CHAPTER 6

ESCAPE FROM THE HUMEAN PREDICAMENT

As we discovered in the previous chapter, a number of critics contend that Quine has no serious reply to the skeptic, and that his view itself leads to skepticism. Quine has an interesting response to this attack: because skeptical doubts themselves arise within science, we can conclude that the skeptic is “overreacting.” The aim of this chapter is to propose a new anti-skeptical strategy inspired by this response, one that can even be used against skeptical elements of Quine’s own philosophy, and at least in part against other prominent forms of skepticism in the empiricist tradition.

One of Quine’s examples of the scientific nature of skeptical doubt concerns forms of the argument from illusion. Quine points out that these cases presuppose an understanding of “veridical” perception of reality, in contrast to which some cases are understood as “illusory.” Another example is the doubt stemming from the discovery of the retinal image, and the worry about how three-dimensional awareness can result from projection on a two-dimensional surface. Quine points out that even this worry presupposes an understanding of a three-dimensional eye in a three-dimensional world.

As I have noted in chapter 5, it looks like Quine does not take these points to establish that skepticism is somehow incoherent. He writes that the skeptic “is quite within his rights in assuming science in order to refute science; this, if carried out, would be a straightforward argument by *reductio ad absurdum*” (1975b, 68). Quine’s claim that he is only criticizing the skeptic’s overreaction (rather than any incoherence) leads critics such as Stroud (1984) and Williams (1996) to suggest that Quine is trying to use science to defuse the skeptic’s *reductio*: if the skeptic has the right to assume science at the outset of a *reductio*, then anti-skeptic also has the right to show how *full* consideration of science blocks this *reductio*, by showing that skeptical absurdities are merely apparent, arising only from scientific assumptions made in isolation from others.

But in chapter 5, I suggested that Stroud and William's interpretation of Quine is incorrect. If Quine intended to respond by defusing the skeptic's *reductio*, then by his own terms it would be an inadequate response, because he takes some of the most important discoveries of the naturalized epistemologist to be that reference is inscrutable and theory underdetermined by evidence. These theses by themselves seem to engender skepticism. If anything this makes the findings of science sound more, not less, absurd. As it happens, I argue that Quine does not mean to dissolve skeptical doubts so much as suggest that they do not matter, because we can always "acquiesce" in our own language and theory, and find consolation in pragmatic rather than logical forms of justification.

However Quine's pragmatism is not necessarily satisfactory. A traditional epistemologist will of course lament the abandonment of traditional goals of explaining the logical justification of science, and it is also not clear if a pragmatic criterion can succeed on its own terms in privileging science in favor of other forms of useful human discourse. So it is important to explore whether the skeptical theses that motivate Quine's pragmatism are true in the first place. At the end of the last chapter, I argued that the indeterminacy of translation and inscrutability of reference theses could each be reduced in large part to versions of the underdetermination thesis. Now indeterminacy, inscrutability and underdetermination are not the only points in the case for naturalizing epistemology. As we saw when we examined Kim's proposed alternatives to naturalized epistemology, behaviorism and extensionalism also help to eliminate traditional epistemological projects and help leave naturalism as the only viable proposal. In the previous chapter, I mentioned in passing strategies that might help undercut both behaviorism and extensionalism, but said I would focus on the underdetermination thesis in the current chapter. I will focus on it by asking the following question: Even if Quine did not intend to dissolve skeptical doubts by pointing to the scientific nature of these doubts, what if *we* were to use that strategy against the very skeptical theses, in particular against the underdetermination thesis, the embrace of which necessitates Quine's acquiescence in skepticism? What if we can abandon the underdetermination thesis itself by determining that, while the thesis appears to arise from some important scientific facts, it dissolves in the light of the wider context of scientific findings?

In this chapter, I hope to show how this strategy can be deployed against the underdetermination thesis by drawing on existing anti-skeptical arguments from a number of philosophers. Drawing on material from Larry Laudan and Jarrett Leplin among others, I present the assumptions about the nature of scientific practice behind the underdetermination thesis, and show what further scientific assumptions may be appealed to in order to defuse it. In my second section, drawing on material from John Norton, I examine a strategy for dissolving another source of radical skepticism, the classical Humean problem of induction, based on consulting the wider context of scientific practice. In my final section, drawing on material from J.J. Gibson and Jesse Prinz, I show how a final anti-skeptical strategy can be deployed against the anti-foundationalism which motivates inductivist skepticism. In each case, I show how skeptical doubts can be opposed by the free use of science within the scope of skeptical reductios, allowing us to escape from the "Humean predicament."

There will be just one catch. Finding the scientific source of doubts is sometimes only a partial solution to skeptical arguments. I say that because in some cases, I think skeptical doubts arise from scientific evidence only in conjunction with particular philosophical assumptions, and these assumptions are not always derived directly from science themselves. As we move on to more and more fundamentally skeptical problems, such as the problem of induction and anti-foundationalism, the doubts that arise will become more philosophical in nature. In these cases, I will try to identify and, in some cases, refute erroneous philosophical assumptions. But I do not need to refute them completely to make my point: the presence of purely philosophic, non-scientific assumptions in the case for naturalizing epistemology shows, I think, that Quine is not entirely correct that skeptical doubts are always scientific ones. But if that is the case, it follows that to the extent that he has relied on any of these doubts to motivate his naturalism and pragmatism, he has also relied on non-naturalistic presuppositions, which would undermine the idea that philosophy can ever be entirely continuous with natural science.

This discussion will, therefore, address the "reciprocal containment" of science and epistemology in three ways. First, showing that the underdetermination thesis and its assumptions dissolve in light of a wider context of scientific evidence will help to show that naturalized epistemology is not contained in natural science in the way Quine has suggested. Second, showing that the case for naturalizing epistemology relies on non-scientific philosophical assumptions will show that naturalized epistemology is not contained *exclusively* by natural science. This means that we can accept science without feeling obliged to abandon epistemological tradition. Finally, as promised in the last chapter, our examination will suggest that while an understanding of the value of science is contained in an empiricist epistemology, a reformed empiricism (one that conforms to wider context of scientific information to be examined) need not champion scientific empirical knowledge as the *only* kind of empirical knowledge. There will be room for first-person pre-scientific knowledge, of the sort that will underpin a foundationalism that challenges the need for naturalizing epistemology.

Understanding the scientific roots of the underdetermination thesis

What exactly is the underdetermination thesis, and why does Quine accept it? Philosophers sometimes distinguish between a "weak" and a "strong" underdetermination thesis (see Newton-Smith 2001). The weak version asserts that it is possible to construct empirically equivalent theories, i.e., theories that have the same empirical consequences or which can be consistent with the same body of empirical evidence. This is usually seen as uncontroversial in comparison with the strong version, which says not only that it is possible to construct empirically equivalent theories, but *also* that these empirically equivalent rivals are also equally well-confirmed (or well-supported or well-justified) by observational evidence.[49] Strictly speaking, I don't think it is proper to call the first thesis an "underdetermination" thesis at all, because it says nothing about *determination* by evidence. The second and distinctive part of the strong thesis does say that. The strong thesis says, in effect: *because* evidence can consistent with multiple theories (there are empirically equivalent rivals), therefore theories are not determined by observational evidence alone (the empirically equivalent rivals are equally well-confirmed by the evidence).

This fuller statement of this thesis takes the form of an argument ("because . . . therefore"), but it is perfectly natural to differentiate certain important conclusions by reference to the distinctive premises that lead to them, particularly when there are different ways to come to the same conclusion. There may be many other philosophers' reasons for thinking that theories cannot be determined by evidence. Traditional Pyhronnian skepticism is one example. The underdetermination thesis comes to same conclusion by way of the point about the existence of empirically equivalent rivals. In attacking the thesis, I will attack the argument implicitly contained within it: the idea that the possibility of empirical equivalence implies the impossibility of evidential determination of theories.

A note on my use of the term "confirmation": In what follows, I am using the term as a generic placeholder term referring to non-pragmatic logical support (whether deductive or inductive). Sometimes advocates of the underdetermination thesis, who will insist that empirically equivalent theories are not logically determined by the evidence, will also say that because of pragmatic differences between the theories, one is still better "confirmed" than the other. For example, in "The Scope and Language of Science" (2004e, 198), Quine says that simplicity, one of his pragmatic virtues

[49] Notice that this problem is much stronger than the traditional problem of induction. It is not a problem about the logical gap between limited observations and the unlimited scope of generalizations. Rather, it is a logical gap between *all possible observations* and theory. According to the underdetermination thesis, no matter how much evidence we collect for a theory, there will always be an empirically equivalent rival that will account for it just as well. Arguing that the problem of underdetermination thesis arises from special problems concerning the scientist's method of hypothesis—which we will shortly examine—Larry Laudan calls Hume's problem the problem of "plebeian induction," whereas underdetermination is the problem of "aristocratic induction" (Laudan 1981, 73).)

of theory choice, in some sense counts as "a kind of evidence;" in "On Simple Theories of a Complex World" (2004c, 211) he suggests that perhaps "the simpler hypothesis stands the better chance of confirmation."[50] So there are pragmatic and non-pragmatic senses of "confirm." In discussing formulations of the underdetermination thesis, I will use the non-pragmatic sense, because that is often how it is described in the literature, including by Quine himself in "Two Dogmas of Empiricism" (1953b). There Quine *denies* that individual statements "admit of confirmation or infirmation at all" and therefore may be accepted as true come what may, or revised at any time, provided that the appropriate auxiliary hypotheses in one's holistic theories are adjusted (41). As we shall see shortly, this provides Quine's most important argument for the underdetermination thesis (the statement of which follows almost immediately in "Two Dogmas").

