirmation, then evidence E confirming theory T will equally well confirm theory T + C (where C is any further claim that does not undermine T's implication of E), thus offering spurious confirmation to C itself. Like Cartesian skepticism, this problem is philosophically serious (indeed, it requires some solution), but it cannot be our only reason for taking underdetermination seriously without simply collapsing the latter problem into the former.

There are, of course, *examples* of empirical equivalents which are neither skeptical fantasies nor trivial variations on a single theory. Perhaps most convincing is Earman's other supporting case: "TN (*sans* absolute space) . . . opposed by a theory which eschews gravitational force in favor of a non-flat affine connection and which predicts exactly the same particle orbits as TN for gravitationally interacting particles" (1993, 31; see also Glymour 1977). Neither theory is a skeptical fantasy, nor is one a trivial variation on the other: treating gravitational attraction as a fundamental force seems substantially different from treating it as manifesting the curvature of spacetime. Other plausible (albeit more controversial) cases include Special Relativity versus Lorentzian Mechanics (controversial because the latter's requirement of systematic expansion and contraction for all our measuring devices (including rods and clocks irrespective of internal composition or construction) when in motion relative to absolute space might be thought a skeptical fantasy) and Bohmian hidden variable versus standard Von Neumann-Dirac formulations of Quantum Mechanics (con-

troversial because it is not clear that we understand Quantum Mechanics well enough to say convincingly what formulations of it count as genuinely different theories).[6]

Of course, the convincing examples are drawn exclusively from the physical sciences, an idiosyncrasy which might lead us to suspect that they form an unrepresentative sample and/or that there is something about the characteristic structure of physical theories (if such there be) which renders them especially susceptible to empirical equivalents—biologists and philosophers of biology, for example, have no idea how they would go about constructing even one (genuinely distinct, nonskeptical) alternative to the modern synthesis of Darwinian evolutionary theory and Mendelian genetics. Much more importantly, however, *none* of these examples is generated by an algorithm or formula and each is a hard-won *particular* alternative to an existing theory (rather than an infinite or indefinite collection) that proved quite difficult to identify and characterize: surely one or even a few such convincing cases do not provide sufficient warrant for concluding that genuine or serious empirical equivalence is a ubiquitous phenomenon! If numerous serious empirical equivalents to virtually any theory could be produced with just a little work and ingenuity, this would certainly ground the worry that an infinite space of alternatives looms over each of our best-confirmed scientific theories, but the profound difficulties and rare success we have encountered in trying to develop even one or a few convincing examples of nonskeptical and genuinely distinct empirical equivalents might sensibly be seen to support just the opposite general conclusion.

Thus, the case for underdetermination from empirical equivalents will simply not support the intoxicating morals that advocates hoped to draw. Algorithms provide proofs of the underdetermination predicament only by transforming the problem into one venerable philosophical chestnut or another, while one or a few convincing examples, dearly purchased and drawn from a single domain of scientific theorizing, are unable to support the sweeping conclusion that there are likely serious empirical equivalents

6. While Eddington, Reichenbach, Schlick, and others have famously agreed that General Relativity is empirically equivalent to a Newtonian gravitational theory with compensating "universal forces," the Newtonian variant has never been given a precise mathematical formulation (the talk of universal forces is invariably left as a promissory note), and it is not at all clear that it can be given one. (David Malament has made this point to me in conversation.) The "forces" in question would have to act in ways no ordinary forces act (including gravitation) or any forces could act insofar as they bear even a family resemblance to ordinary ones: in the end, such "forces" are no better than "phantom effects" and we are left with just another skeptical fantasy. At a minimum, defenders of this example have not done the work needed to show that we are faced with a credible case of nonskeptical empirical equivalence. Were this example to be accepted as genuine, however, it would not affect the status of my conclusion below.

to most theories in most domains of scientific inquiry. Scientists and philosophers concerned with a particular theory should surely worry about whether *that theory* has genuine empirical equivalents, but critics of underdetermination are well within their rights to demand that serious empirical equivalents to a theory actually be produced before they withhold belief in it, refusing to *presume* that such equivalents exist when none can be identified (see Kitcher 1993; Leplin 1997).