So even though it is possible to deny that empirically equivalent rivals are necessarily equally well-confirmed in the pragmatic sense, this is consistent with the idea that theory is not determined by observational evidence, which is what makes the underdetermination thesis an effective tool in motivating naturalized epistemology. This is because the idea that the evidence does not help logically (deductively or inductively) determine theory is cause enough for concern among traditional epistemologists. If we hold one theory, while we know another actual or even *possible* rival theory is equally well-confirmed by the evidence, it follows that the first theory is not evidentially superior to the rival, and it would seem to be irrational to accept it (Bergstrom 1993), at least as long as evidence is our only measure of rationality or epistemic justification.[51] Even if we take pragmatic factors such as simplicity or predictive power as further components of rationality, it not uncontroversial to wonder why pragmatic rationality is relevant to the *truth* of our hypotheses—especially because of the difficulty involved in giving an objective account of the pragmatic values involved in this kind of "rationality."

---

[50] See also Bergstrom (1990, 44).
[51] There are further complications for how underdetermination leads to skepticism, depending upon whether or not the original theory and its rival are taken to be logically compatible. For details on how these complications are handled, see Bergstrom (2004) and chapter 5.

Having clarified the meaning of the underdetermination thesis and the meaning of some of its terms, we are now in a position to see why Quine and others think it is true. First we need to consider reasons for accepting the premise that there are empirically equivalent rivals for every theory. Then we need to consider whether the conclusion about evidential determination follows from it.

There are at least two different reasons for believing in the existence of empirically equivalent rivals. These are not always articulated by Quine, but have been made explicit by a Quine commentator, Lars Bergstrom (1990), without Quine's objection (1990). The first reason Bergstrom calls the "argument from logic": "a given set of sentences is implied by different, and even logically incompatible sentences; hence an empirical content is implied by different theory formulations."[52] According to this argument, any premises we can construct which imply the same empirical consequences are empirically equivalent rivals. Because we can construct them for any theory, empirically equivalent rivals exist.

The second reason to believe in the existence of empirically equivalent rivals is what Bergstrom calls the argument from confirmation holism. Confirmation holism is the thesis that Quine, in "Two Dogmas of Empiricism," claims to share with Duhem: the idea that it is only whole blocks of theory, not individual hypotheses, that are subject to confirmation or disconfirmation.[53] Confirmation holism results from the fairly obvious point that hypotheses imply empirical consequences only with the assistance of auxiliary hypotheses, and the point that in the face of new observations, one can always either retain or reject the hypothesis in question, depending upon whether one rejects or retains associated auxiliary hypotheses. Each possible set of hypotheses and auxiliaries consistent with observational evidence is itself a block of theory, and since there are many such sets possible given the observational evidence, there are therefore many empirically equivalent rivals. I note that the argument

---

[52] One might, for example, argue for "If A then B" with either set of premises: "If A then C, If C then B," or "If A then D, If D then B." Since "If A then B" could be understood as an empirical consequence--Quine himself treats empirical consequences as "observation conditionals" or "observation categoricals"--this shows how the same empirical consequence can derive from different theoretical premises.

[53] It does appear, however, that Quine was originally incorrect to ascribe this view to Duhem. See Ariew (1984).

from confirmation holism is more attuned to scientific practice than the argument from logic, because it suggests that empirical equivalents are generated not by completely artificial means, but from transformations of existing bodies of scientific belief. Quine himself (1975a, 313) draws the connection between confirmation holism and underdetermination:

> This holism thesis lends credence to the under-determination theses. If in the face of adverse observations we are free always to choose among various adequate modifications of our theory, then presumably all possible observations are insufficient to determine theory uniquely.[54]

Of course the quote from Quine above is not just addressing the question of the existence of empirically equivalent rivals, but also their significance for the question of evidential determination. We are now in a position to see this significance in light of an obvious problem with the argument for alluded to so far: the conclusion that evidence does not determine theory simply does not follow from the premise that there are empirically equivalent rivals. Some intermediate premise is needed to connect the two points. The likely intermediate premise is the second in the following argument:

> 1. For every theory, there is an empirically equivalent rival.
> 2. Empirically equivalent rivals are equally well-confirmed.
> 3. Therefore, for every theory, there is an equally well-confirmed rival.

The important question, then, is why we should accept Premise 2.

As it turns out, Quine seems to provide his rationale for the premise by a naturalistic appeal to a fact about scientific practice: that scientists rely on the hypothetico-deductive method to confirm their theories. Quine's clearest statement of the connection between hypothetico-deductivism and underdetermination is found in his essay "On Empirically Equivalent Systems of the World" (1975a, 313):

> If all observable events can be accounted for in one comprehensive scientific theory— one system of the world . . . —then we may expect that they can all be accounted for equally by another, conflicting system of the world. We may expect this because of how scientists work. For they do not rest with mere inductive generalizations of their observations: mere extrapolation to observable events from similar observed events.

[54] See also Quine (1981b: 9)

> Scientists invent hypotheses that talk of things beyond the reach of observation. The hypotheses are related to observation only by a kind of one-way implication; namely, the events we observe are what a belief in the hypotheses would have led us to expect. These observable consequences of the hypotheses do not, conversely, imply the hypotheses. Surely there are alternative hypothetical substructures that would surface in the observable ways.

Quine's allegation that observations can be "accounted for equally" by more than one theory "because of how scientists work" is as close to an explicit statement of the source of Premise 2 as we are likely to find.[55]

It should come as no surprise that this is his view about how scientists work, as he has made the point and its signification for confirmation elsewhere in his body of work. In "The Nature of Natural Knowledge" (2004d, 291) he says that the essence of the scientific method is hypothetico-deductivism:

> Science is a ponderous linguistic structure, fabricated of theoretical terms linked by fabricated hypotheses, and keyed to observable events here and there. Indirectly . . . the scientist predicts future observations on the basis of past ones; and he may revise the superstructure when the observations fail. It is no longer simple induction. It is the hypothetico-deductive method.

And in "Naturalism; or, Living within One's Means" (2004b, 276), he describes the implications of this method for a view of confirmation:

> [T]he most we can reasonably seek in support of an inventory and description of reality is testability of its observable consequences in the time-honored hypothetico-deductive way—whereof more anon.

---

[55] This passage should also put to rest any worries that the underdetermination thesis is *prima facie* non-naturalistic, because it uses the concept of "confirmation" (or justification, or support) in alleging that multiple theories are equally well-confirmed. There is no reason to think that the *concept* of confirmation or justification is unacceptable to the naturalist. Even if the concept has been the subject of many a conceptual analysis, naturalists are perfectly happy to explicate the concept for their own purposes. Clearly they are happy to take both deductive and hypothetico-deductive justification as paradigmatic forms of logical justification, and by assuming that there are no other forms of logical justification, argue that theories are equally logically justified if they have the same empirical consequences. To be sure, they can widen their naturalistic concept of justification to include pragmatic justification, but that is *as a result* of their conclusion about underdetermination of theory by evidence, i.e., the equal logical justification of empirically equivalent rivals. If there is something non-naturalistic about the underdetermination thesis, it is not something about its explicit formulation. Rather, later I will argue that the non-naturalistic assumptions are buried within the justification for Premise 2.

And: “[T]he deduction and checking of observation categoricals is the essence, surely, of the experimental method, the hypothetico-deductive method” (280).

Quine himself eventually distances himself from the idea that the global underdetermination thesis is epistemologically significant, or even coherent (1975a; 1976; 1979). But he does seem to take thesis as significant in “Two Dogmas of Empiricism” (1953b, 42–3) where he writes:

> The totality of our so-called knowledge or beliefs . . . is a man-made fabric which impinges on experience only along the edges. . . . [T]he 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.

It is *this* brand of underdetermination that appears to inform his derivation of the indeterminacy of translation thesis and motivate his naturalized epistemology, both of which, as we have shown, are definitely epistemologically significant.

In any case, whatever Quine’s ultimate view about underdetermination or its sources, we can take his view about the connection between the method of hypothesis and underdetermination as representative of a view that has been widely accepted throughout the history of science. Quine is not the first observer to conclude that an underdetermination problem seems to follow from facts about scientific practice related to the hypothetico-deductive method. In his survey of the history of the method of hypothesis, *Science and Hypothesis* (1981), Larry Laudan argues that the idea of *local* underdetermination is at least as old as the 3rd century B.C., when Apollonius first showed that motion of a planet on an epicycle was empirically equivalent to its motion on an eccentric. This equivalence, of course, played a central role in the later controversy over geocentric and heliocentric models of the solar system. Laudan notes that early religious commentators on the heliocentric hypothesis, such as Aquinas, Agostino Nifo, and Osiander (in his preface to Copernicus’ *On the Revolutions*) all argued that because different hypotheses could imply the same astronomical consequences, the evidence by itself could not lead us to believe any of them as literally true—an early version of the argument for anti-realism from underdetermination, except as applied to only to a local domain.