3. Recurrent, Transient Underdetermination and a New Induction over the History of Science. Of course, searching for empirical equivalents was only the most promising strategy for trying to *prove* that underdetermination obtains. It is therefore alarming that the connection between the two issues has become so firmly established that the most influential (and ostensibly general) recent attack on underdetermination (Laudan and Leplin 1991) and its most influential (and ostensibly general) recent defense (Earman 1993) *both* proceed *solely* by addressing the existence and status of putative empirical equivalents.[7] Furthermore, the lack of any convincing case for the widespread existence of genuine empirical equivalents simply does not settle the seriousness with which we should regard the threat of underdetermination. To see why, notice that Duhem's original worry did not concern the possibility that we might identify empirical equivalents to our best theories that are indistinguishable by any *possible* evidence at all; it was instead that there might simply be garden-variety alternative hypotheses, *not yet even imagined or entertained by us*, but nonetheless consistent with or even equally well-confirmed by all of the *actual* evidence *we happen to have in hand.* Following Sklar (1975), we might call this a *transient* underdetermination predicament: one in which the underdetermined theories are empirically *in*equivalent and could therefore be differentially confirmed by the accumulation of further evidence. Little-noticed in the crossfire over empirical equivalents is the fact that even such a transient underdetermination predicament undermines our justification for believing present theories, so long as we have some reason to think that it is also *recurrent*: that is, that there is (probably) at least one such alternative available (and thus this transient predicament rearises) *whenever* we must decide whether to believe a given theory on the strength of a given body of evidence.

The tough question, once again, is how to decide whether this predic-

7. Kukla (1993, 5–6) accuses Laudan and Leplin of presuming that the case for underdetermination rests upon empirical equivalents alone. Leplin and Laudan (1993, 16) deny this, but insist that their joining of the two doctrines was "not capricious," for "the philosophers whose derogations of the epistemic enterprise we have been concerned to redress (e.g., Quine and Rorty) . . . come to [the underdetermination thesis] . . . through [their belief in empirical equivalents]."

ament of recurrent, transient underdetermination is our own, that is, whether there really *are* typically unconceived competitors to our best scientific theories equally well-confirmed by all the actual evidence we have in hand. Sklar represents the most notable exception to the current general neglect of the threat of recurrent, transient underdetermination, but he finds it reasonable to simply *assume* in light of "the limitations of our scientific imagination" (and, at least in part, "reflection upon historical scientific experience," a suggestion I will try to flesh out below) "that there are vast numbers of perfectly respectable scientific hypotheses . . . we just haven't yet brought to mind," including "innumerable alternatives to our best present theories . . . which would save the current data equally well" and probably even some "more plausible than our own theories relative to present observational facts" (1981, 18–19). Elsewhere (1975, 381) he simply supposes without defense that even those who are skeptical of empirical equivalence "are likely to admit that transient underdetermination is a fact of epistemic life."[8] But of course, *these are just the claims with which critics of underdetermination will take issue on any nontrivial reading* (see Kitcher 1993; Leplin 1997).

While it is obviously difficult to acquire convincing evidence regarding the likely existence of presently unconceived theories, I think there is a genuine argument to be made out of Sklar's brief but tantalizing suggestion that the history of science itself bears on this question. Indeed, I will suggest that the historical record of scientific inquiry provides compelling evidence that recurrent, transient underdetermination is our actual epistemic predicament rather than a speculative possibility. Moreover, I submit that this very historical record contradicts Sklar's further suggestion (1981, 22ff.) that the threat can be substantiated only for fundamental physical or cosmological theories, as well as his more recent efforts to soften the sting of this worry with the suggestions that the historical progression of such theories is largely "one in which each successor theory is framed in concepts that are refinements and deepenings of the concepts of the theory that preceded it" (2000, 94) and that well-confirmed past theories can be seen as having been "pointing towards" later alternatives or "heading in the right direction" (2000, Ch. 4, Section I, passim.).

Of course, the most famous effort to extract philosophical mileage from the history of science is Laudan's (1981) Pessimistic Induction, which notes simply that past successful theories have turned out to be false and suggests that we have no reason to think that present successful theories will

8. Sklar is, of course, in good company: Quine's classic (1975) paper, for example, so often cited as providing evidence for an important underdetermination predicament, simply blusters, "Surely there are alternative hypothetical substructures that would surface in the same observable ways" (313).