Laudan argues in the first half of the 17^{th} century, when science began to hypothesize more and more unobservable entities, a more global skeptical problem began to arise. Some scientists of the late 18^{th} century (such as David Hartley and George LeSage) began to think of hypothetico-deduction as the exclusive source of scientific confirmation, and argued not even inductive generalization played the role originally defended by Newton and Reid. Underdetermination, once confined to the local domains to which hypothetico-deductivism was confined, such as astronomy, would now become more global. Of course none of these thinkers thought that all hypotheses were equally good, but they did think that they were equally well-confirmed by the evidence. As long as they regarded pragmatic criteria like simplicity as components of rationality, they did not become skeptics. Later developments would make this conclusion more difficult to avoid. And it was, perhaps, not a coincidence that Descartes, who championed the hypothetico-deductive method in science, also made an underdetermination-style skeptical argument concerning the evil demon (Okasha 2003).

Laudan emphasizes that these philosophers embraced the hypothetico-deductivist conception of confirmation because it was the conception—in contrast with inductivism—that best captured some of the most successful scientific theorizing of the day. They were unwilling, for example, to regard the wave theory of light as inadmissible simply because it posited unobservable entities (as Reid and other inductivists did), not when this theory met with such amazing predictive success (such as Fresnel and Poisson's prediction that a bright spot would appear under the shadow of a circular disk if the wave theory was true).These facts about successful scientific practice enriched their conception of confirmation—but paradoxically, the same conception also led them to regard the deliverances of science as subject to a form of skepticism.

We should, however, clarify just precisely how it is that hypothetico-deductivism implies Premise 2, that empirically equivalent rivals are equally well-confirmed. It is not enough that scientists sometimes find evidence in deduced empirical consequences. A stronger connection between hypothetico-deductive practice and the concept of confirmation is required. To say that two theories are empirically equivalent is to say that they have identical empirical content, i.e. the same empirical consequences. Stated loosely, hypothetico-deductivism is the idea that the empirical consequences of a hypothesis are a source of confirmation. But something stronger than that is needed for it to be the case that if two theories have the same empirical consequences, they therefore have the same degree of confirmation.

One way to support the idea that empirically equivalent rivals are *ipso facto* equally well-confirmed is to suppose that there is no way to confirm a given hypothesis H other than by observing *its* empirical consequences. In other words, observed empirical consequences of H are *uniquely* necessary for confirming H. After all, if other sources of confirmation are available, two empirically equivalent theories may have the same hypothetico-deductive confirmation, but differ in degrees of *other* forms of confirmation. In that case, empirical equivalence would not imply equal degrees of confirmation. Now as we shall see, it is highly implausible to think that a hypothesis can only be confirmed by *its* empirical consequences. Even the hypothetico-deductivist should acknowledge that a hypothesis might be confirmed in a straightforwardly deductive fashion, by deducing it *from* a better known proposition. But if that is the case and there is more to confirmation than observing a hypotheses' own empirical consequences, then more is needed to support Premise 2. *On its own*, Premise 2 does not seem to be true unless we also assume the implausible point we have just mentioned. Without that point, Premise 2 would be true only if the two theories also do not differ in their deductive confirmation, and it is an open question whether we can assume that. In the next section, we will consider arguments concerning whether or not the possibility of differential deductive confirmation is a problem for the underdetermination thesis. We will also consider whether there are forms of confirmation beyond deductive and hypothetico-deductive.

Premise 2 also presupposes a point that may seem much less controversial to many: that observing empirical consequences of a hypothesis *suffices* to confirm the hypothesis to some degree: once the empirical consequences of a hypothesis H are observed, this by itself is enough to provide H with some confirmation. So on this interpretation, if two hypotheses are empirically equivalent, they are both confirmed in some sense. If the mere ability to deduce an observable consequence from H did *not* suffice to confirm it, then the existence of equally confirmed rivals could not be assumed simply because of the construction of empirically equivalent rivals (through pure logic or through adjustment of auxiliaries): some empirical consequences might not be relevant to the confirmation of the hypothesis in question. So even if empirically equivalent rivals exist, and empirical consequences of hypotheses *often* confirm hypotheses, it would not follow that empirical consequences *always* confirm the hypotheses from which they are derived.

We now have a range of options before us concerning how to challenge the underdetermination thesis. We can challenge Premise 1 directly by challenging the idea that for any theory, there is always an empirically equivalent rival. And we can challenge Premise 2 in three separate ways: by challenging the idea that empirical equivalence implies an equivalence of deductive confirmation, by challenging the idea that hypothetico-deductive and deductive confirmation exhaust our forms of confirmation, and by challenging the idea that empirical consequences suffice to confirm. In the following section, I will argue that each of these premises can be called into question when we examine a wider context of scientific evidence that calls into question our original naturalistic sources of doubt.

Undermining underdetermination: the scientific roots in context

Premise 1: Are there always empirically equivalent rivals?

In an influential article, "Empirical Equivalence and Underdetermination," Larry Laudan and Jarrett Leplin (1991) challenge the idea that there *are* empirically equivalent rivals for every theory. They argue that theories with identical logical consequences are not necessarily identical in *empirical* consequences, because what it is to be an *empirical* consequence is relative to one's background theory, and varies from theory to theory. Changing auxiliary hypotheses in the manner envisioned by the confirmation holist may have the effect of changing a theory's class of empirical consequences, so it is not trivial that adjusting one theory when confronted will recalcitrant "observations" will result in another theory that is empirically equivalent. Much of the debate on underdetermination in the literature results from Laudan and Leplin's paper, and concerns this point about whether rivals generated by logical algorithm should count as empirically equivalent.[56] I should note, incidentally, that even if Laudan and Leplin's objection here misses the mark, it does show that what Bergstrom calls the "argument from logic" is irrelevant in the end to supporting Premise 1. The fact that we can find different hypotheses with the same logical consequences is unrelated to the question of underdetermination by *evidence,* because it does not address the question of *empirical* consequences. So clearly the argument from confirmation holism is the only serious argument for underdetermination.

John Norton (1994) raises another question about the existence of empirically equivalent rivals, this time more attuned to the naturalistic challenge of the argument from confirmation holism: "[O]ne rarely finds the scientist has such an embarrassment of riches. Typically, a scientist is pleased to find even one theory that is acceptable for a given body of evidence" (4). Of course it is possible that *some* actual empirically equivalent pairs may be found in the history of science. Alleged examples of these have included geocentric vs. heliocentric theories of the solar system, Newtonian mechanical theories differing in state of absolute rest, and Einsteinian relativity vs. Lorentz ether theory.[57] Of course not all of these remained empirically equivalent with the advance of theory, and even assuming that they offer actual examples of *local* underdetermination, it is harder to see that they threaten to establish any epistemologically significant, *global* underdetermination thesis.

[56] See Kukla (1993), Laudan and Leplin (1993), and Kukla (1996).

[57] Though even here, we note, with Laudan and Leplin (1991, 459) that it is interesting that all of the historical examples seem to stem from the relativity of motion. This strongly suggests that underdetermination effects can be isolated to a given domain of theory. This still leaves the question open as to how one is to choose among such rivals, but that open question would still seem to have little general epistemological significance.

In his later writings, Quine (1975a) became concerned that even if we find empirically equivalent rivals, they may not be the kind needed to generate a significant underdetermination thesis. This is because it is plausible that many cases that empirically equivalent theory formulations that are logically incompatible may be mere notational variants of each other, which can be rendered logically equivalent through an appropriate reconstrual of predicates. In the end Quine decided that there is no practical way for human beings to know that there are such reconstruals available, so rival theories may appear to be underdetermined. But he says it is also impossible to rule out the existence of such reconstruals, so whether theories are underdetermined amounts to an "open question." John Norton (2003b) makes a similar observation when he offers an interesting argument for why any empirically equivalent theory easy enough to construct in the space of a single journal article is probably just a notational variant of the original theory, not a genuine rival.

Whether or not there are empirically equivalent but logically incompatible *rivals* is an interesting question. If there are not, then not even the "weak" underdetermination thesis would be true; surely then the "strong" thesis would not be, either. However I think that the issues involved in resolving this debate relate little to dealing with the questions that originally gave rise to the underdetermination thesis. Here I have in mind more basic questions about the role of empirical consequences in confirmation in the first place. So we should ask about whether the existence of empirically equivalent rivals also implies that these rivals are equally well-confirmed, and question Premise 2. To do this, we should look to see if the wider context of scientific practice supports the hypothetico-deductivist conception of confirmation. We should ask: do scientists actually treat the deduction of an observed empirical consequence from a hypothesis as both a uniquely necessary and sufficient for confirming that hypothesis? I believe the answer is a strong "no," if counterexamples from the history of science presented by Laudan and Leplin (1991) and others are convincing.

Premise 2: Equal deductive confirmation?

Challenging the admittedly implausible claim that observing the empirical consequences of a hypothesis is the *only* way to confirm that hypothesis, Laudan and Leplin cite a historical example about the development of the theory of continental drift, call it T. This theory holds that every part of the earth has at one time occupied a different location, and implies the following hypotheses: H_1: The climate of each part of the earth has differed in the past; H_2: The current magnetic alignment of each part of the earth will differ with the alignment of older magnetic rocks. Laudan and Leplin then cite the fact that in the 1950s and 1960s, evidence *e* for H_2 was acquired in the normal hypothetico-deductive way (H_2 entailed *e*, and *e*—a set of facts about changing magnetic polarities in adjacent sections of the mid-Atlantic trench—was then observed). Interestingly, however, *e* thereby supported H_1, even though *e* was not an empirical consequence of H_1. The climate hypothesis H_1 was supported not because of *its* empirical consequences, but because it was a consequence of the theory of continental drift, T, which was hypothetico-deductively supported by H_2. Further historical examples of this kind of "indirect" confirmation are cited: the discovery of Brownian motion supported the atomic theory; the observation of the increase of mass with velocity supported relativistic principles; Thompson's cathode ray experiments supported hypotheses about electricity; Maxwell's observations of heat radiation supported the kinetic molecular theory; and the discover of heritable variation supported Darwin's hypothesis about the age of the earth (1991, 462). In none of these cases were the new discoveries empirical consequences of the hypotheses in question.[58]

The significance of these counterexamples is that while in some cases (as in *e*'s support for H_2), a hypothesis is supported by its empirical consequences, in other cases (as in *e*'s indirect support for H_1 via a theory which implies H_1), a hypothesis can be supported by what it derives *from*. This is

[58] Okasha (1997) objects that Laudan and Leplin do not specify the circumstances in which evidence flows upwards and those in which it flows downwards (253-4). That may be true, but I do not think it is incumbent on Laudan and Leplin to perform this task, particularly if, as naturalists, they are concerned with scientific practice. It seems to be sufficient that scientists sometimes regard one direction of flow as evidential, and sometimes the other—and that they meet with great success when they do this.

what I have referred to as "deductive confirmation" in the previous section. Laudan and Leplin call it "indirect confirmation" because the H_1 is supported not by its empirical consequences, but indirectly by virtue of deducibility from T, which is in turn supported by *its* empirical consequences (H_2, which is supported by its consequence, e). Earlier I had said that it seems unlikely that naturalists would discount this as an additional form of confirmation. But the existence of alternate forms of confirmation then implies that Premise 2—that empirically equivalent rivals are equally well-confirmed—cannot be supported unless we assume the further premise that empirically equivalent rivals are not unequal in deductive confirmation. We should now consider whether there is any good reason to assume this.

One objection to Laudan and Leplin's argument, raised by Bangu (2006), raises a point that might suggest why equal deductive confirmation of rival hypotheses might be assumed—thereby supporting Premise 2 of the argument for underdetermination. Suppose that our rival hypotheses H_a and H_b both imply the same empirical consequences, E. We shall consider whether they can be said to be equally well-confirmed in virtue of this empirical equivalence. Now one could use deductive confirmation to confirm H_a better than H_b, but only if there is a theory T_a (supported hypothetico-deductively by its own evidence *e*) which implies H_a but does not imply H_b. Bangu objects that the existence of such a T_a would not solve the underdetermination problem, because nothing rules out the existence of another theory, T_b, supported likewise by evidence *e*, the same evidence supporting T_a, which *does* imply $H_{b,}$ but not H_a. So presumably whatever allows us to find the empirically equivalent pair, H_a and H_b also allows us to find a T_a and T_b which are empirically equivalent vis-à-vis evidence *e*. As long as there can be a T_a and T_b which imply different H's, we can assume, in effect, that H_a and H_b are equally well-confirmed just in virtue of being empirically equivalent. True, relative to T_a, H_a is better supported than H_b. But relative to T_b, H_a is better supported. Since there is no way to choose between T_a and T_b based on *e*, it seems that H_a and H_b are tied once again.

There are several responses to Bangu. First of all, I note that he admits that the existence of T_b is only a "logical possibility," not an inevitable fact of scientific practice. Yet I have emphasized that we should address ourselves to scientific practice to assess this question. If his response to Laudan and Leplin is to be a serious naturalistic source of underdetermination, there should be actual examples of T_a and T_b in the history of science—but none are presented. What's more, nothing about the logic of confirmation holism guarantees the existence of the higher-level theories needed to provide the necessary deductive confirmation for *any* hypotheses, to say nothing about multiple theories which differentially confirm rival hypotheses. At best, confirmation holism permits a way to find empirically equivalent hypotheses by permitting the arbitrary reshuffling of existing auxiliary assumptions. Once these have been reshuffled, it becomes deductively obvious what should be done with the original hypothesis. But finding new theories from which new hypotheses are to be deduced is quite another thing, involving not just the reshuffling of auxiliaries, but the discovery of new ones. Kepler's laws, for example, required Newton's discovery of calculus before they could be deductively confirmed by Newton's laws.

Second, Bangu considers an objection to his proposal which I believe he does not fully answer. Some would argue that that there could be no such thing as a T_b which is supported by the same evidence as supports T_a (that is, e), but which implies a different hypothesis than T_b (H_a rather than H_b).Bangu says it is trivial that two different theories, each entailing different hypotheses, can be supported by the same *e*. His example involves T_a, "All swans are white," and T_b, "All swans are white and all crows are black," each of which entail and are therefore supported by *e*, "All observed swans were white" (274).[59] There is a hypothesis, H_b, "All crows are black," which is entailed by T_b but not by T_a. Now although Bangu shows how it is possible for two different theories to entail different hypotheses while being supported by the same hypothetico-deductive evidence, he does not

[59] I note in passing that this T_b includes an irrelevant conjunct, which is not clearly supported by the evidence *e*. We will examine this issue later when discussing whether deriving an empirical consequence is sufficient for confirmation.

shown that the two different entailed hypotheses are *logically incompatible rivals*, which is the usual stipulation of an underdetermination problem. If H_b is "All crows are black," what H_a might be the different hypothesis supported by T_b? Presumably only "All swans are white." But "All swans are white" and "All crows are black" are not logically incompatible. So while the choice between H_a and H_b is underdetermined by hypothetico-deductive evidence, there is no reason to think a choice between them must be made! As we have already seen, two theories must be logically incompatible to be the kind of rivals that are interestingly underdetermined.

Presumably, however, it is a trivial matter to supply a pair of theories other than Bangu's which do imply incompatible hypotheses. Suppose, for example, that we selected as T_a, "All swans are white and All crows are black," and selected as T_b, "All swans are white and All crows are white." Each of these do indeed entail "All observed swans are white" and are therefore supported by it. But each can also be made to imply incompatible hypotheses: H_a, "All crows are black," and H_b, "All crows are white." Very well. But now there is a new problem. H_a and H_b achieve their incompatibility at the cost of no longer obviously entail the same evidence E, which is the original stipulation of the idea that they are empirically equivalent and therefore underdetermined by that evidence. Perhaps, however, this was a problem for Bangu's example to begin with, since it would also not have been obvious how "All swans are white" and "All crows are black" might have been entailed and supported by the same evidence. At least in that case, however, we were not dealing with hypotheses predicating opposite colors of the same species. Perhaps a story could be told about why white swans explain some piece of evidence, while black swans could also explain it. This overall objection—that Bangu's example does not meet all of the stipulations of the standard underdetermination problem—does not imply that there is no way to concoct the proper example. But it does show that it is not "trivially true" that two different theories, each entailing *logically incompatible* but *empirically equivalent* hypotheses, can nevertheless be supported by the same evidence. If it is not trivially true, then particularly because Bangu has produced no examples from actual scientific practice, there is no reason to think we can always assume that there always exist a T_a and T_b from which H_a and H_b might be deduced, rendering them equally confirmed and supporting Premise 2 (that empirically equivalent rivals are equally well-confirmed).

It is also noteworthy that Bangu assumes that the only source of support for T_a and T_b is itself of a hypothetico-deductive nature. This is a fair assumption, since Quine and other major parties to the discussion assume this. "Hypotheses," Quine says, are "related to observation *only* by a kind of one-way implication"—from hypothesis to observation, not vice versa. But it is worth thinking about whether we should assume this. If there are other ways for our theoretical knowledge to interface with experience, this would mean that there could be ways aside from hypothetico-deduction in which hypotheses can be empirically supported. If there is, then there is another way in which Premise 2 might fail: two hypotheses with identical empirical consequences might still be differently confirmed.

Premise 2: Alternative sources of empirical confirmation?

Laudan and Leplin themselves entertain one additional source of independent empirical support: analogy. They identify important examples of reasoning by analogy in the history of science. They point out how Maxwell drew an analogy between an enclosed gas and an enclosed set of colliding particles, supporting the molecular theory of gases, and how Einstein drew an analogy between the entropy of certain kinds of radiation with that of ideal gases, supporting a quantum view of radiation (1991, 464-465). Surely there are others. If full-fledged theories can be confirmed by analogy, then surely perhaps individual hypotheses might be, as well.

Of course analogy is not always an independent source of confirmation. An analogy drawn with the behavior of ideal gases would be only as strong as our knowledge of these gases, which undoubtedly requires much independent confirmation of its own. Consider, however, Maxwell's analogy between gas molecules and colliding macroscopic particles. Confirming our knowledge of the properties of colliding particles more clearly involves direct observation, and plausibly, inductive generalization. Understanding the way analogy can draw on inductive generalization is particularly important for understanding a common criticism of inductive generalization made by exclusive advocates of hypothetico-deductivism: that inductive generalization cannot *directly* generate hypotheses about unobservable entities, like light waves or atoms. This fact has been used by some hypothetico-deductivists to argue that, contrary to inductivists, the inductive generation of hypotheses is irrelevant to their confirmation. Laudan (1981) reports that some hypothetico-deductivists (here I have in mind Auguste Comte and possibly William Whewell) went so far as to suggest that inductive generalization has no role to play at all—that *only* hypothetico-deduction confirms hypotheses. But even if it is true that there can be no direct inductive generalizations about unobservables, it is possible that many hypotheses supported by analogies to entities we *can* generalize about.