not suffer the same fate. By contrast, I propose the following New Induction over the History of Science: 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 the subsequent history of inquiry has routinely (if not invariably) revealed 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.[9] For example, in the historical progression from Aristotelian to Cartesian to Newtonian to contemporary mechanical theories, the evidence available at the time each earlier theory was accepted offered equally strong support to each of the (then-unimagined) later alternatives. The same pattern would seem to obtain in the historical progression from elemental to early corpuscularian chemistry to Stahl's phlogiston theory to Lavoisier's oxygen chemistry to Daltonian atomic and contemporary physical chemistry; from various versions of preformationism to epigenetic theories of embryology; from the caloric theory of heat to later and ultimately contemporary thermodynamic theories; from effluvial theories of electricity and magnetism to theories of the electromagnetic ether and contemporary electromagnetism; from humoral imbalance to miasmatic to contagion and ultimately germ theories of disease; from 18th Century corpuscular theories of light to 19th Century wave theories to the contemporary quantum mechanical conception; from Hippocrates's pangenesis to Darwin's blending theory of inheritance (and his own 'gemmule' version of pangenesis) to Weismann's germ-plasm theory and Mendelian and contemporary molecular genetics; from Cuvier's theory of functionally integrated and necessarily static biological species or Lamarck's autogenesis to Darwinian evolutionary theory; and so on in a seemingly endless array of theories, the evidence for each of which ultimately turned out to support one or more unimagined competitors just as well. Thus, the history of scientific inquiry offers a straightforward inductive 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.

We cannot respond to these examples by noting that theories in the same general family or category as a later alternative (say atomism) sometimes *had* already been entertained by the time of an earlier theory's ex-

9. Strictly speaking, of course, the case for recurrent, transient underdetermination requires only that there have routinely been (nonskeptical, nontrival) unconsidered alternatives *not ruled out by* the evidence. I would suggest, however, that the stronger existence claim of unconsidered alternatives at least roughly equally well-confirmed by the available evidence is historically defensible, and it deflects any suggestion that such alternatives were ignored on evidential grounds rather than simply unconceived.

clusive dominance, for our confidence in the truth of our present theories cannot survive an inductive rationale for thinking that present evidence likely supports a presently unconceived detailed *version* of a theory from an existing family or type just as well as it supports the present alternative we accept on the strength of that evidence.

It will surely be objected, however, that in at least some of these cases changes in accepted auxiliary hypotheses were required before the alternatives could rightly be regarded as equally well-confirmed by the available evidence as the accepted theory. This is so, but misses the point that in such cases the needed alternative auxiliary hypotheses (often unconceived at the time) are typically *themselves* ones for which the available evidence provided equally compelling support. In other words, the *totality* of evidence available at the time of an earlier theory's acceptance typically offers equally compelling support for the *combination* of a later accepted alternative to that theory together with the requisite changes in auxiliary hypotheses that would later be accepted. And surely such a combination must be regarded as a scientifically serious alternative possibility, rather than a mere skeptical fantasy, for it is ultimately accepted by some actual scientific community.

Of course, a theory need not explain or accommodate *all* the existing data in order to be well-confirmed: evidential anomalies are allowed. The point is that we have repeatedly been able to conceive of only a single theory that was *well*-supported by all of the available evidence when there were indeed alternative possibilities *equally* well-supported by that evidence. Nor, therefore, does this argument ignore the phenomenon of explanatory losses in the transition from an earlier theory to a later one: a theory need not explain everything that a competitor explains in order to be as well-supported by the totality of available evidence: the theories may simply have *different* evidential anomalies. The judgment that alternatives not yet conceived were at least roughly as well-supported by the available evidence as earlier competitors will, however, require us to reject the most radical forms of Kuhnian incommensurability, on which the very phenomena themselves literally do not *exist* in any way that permits their identification across theories or theoretical paradigms[10] (although the triumph of such radical incommensurability would seem to offer scant comfort to realist critics of underdetermination in any case).

Thus, a full defense of the New Induction will require not only careful elaboration of historical cases, but also replies to any number of worries

10. It *is* sufficient to ground the New Induction, however, if we grant that the later theory is confirmed by the earlier phenomena *as those phenomena would be described or conceived by the later theory itself* just as well as the earlier theory was confirmed by those same phenomena under its own description or conception of them.