Consider another example concerning unobservable entities: even if we cannot observe and generalize about atomic nuclei, we *can* observe and generalize about larger massive objects. When Rutherford bombarded gold foil with alpha particles, and found that the majority passed straight through while a few were deflected at drastic angles, he is reported to have compared the discovery to naval shells being deflected through tissue paper. Presumably only tiny but massive entities scattered throughout tissue paper could account for naval shell deflection, so Rutherford was led to conclude that there must be proportionally similar objects embedded in the gold foil. Famously he concluded that the bulk of atomic mass was concentrated in a nucleus; much of the rest of the atom was empty space. Now of course Rutherford had probably never made any special study of tissue paper and naval shells, but he doubtless had commonsense *generalized* knowledge about the properties of macroscropic objects such as: massive bodies ordinarily pass through thin sheets of flimsy material. And granted, he probably never bothered to explicitly consider or write down such an obvious piece of general knowledge, but it is difficult to imagine that it was not an important element of background knowledge leading to the generation of his original nuclear hypothesis (and as we shall see in the final section of the chapter, there is no good reason to think that every justification must be written down).

It is widely recognized that *arguments* by analogy can play a justificatory role, but a space needs to be made for the possibility of implicit analogical justification deriving from background inductive generalization. Of course the support analogy provides is not complete support: once the hypothesis is proposed, rigorous hypothetico-deductive testing along various parameters would be necessary. But allowing that background analogies can offer support would also allow us to see hypothesis formation as more than merely psychological, but also as logical. Premise 2 would then suffer: identical empirical consequences would not render two hypotheses equally well-confirmed.

All of this, of course, depends on the justificatory power of inductive generalization, and here Hume's problems loom. In the next major section of this chapter, I will turn directly to Hume's problems, and possible scientific solutions to them. But for the moment it is worth addressing some preliminary worries about how inductive generalization might serve as an independent support for hypotheses. Sarkar (2000) considers just this question, in relation to assessing Leplin's (1997) critique of the underdetermination thesis. Sarkar (2000, 190) notes a tension between the possibility that auxiliaries might be confirmed inductively, and Leplin's view that "brute inductive generalization" is itself not justifiable in the absence of some nomic or causal explanation of the generalization (1997, 24, 112–115). And, since nomic/causal explanations would need to be justified hypothetico-deductively themselves, inductive generalization could never be a source of independent support, at least on Leplin's own premises.

Sarkar is correct that there is a tension here, but I believe that it stems from Leplin's failure to articulate his own views precisely. In an interesting passage, Leplin notes that inductive generalizations may be supported by "explanatory" connections which themselves do not depend on "antecedent inductions." He writes:

> There is more to causal agency than the observed concomitances that Hume thought exhausted our perceptual access to agency. The idea of agency is supported by elementary impressions of agency, as when we experience effort or pressure. Such experience is the primitive basis of the idea of force, which is undefined, or only implicitly defined, in classical physics. A chalk mark made on the blackboard is not simply the concomitance of the chalk's motion with the appearance of the mark. One

> is directly aware of *making* the mark, of *imposing* that chalk upon the board…[T]he possibility of an experience of agency unmediated by generalization from instances supports, at least, an epistemic parity between inducing and explaining. (1997, 117–8)

Here Leplin is appealing to a kind of direct realist justification of generalization that is highly compatible with the account of the empirically given that I will propose at the end of this chapter, an account which relies in part on the idea that causal relations might be directly perceived, thereby allowing for the possibility of inductivist foundationalism. If there is any causal relation directly perceived, Leplin is surely correct that one's own agency may be the most fundamental, whereupon other forms of causality are attributed to nature by analogy. But in contrasting this view of "explanation" with inductive generalization, Leplin is assuming that induction begins with Humean regularities, stated in the form of singular propositions. One need not accept this if one thinks that even singular propositions, insofar as they involve the use of concepts, already involve generalizations. This is, in fact, the position of Norton (2003a) and Lange (2004), to which I shall turn at the end of the chapter. If inductive generalization itself is non-Humean, and is *constituted* by this direct realist "explanation," it is not the case that inductive generalization needs some separate explanatory justification.

In the last few paragraphs, I have spoken about forms of confirmation beyond hypothetico-deductive confirmation, in order to directly challenge the supposition that confirmation is ultimately hypothetico-deductive, and therefore to show that equality of empirical consequences does not imply equal confirmation. But there may still be some question as to how the use of alternative sources of confirmation works in practice, especially as regards using these sources to independently confirm or disconfirm the auxiliary hypotheses that are otherwise subject to the argument from confirmation holism. The idea that the independent confirmation of auxiliaries is a solution to the underdetermination problem is shared by Leplin (1997, 152–8) and also by Elliot Sober (1999). But it might be objected that expecting all of the relevant auxiliaries to be fixed independently is unrealistic, and therefore not a plausible solution to underdetermination.

We can only grapple with this question by breaking down types of auxiliaries into categories. The simplest type is discussed by Sober: assumptions about the reliability of particular scientific procedures. He considers the assumption that a pregnancy test is reliable: typically, he notes, the outcome of a pregnancy test does not alter one's confidence in the reliability of the test. Very different tests would be needed to judge this matter. Or consider my example: Galileo's telescope. Galileo's confidence that the scope allows him to see mountains on the moon is clearly based on prior calibration of the instrument with already-known terrestrial observations. These kinds of auxiliaries can clearly be supported independently by straight induction, together with *some* understanding of the causal mechanism involved. True, understanding the optics of the telescope is based on some prior theory, but it is theory that is closer to direct inductive generalizations than the theories he uses his telescope to support. Galileo did not need a wave theory of light, for example, to understand the relevant optics of the telescope, only the ray optics of the lenses and how they combine to form a certain image.

A second type of typical auxiliary is the *ceteris paribus* clause. Consider the relevance of such clauses in the testing of Newton's theory of gravitation: it was originally used to predict the orbits of both Uranus and Mercury to be different than they were actually observed to be. The original (false) Newtonian prediction about planetary orbits could be seen as following with the assistance of the eventually discarded auxiliary assumption, "There are no trans-Uranian planets." So the original prediction assumed that all things were equal, that there weren't interfering planets. But the prediction was false, so scientists called into question their assumption that all things were equal. A new auxiliary assumption was considered: that there *was* a trans-Uranian planet causing the perturbations. When such a planet was actually discovered (Neptune), the new auxiliary was independently confirmed, and Newtonian theory was saved from refutation. Before it was independently confirmed, it was merely a working hypothesis, not actual grounds for retaining Newtonian theory.

Of course the outcome was quite different in the case of Mercury. In that case, the original auxiliary (that there was no intra-Mercurial matter) was retained, while Newtonian theory was jettisoned in favor of general relativity. The difference between these two cases is often cited as evidence that the choice of hypothesis/auxiliary combination is pragmatic, even a result of social factors (such as the pressure to adopt a mathematically elegant new theory rather than retain the increasingly difficult to defend old theory). But there are conceivably good independent reasons that that the original auxiliary was retained in the case of Mercury, while it was discarded in the case of Uranus. It is important, for example, that whereas a trans-Uranian planet was actually found, intra-Mercurial never was, even though the latter, being confined to a smaller space closer to the Earth, should have been easier to detect, even by scientists of the day. That it should have been easier was itself a further auxiliary hypothesis, much like our assumption that we would see a large elephant in the room if it were here: since we do not see it, we conclude there is no such thing. Both of these, I think, derive from the general background knowledge that things are easier to find in smaller spaces. The same auxiliary would not have been applicable in the case of Neptune—more like a needle in a haystack than an elephant in a room—so there would have been reason to keep looking for a longer time without officially discarding Newtonian theory.

Other auxiliaries concerning unobservables are not so easily independently confirmed. An interesting case, also discussed by Leplin (1982), is Pauli's postulation of the neutrino to account for beta decay in a manner consistent with the conservation of energy. In this case, neutrinos would be in principle unobservable (unlike distant planets), so there is a serious question about how any auxiliary hypothesis used to predict facts about beta decay could be supported independently, rather than simply giving up the conservation of energy principle. Now I think that the neutrino hypothesis itself is formed by analogy to other observable cases in which conservation laws are preserved pending the discovery of some new interacting agent (for example, Lavoisier's early conservation of mass principles). Apart from that, the hypothesis was also eventually supported by *independent* hypothetico-deductive evidence.[60] Perhaps before this independent support was offered, energy conservation of beta decay was temporarily called into question in the intervening period, rather than accepted.[61]

So, yes, hypothetico-deduction is a crucial source of scientific confirmation. But it is not the only source or the *ultimate* source. As long as there are other sources, there are independent ways to confirm both auxiliaries and stand-alone hypotheses, and therefore ways to confirm empirically equivalent pairs differentially.

Premise 2: Do empirical consequences always confirm?

Once we accept scientific practice as a standard for evaluating conceptions of confirmation, I think it is also easier to show how observing an empirical consequence of a hypothesis does not always suffice to confirm it. If this is true, then simply from the fact that we might find a hypothesis that is empirically equivalent with another, we cannot infer that it is equally well-confirmed, because we cannot conclude that it is even confirmed at all.

To show this, we need only give examples of empirical consequences that no scientist would take as confirming any hypothesis. One traditional objection to hypothetico-deductivism on this front is the "problem of irrelevant conjunction," which notes that if observing an empirical consequence suffices to confirm a hypothesis, and if *e* is a logical consequence of H, then the conjunction of H with any arbitrary X is also confirmed by *e*, since anything that is a logical consequence of H is also a logical consequence of H&X. Of course no scientist would take any arbitrary conjunction as confirmed by the consequences of only one of the conjuncts.

There are more realistic examples of hypotheses that are not confirmed by every one of their logical consequences. Laudan and Leplin note that observing young men entering puberty after reading scripture does not confirm the strange hypothesis that scripture-reading induces puberty, and

[60] In 1956 the important Cowan-Reines experiment detected neutrons and positrons that should have been produced by the interaction of free neutrinos with protons.
[61] For more on how theories involving auxiliaries about unobservables might overcome underdetermination problems, see Michela Massimi's (2004) application of the method of demonstrative induction, rather than hypothetico-deduction, to aspects of quantum theory.

that the point of controlled experiments is precisely to rule out those empirical consequences which are not confirmatory (1991, 465–6). Even for theories not open to experimental testing, serious scientists do not consider theories with identical observational consequences equally well-confirmed. A good example from John Norton (2003b, 15) is the empirical equivalence between standard geology and young Earth creationism, as regards the fossil record: provided that the young Earth creationists tell the right story about how God put the fossils there to test our faith, the evidence we observe may be exactly what the latter theory predicts. Yet scientists do not take the second view seriously.

Bayesian epistemologists may argue that the objection above confuses two senses of "confirmation." They might say that while it is true that scientists do not take these arbitrary hypotheses as *absolutely* confirmed (i.e., having a probability above a certain threshold, warranting acceptability), this does not mean they do not take them as *incrementally* confirmed (i.e., having probabilities given the evidence which are *higher* than the prior probability of the hypotheses).[62] On this view, observing an empirical consequences does raise the probability of a hypothesis—even if only slightly—and so observing empirical consequences always suffices to confirm hypotheses incrementally. This is essentially the solution Maher (2004) presents for the problem of irrelevant conjunctions, which is analogous to the current problem: both objections suggest that hypothetico-deductivism is committed to the confirmation of strange hypotheses (either intrinsically strange ones, or ones paired with strange conjuncts). Maher suggests that an irrelevant conjunction *is* confirmed by *e*, though it is only incrementally confirmed. If scientists do not accept that these kinds of hypotheses are confirmed, then they are simply talking about a different kind of confirmation than incremental confirmation.

I am of two minds about this response. On the one hand, even if it is correct that hypotheses are always confirmed in this incremental sense by the observation of their empirical consequences, this does not yield the *kind* of sufficiency requirement needed by the underdetermination thesis. Just

[62] See Maher (2006) for the distinction.

because two hypothesis, one respectably scientific, the other a creation of fantasy, are both incrementally confirmed to some degree does not mean that they are incrementally confirmed to the *same* degree. Many measures of incremental confirmation are defined in terms of the prior unconditioned probabilities of the hypotheses, and if these are different, their probability on the evidence will be different as well. Other measures in terms of the probability of the evidence on the negation of the hypothesis have a similar outcome. So I could readily concede this kind of sufficiency clause, but show that it defeats Premise 2, and thereby underdetermination.

At the same time, I think there are reasons to question the claim that hypotheses are always confirmed incrementally by observation of their empirical consequences. In Bayesian epistemology, a hypothesis is incrementally confirmed on the assumption that the probability of the hypothesis given our background knowledge is as follows: $0 < P(H/D) < 1$. It has been traditionally assumed that as long as H (and D) involve no outright contradictions, $P(H/D)$ is indeed greater than 0. Perhaps it is this assumption that must be rejected. Perhaps to have some degree of probability given background knowledge, it is not enough that no contradictions obtain, that the hypothesis be conceivable. Perhaps there needs to be some special reason to think a hypothesis is possibly true—the kind of special reason often suggested by analogies or inductive evidence. Indeed it is not at all clear why *deductive* logical possibility should be a guide to empirical *inductive* probability.

In general, I think there is good reason to believe that defining confirmation in terms of probability is also putting the cart before the horse. I think "confirmation" is a more basic concept than "probability." We say that a proposition has a high-probability if it has been well-confirmed. Probability is a measure of the degree of confirmation; it isn't that confirmation is a measure of the degree of probability-raising. This is consistent with my suggestion above that some evidence—some confirmation—is needed before we can say that a hypothesis has even a tiny degree of probability. I doubt that "probability" is a primitive concept. Our idea of degrees of probability derives from the fact that we know that we can have different amounts of evidence for the same conclusion as we collect new facts. We come to this realization only after we have collected all of the evidence we think possible, and then realize that we didn't always have this much evidence. Probability is also clearly not primitive in that it is a modal modifier: a claim can be probably *true,* or possibly *true*, and we see these two different modalities of truth as falling along a continuum. The etymology of "probability" is also suggestive: probability derives from *provability,* and different degrees of probability do indeed reflect degrees to which a claim has been proved. Some facts about the history of science also support this. Philosophers and scientists debated theories of confirmation well before they had theories of probability. They also did not regard irrelevant conjunctions as confirmed. Now it is possible that when probability theory entered the field, and began to inform scientists' decision-making, there was a Kuhnian paradigm shift, and the old concept of confirmation was simply incommensurable with the new concept. But I doubt this. I have already given some reason to think that forms of confirmation that at first appear to be discontinuous (such as inductive generalization, analogy, and hypothetico-deduction) may in fact be related as fundamentals to derivatives. I suspect something like this is also true about the relation between confirmation theory and probability theory, though this has not yet been worked out (Norton will suggest some leads in our next section).

Of course I do not even need to defend against incremental confirmation to undermine the underdetermination thesis, and that is my primary task. Even if arbitrary hypotheses are incrementally confirmed, rival hypotheses are not necessarily equally incrementally confirmed, so there is no underdetermination problem that results from this formulation of the sufficiency point. And even if the sufficiency requirement cannot be refuted, the other attacks on Premise 2 (and on Premise 1) are still relevant. In any case, looking at the wider context of scientific practice shows us how to dissolve a skeptical problem that at first appeared to arise from scientific practice. We have used Quine's own putative strategy, of urging the skeptic not to "overreact," to disavow underdetermination, one of Quine's most basic skeptical theses.

Concluding notes on pragmatism and confirmation

In the above, I have surveyed some aspects of scientific practice that I take to be relevant in showing the impoverishment of the idea that all and only empirical consequences confirm. One question worth asking, however, is why I choose to admit the examples of scientific practice that I do. Some might object that the only criterion I could appeal to is a pragmatic one: scientists regard this and that as a source of confirmation because achieves successful prediction and explanation, etc. If that is in fact my criterion, then there is something disingenuous here to Quine. I noted at the very beginning that I would distinguish two senses of "confirm": a pragmatic and a non-pragmatic one. I acknowledged that Quine held that many theories are better *pragmatically* confirmed than empirically equivalent rivals, but wanted to decide if he was correct that theories were not determined by *logically* confirming evidence. If now I claim to show that other sources of confirmation are available—but only because of examples that have been chosen because of their pragmatic significance, then this might be taken as sleight of hand. I would be arguing that additional sources of logical confirmation are possible because pragmatic confirmation is possible, which would neglect the very distinction I attempted to begin with.

It is true that we could understand the choice of scientific examples as stemming from pragmatic considerations. But we do not need to. At minimum, I could say that I pick examples of scientific theories that are widely accepted by my audience for a variety of reasons. The question is *why* they are accepted. Some might accept them because of the pragmatic value they exhibit, others not. But in the course of presenting some of these theories, the reason these theories are accepted is not pragmatic at all. They may be based ultimately on analogy and induction, which are evidence-based forms of confirmation. Now it might be thought that this is a question-begging strategy. Whether or not evidence-based confirmation is possible is the very thing in question, so I should not simply propose that it is possible and show where more of it might be. But this charge is to misunderstand that logical confirmation is being question by a *reductio ad absurdum*, which permits us to accept certain premises as true in the interim, only to see that they yield unacceptable consequences and must be rejected. In the above, I have simply provisionally assumed that more forms of evidence are confirmatory than the hypothetico-deductivist is originally willing to assume. In the same way, the underdetermination skeptic would assume a smaller range of forms of confirmation in order to test whether such confirmation helps to determine theory choice. (Quine, for example, would assume that observation sentences can be accepted as true; the question is whether they yield the truth of theoretical sentences.) In this way, I help to dissolve the skeptic's *reduction* by showing that the absurdities result only from impoverished original assumptions.