that arise in trying to compare the confirmation of distinct theories (including the account of confirmation at issue and radical incommensurability of evidential standards as well as phenomena). Perhaps most important of all, it will require a response to the recent realist reading of the historical record, which claims that the available evidence confirmed only those *parts*, *features*, or *aspects* of past theories which have turned out to be *true*. But it is well worth noting that a number of the most influential objections offered against the classic Pessimistic Induction simply do not weigh against the New Induction I propose. The New Induction is not open, for example, to the objection that its inductive basis includes theories drawn from the 'immature' periods of sciences and/or theories that did not enjoy some particularly important *kind* of empirical success (typically success in predicting *novel* phenomena), for its point is not that past successful theories were ultimately found false or otherwise wanting in some way, but instead that they were at one time *the best or only theories we could come up with*, notwithstanding the *availability* of equally well-confirmed alternatives. To pursue this line of criticism against the New Induction, then, opponents would have to argue that achieving a 'mature' theory or one that enjoys the right kind of success somehow and suddenly enables our imaginative capabilities to exhaust the space of serious possibilities.[11]

The New Induction will nonetheless disappoint a great many champions of underdetermination, for the historical record offers at best fallible evidence that we occupy a significant underdetermination predicament, rather than the sort of proof that advocates have traditionally sought (and I have been unable to do more here than suggest that this is indeed the verdict of the historical record). Furthermore, unlike constructing empirical equivalents, it does not allow us to say just *which* actual theories are underdetermined by the evidence, nor anything about what the (unconceived) competitors to present theories look like. But I have tried to suggest that empirical equivalents have proved to be a Devil's bargain for advocates of underdetermination—providing convincing evidence of an underdetermination predicament only where they have transformed the problem into one or another familiar philosophical puzzle—and I would

11. I would suggest that the distinctively holist version of underdetermination (see note 1) is also better defended by appeal to an analogue of the New Induction (asserting the historical ubiquity of plausible modifications of theories and challenges to background assumptions, typically unconceived in advance of anomalous evidence) than by familiar Cartesian strategies (e.g., pleading hallucination) or a shrill insistence that the legitimacy of absolutely any such modification or challenge is universally subject to "social negotiation," a claim which social constructivist case studies have never rendered remotely plausible and which threatens to bury any important holist case for underdetermination in a point of no more than Cartesian epistemological significance.

suggest that we start worrying instead about the kind of underdetermination that the history of science reveals to be a distinctive and genuine threat to even our best scientific theories. At the end of the day, clearly distinguishing a distinctive problem of scientific underdetermination from such familiar worries as Cartesian skepticism and spurious confirmation would be of little significance if not for the fact that the historical record offers compelling grounds for concern about this distinctive predicament in the case of scientific theorizing that survive solution (or even outright dismissal) of Cartesian fantasies, spurious confirmation, and the handful of convincing empirical equivalents we have managed to produce.

REFERENCES

Duhem, Pierre ([1914] 1954), *The Aim and Structure of Physical Theory*, 2nd ed. Translated by P. Wiener. Originally published as *La Théorie Physique: Son Objet, et sa Structure* (Paris: Marcel Rivièra & Cie). Princeton: Princeton University Press.

Earman, John (1993), "Underdetermination, Realism and Reason", *Midwest Studies in Philosophy* 18: 19–38.

Glymour, Clark (1977), "The Epistemology of Geometry", *Nous* 11: 227–251.

Hoefer, Carl and Alexander Rosenberg (1994), "Empirical Equivalence, Underdetermination, and Systems of the World", *Philosophy of Science* 61: 592–608.

Kitcher, Philip (1993), *The Advancement of Science.* Oxford: Oxford University Press.

Kukla, Andre (1993), "Laudan, Leplin, Empirical Equivalence and Underdetermination", *Analysis* 53: 1–7.

——— (1996), "Does Every Theory Have Empirically Equivalent Rivals?", *Erkenntnis* 44: 137–166.

Laudan, Larry (1981), "A Confutation of Convergent Realism", *Philosophy of Science* 48: 19–48.

Laudan, Larry and Jarrett Leplin (1991), "Empirical Equivalence and Underdetermination", *Journal of Philosophy* 88: 449–472.

Leplin, Jarrett (1997), *A Novel Defense of Scientific Realism.* New York, Oxford University Press.

Leplin, Jarrett and Larry Laudan (1993), "Determination Undeterred: Reply to Kukla", *Analysis* 53: 8–16.

Quine, Willard Van Orman (1975), "On Empirically Equivalent Systems of the World", *Erkenntnis* 9: 313–328.

Sklar, Lawrence (1975), "Methodological Conservatism", *Philosophical Review* 84: 384–400.

——— (1981), "Do Unborn Hypotheses Have Rights?", *Pacific Philosophical Quarterly* 62: 17–29.

——— (2000), *Theory and Truth.* New York: Oxford University Press.

Van Fraassen, Bas (1980), *The Scientific Image.* Oxford: Clarendon Press.