Of course there could still be independent reasons for doubting even the evidential relevance of induction, for instance, and insisting that hypothetico-deduction is the only possible form of confirmation. Interestingly, however, the deeper the skepticism becomes, the harder it is to see what kind of scientific problems might have prompted these doubts. It therefore becomes increasingly plausible that purely philosophical assumptions are at work. Laudan and Leplin suggest, for instance, that the 20^{th} century commitment to hypothetico-deductivism derives from central problems in logical positivism. They note that the positivists' project was motivated by the wish to dismiss favorite targets (such as metaphysics, religion, or ethics) by proposing criteria according to which these targets counted as meaningless. Their goal was primarily epistemological: to characterize a certain kind of belief as irrational. But their means was semantic: to devise formal criteria "demarcating" the meaningful from the meaningless. The epistemic goal was then reached through the semantic injunction that the meaningless not be entertained. As a form of deductivism, hypothetico-deductivism helped flesh out this semantic injunction (1991, 466–8). First, positivists sought to define meaningfulness in terms of deducibility *from* observation statements ("strong verifiability"). As Soames (2005) outlines, it was only because this standard implied that too much ordinary scientific discourse would be meaningless that hypothetico-deductivism ("weak verifiability") was proposed: much more would count as meaningful on this latter view because it was far easier to find particular empirical consequences of scientific theories than it was to translate them wholly into sensory language.

With all of this in mind, it is worth wondering: *why* fight an epistemic war with logico-semantic proxies? Why not engage the enemies directly, by defining standards of epistemic justification in terms of something else? Why not, for example, define an *inductivist* standard of epistemic justification, rather than so many deductivist varieties? This is a question Quine answered explicitly, in his presentation and critique of Carnap in “Epistemology Naturalized.” Carnap’s project of offering translations of scientific statements in terms of observation language was what Quine called a “conceptual project,” specifying a logico-semantic criteria of cognitive value, offered for the sake of a wider “doctrinal” goal, i.e., legitimating scientific discourse by showing how its language could be eliminated in favor of respectable logical and set theoretical terms. Of course empiricist doctrinal goals *had* been more ambitious. Rather than simply *exculpating* scientific discourse—showing that it was at least not guilty of any major crimes concerning meaning—empiricist tradition had once sought to *justify* scientific knowledge. But Quine believes this effort had originally met its match when it faced Hume’s challenge: generalizations and singular statements about the future proved to be impossible to deduce from experience. Carnap’s project never advanced much beyond this “Humean predicament” because these same statements remained untranslatable and Carnap looked only for ways of stating their consequences. So, the reason that empiricists fought their battle with logico-semantical proxies was their prior conclusion that a full-frontal epistemological assault was futile. Quine, famously, extended the devastating implications of the Humean predicament by urging that any further attempt even at merely exculpatory translation, had to fail—because of his derivation of the indeterminacy thesis from confirmation holism and/or the underdetermination thesis.

But now we are in a unique position to reassess the original Humean predicament. Having dispensed with underdetermination-style skepticism by resorting to the free (non-question-begging) use of science, could we apply a similar strategy to more traditional anti-inductivist skepticism? To see whether we could, we first need to survey some possible theories of induction. Having examined these, we can revisit the standard Humean doubts to see whether they, too, arise as scientific doubts. If they do, we may be able to make further free use of science to block yet another *reductio*. But, if it turns out that Humean doubts themselves make some questionable purely philosophic assumptions, then Quine is wrong that all skeptical doubts arise in a scientific context. If he is wrong about that, it is devastating to his naturalism, because it will mean that Quinean naturalized epistemology is motivated in part by purely philosophic concerns—not scientific problems.

Actually, before moving on to examine any of these questions, I want to draw a line between what I have accomplished above and what I suggest below. While it is true that a full solution to the underdetermination problem requires showing that there are evidence-based sources of confirmation apart from hypothetico-deduction, it is fairly plausible that debates about whether these sources are available will be more purely philosophical than scientific. So I think that once I have shown that one of Quinean naturalism's main motivations—the underdetermination thesis—rests on more basic Humean questions about induction, then regardless of how those questions are to be answered, much of the critique of Quinean naturalism is already complete. Even if there are unchallengeable reasons to be skeptical about induction, the fact that these are likely to be philosophic rather than scientific reasons already undermines the naturalistic claim that skeptical doubts are prompted by science—and hence the claim that philosophy itself is continuous with science. For this reason, most of what follows should be considered as merely suggested alternatives to Humean doubts. This is not a dissertation on induction or perception, and I cannot hope to defend the views I suggest below in much detail.

A scientific solution to Humean doubts?

Since we are curious about the possibility of confronting Humean doubts head-on, we should look for accounts of induction that a) do not rely on any logico-semantic criteria, and b) that draw inspiration from the history of science. An ideal candidate for such an account was proposed by John D. Norton in his recent, fascinating paper, "A Material Theory of Induction" (2003).

Norton argues that attempts to understand inductive confirmation in formal terms face an inevitable tension between the universality of their scope and their success in describing inductive relevance. An example of this tension is seen in the view that induction proceeds by simple enumeration. Norton gives an example from Mill, noting the difference between the inference from "Some samples of element bismuth melt at 271°C" to "All samples of the element bismuth melt at 271°C" on the one hand, and the inference from "Some samples of wax melt at 91°C" to "All samples of wax melt at 91°C" on the other. One way to state a universal formula for induction by simple enumeration might be to say that we can always infer "All S is P" from "Some S is P." But a formality like this does not explain what makes the inference about bismuth superior to the inference about wax.

In order to state standards that underwrite successful inductions, Norton thinks that a criterion must not be formal or universal. This is why he thinks "All induction is local," i.e., all inductive inferences are licensed by knowledge of particular matters of fact, relevant within particular domains, rather than by a formal scheme applicable universally. His positive case results from surveying various theories of induction, and concluding that each theory usually does describe one aspect of inductive practice, but that the successful use of this practice depends on specific matters of fact related to the subject matter. For example, he notes that the usual way of explaining the superiority of the bismuth inference over the wax inference is comparable to Quine's account of the superiority of "green" vs. "grue" inferences (1969b): the first is regarded as a natural kind concept whereas the second is not. Whereas wax is a more variegated family of substances, bismuth is a chemical element, bringing with it the expectation of more uniform associated characteristics. But whether a particular concept designates a natural kind is itself a question depending on our local scientific (e.g. chemical) theories. To embellish Norton's point: "wax" may have good inductive potential within certain domains: from the fact that some wax melts it is reasonable to conclude that all or even much of it does, without specifying a melting point. But only special scientific theory gives us confidence in generalizing about something as specific as a melting *point*, e.g. the specific chemical knowledge contained in our concept of "bismuth." So the inductive relevance of simple enumeration—if there is any—is not a merely formal matter, but a matter of specific facts about the concepts involved in enumerative inductions.

Norton goes on to make parallel points about other theories of induction, especially the hypothetico-deductivist and probabilist views. We have already encountered the pitfalls of the formalist approach to hypothetico-deductivism. Understood purely formally, this view sanctions "promiscuous" inductions, in such a way that *any* hypothesis, no matter how arbitrarily selected, would count as incrementally confirmed just in case some actually observed empirical consequences are deducible from it. The problem of "irrelevant conjunctions" is one such example. If my development above is correct, and this view of confirmation is unacceptable by reference to actual scientific practice, then to the extent that empirical consequences *are* ever confirmatory, some constraint must be placed on which conjunctions would count as confirmed. Norton points out that a judgment about which conjuncts would count as "relevant" would, again, be a matter of local fact as judged by our background knowledge. Similar constraints could prevent the consideration of entirely arbitrary hypotheses, such as the one about scripture reading causing puberty.

Norton then surveys a variety of theories of confirmation intended to be constraints on the original hypothetico-deductive model. Exclusionary accounts, for example, supplement hypothetico-deductive confirmation with some demonstration that the *falsity* of the hypothesis would make it unlikely to observe the empirical consequences. Demonstrating this, Norton thinks, would depend on facts of the matter concerning, for example, what counts as the randomization of the sample in a particular controlled-experimental setup. Simplicity accounts try to show that only "simple" hypotheses are confirmed by empirical consequences. But Norton thinks that judgments of simplicity (as in the stock example of curve fitting) are essentially judgments about whether or not to consider some causal factor as operative in the particular domain of facts (as when deciding to raise the power of the polynomial equation fitting the data points). The inference to the best explanation version of hypothetico-deductivism is similar: empirical consequences confirm a hypothesis only if the truth of the hypothesis would also help explain them. But what it is to explain in a particular domain (e.g. physical vs. biological vs. social), and what factors count as relevant to explanation (e.g. which types of causes, if explanation is about displaying causes), are judgments concerning local matters of fact. Finally, much the same is true about judgments of reliability involved in reliabilist views.

Norton thinks that a final class of confirmation theories—probabilist ones—are the most successful in making a case for a universal inductive scheme. Most of his discussion concerns subjective Bayesianism, which he critiques for reasons unimportant here. His most important comments concern Bayesianism as such. However probability is to be understood (as degree of belief or some objective relation to evidence), Norton urges that the calculation of probability conditional on evidence is determined by "factual properties of the relevant system." He thinks this holds even when we try to determine if the hypothesis *H,* "All swans are white," is confirmed by the evidence *e* that some particular swan is white. In this case, *H* entails *e*, and $P(e|H)=1$, but we must assess $P(e|{\sim}H)$, to know the probability of finding white swans if not all swans are white, and this depends on the prior probability $P(e\&{\sim}H)/P({\sim}H)$. Presumably he thinks this and other priors could only be assessed by reference to "factual properties of the relevant system." [63] Presumably he would make the same point about the dependence of $P(H|e)$ on $P(e|H)$, where $P(e|H)$ is to be determined by our knowledge of how the world would be if *H* were true of it.

Having surveyed various views of induction, to show that each has its place depending on the domain of facts in question, Norton presents the most remarkable part of his paper: a discussion of how a material theory of induction avoids Hume's problem of induction. Famously, Hume's problem

[63] It is, of course, not uncontroversial to think that prior probabilities must be factual in some way. Maher (2006) denies this, holding that they must be understood in a purely logical, *a priori* manner. I address these concerns in an online appendix at http://www.benbayer.com/blog/archives/2007/03/do_probability.html.

claims that there can be no justification of induction. Such a justification would be either deductive or inductive. If the justification is deductive, it fails to provide ampliative inferences beyond observation reports, as universal generalizations cannot be deduced from singular statements. If the justification is inductive, it either generates circularity or an infinite regress: circularity, because relying on the past success of induction to infer its future success would clearly be using induction; an infinite regress, because citing some principle like "the uniformity of nature" to bolster our inductions would then require some further principle as justification (how do we know nature is uniform, and in the relevant respects?). Norton observes that many of these problems might also befall a material theory of induction. Just as a deductive justification of induction would fail to ground its special inductive character, so some *a priori* postulate of a material fact of the matter would likewise sunder material induction. Just as the justification for a formal inductive schema would be circular if it applied the very same schema, a justification of a material induction would be circular if it used the same material fact referenced in its conclusion. The chief difference between formal and material views of induction concerns the last possibility, regress.

A formal schema of inductive inference could justify itself by reference to some different inductive schema (like the uniformity of nature), and further meta-justifications of this schema would likely be more artificial and abstract. A material induction faces no such problem. If the conclusion of one material induction is licensed by some different *particular* fact, the conclusion of an earlier material induction, there is a regress involved, but there is no reason to think it must be infinite. Rather than becoming more "fanciful" and abstract, the regress of material induction could gradually become more familiar and concrete. Norton (2003a, 668) concludes:

> The regress described here…merely describes the routine inductive explorations in science. Facts are inductively grounded on other facts; and those in yet other facts; and so on. As we trace back the justifications of justifications of inductions, we are simply engaged in the repeated exercise of displaying the reasons for why we believe this or that fact within our sciences.

Consider Norton's example, "All samples of the element bismuth melt at 271°C." This generalization was not supported by a simple formal rule, such as the general relevance of inferring "All S is P" from "Some S is P" (or "All observed S is P," etc.). What gives us confidence in this generalization, unlike the similar generalization about the melting point of wax is the specific knowledge that bismuth is a separate chemical element, plus our general knowledge about the uniform nature of chemical elements. Bismuth's status as an element is justified by isolating it from its oxides (bismite, bismuthinite or bismutite), and from the subsequent inability to isolate it further by chemical means. The bulk of our confidence in the induction comes, of course, from subsuming bismuth under our concept of "element," which is informed by quantum theory, atomic theory, and early modern chemistry, each of which were founded on a series of observations and experiments. We are confident in our bismuth induction because we are confident that quantum theory explains the properties of elements by reference to atomic structure. We are confident in the existence of atomic structure because of the work of by Thomson, Rutherford, Dalton, etc. We are confident that this work relates to standard chemical methods of isolating elements because of Lavoisier's early work with the most basic of elements (hydrogen, oxygen, etc.), establishing chemistry as a discipline. Each of these theories is in turn supported by a series of experiments on particular substances. As our justifications take us back into the historical foundations of chemistry, the facts to which we appeal become more and more concrete, like Dalton's observations about the combination of elements in multiple, integral proportions, Lavoisier's observations about the conservation of mass of the reactants in combustion, etc. In the material theory of induction, the context of discovery and the context of justification begin to intertwine.

It is noteworthy that with the material theory, a source of ambiguity that exacerbated Hume's problem is exposed: that between the attempt to justify *induction* and the attempt to justify *particular inductions*. Treating inductive justification as the first leads to a formalist conception of induction, and with this are associated the usual Humean problems. But treating inductive justification as the need to justify particular *inductions*, i.e. particular inferences, not only offers the hope of avoiding the infinite regress, but also the promise of doing so by attending seriously to actual scientific practice—just what our anti-skeptical strategy recommended.

Norton does face one serious problem. While the conception of material induction does at least make it more plausible that there need be no infinite regress of justification, it is not clear how the regress would terminate. Norton himself is not entirely sure. He suggests that it is possible that the regress could terminate in beliefs only admitting of circular justification, but they could also end in "brute facts of experience" in no further need of justification. Norton takes the uncertainty here as good news: as long as the foundationalist option is open, there is no reason to think that Hume has offered a knock-down argument against the possibility of justifying inductions. I think Norton is right about Hume, but it is worth considering whether latter-day empiricists may have gone any further. The idea that termini in "brute facts" of experience could provide epistemic justification ("empirical foundationalism") is precisely what is denied by coherentist empiricists like Wilfred Sellars (1963), and Donald Davidson (1983), both of whom agree that treating the "empirically given" as a source of justification is a "myth." The Sellars-Davidson objections to the given usually involve the following two points: 1) experience at best causes our beliefs, but mere causation is not sufficient for justification; 2) even if experience of "brute facts" gives us awareness of "bare particulars," it does not give us the perspective to judge the *authority* of our evidence, the relevance of the observation of a particular to propositional judgments involving the application of concepts. The upshot is that "nothing can count as a reason for holding a belief except another belief" (Davidson 1983, 426). As a result, Humean skepticism about the justification of our inductions is still a looming threat.[64] Clearly, then, if we are to dispel this threat, we must find an answer to the two Sellars-Davidson points.

[64] Coherentists of this stripe are of course quite content with this result, and do not think skepticism is necessary, provided that any resulting circularity is "virtuous," or any infinity is "progressive." This is no place to assess the case for non-skeptical coherentism, of course. We can only observe that the case is a difficult one to make, and that some of its most ambitious advocates in recent years, such as BonJour (1985), have given up the fight (BonJour 2000).

Notice, however, that in the course of searching for solutions to Humean doubts, while we have looked to scientific practice for an account of inductive justification, we have not yet located the source of Humean *doubts* in scientific practice. In fact it is quite plausible that doubts about the epistemic relevance of the empirically given—even if not in the same form as Sellars and Davidson—were also relevant in Humean skepticism about induction. Hume's sensationalism, for instance, entailed the impossibility of perceiving causal relationships. We will soon see that aspects of his view of the senses had scientific sources, but as our exploration deepens we will see that more and more philosophical presuppositions begin to creep in. In our next and final section, I show that our anti-skeptical strategy of making free use of science to answer skeptical reductios can pay some dividends even on the question of skepticism about the senses. There is a substantial body of scientific evidence, particularly in the field of psychology, which can help dispel the doubts summarized in the two Sellars-Davidson points, and others relevant to Hume's views about the perception of causal relationships. But these dividends are limited. Science will only take us so far; we will have to make some normal philosophic arguments to dispel doubts past a certain point. But by the same token, these same doubts are not the product of science. This will suggest that Quinean naturalism's roots in the Humean predicament are not themselves fully naturalistic.

Clearing the naturalistic ground for inductive foundations

In this section, I address each of the two points of the Sellars-Davidson objection to empirical foundationalism, by indicating the kind of scientific evidence they seem to derive from and the wider body of evidence that should cancel their significance. In particular, I show how the wider evidence answering the first objection also sets the stage for addressing the second. In both points, however, we will find philosophical assumptions buried beneath scientific concerns.

I begin with the first anti-foundationalist objection: that experience cannot help justify our beliefs because at best it can only help *explain* them, causally. As with previous skeptical problems, there is an important kernel of truth here. Anti-foundationalists are correct to think that *mere* explanation would probably not suffice for justification. So the question is whether experience *merely* causally explains our beliefs, or perhaps does this plus something more.

A further kernel of truth in the first anti-foundationalist objection is that the conventional empiricist picture of experience does lend credence to the idea that experience merely causes our beliefs. This can be seen in both classical and modern versions of empiricism. In the classical empiricism of Locke and Hume, experience is conceived as an awareness of "ideas" or "perceptions" in an inner mental theater. At best experience offers an indirect, representational awareness of the world if there is some inference one can make from the inner mental theater to the external world. Famously, of course, Berkeley and Hume show that such an inference is difficult if not impossible or incoherent. Modern empiricists (the logical positivists) dispensed with the idea of internal mental entities, and instead sought to define or translate knowledge claims in terms of sensory *language.* This met with failures we have already discussed. In this regard Quine was to the positivists as Hume was to the earlier classical empiricists, drawing as he did skeptical conclusions from shared basic premises of his predecessors.

Both Hume and Quine were informed in this effort by the traditional empiricist view about experience. Quine thought that empiricism advanced with the focus on language rather than "ideas" (1975b), but interestingly, the classical empiricist view was not so far removed from Quine's (1969a) idea that the "stimulus meaning" of an observation statement is determined by "certain patterns of irradiation in assorted frequencies" (1969a, 83), or by "the impact of light rays and molecules on our sensory surfaces" (2004d, 288). The classical empiricists did not arrive at their theory of the inner mental theater through sheer fancy, but in large part from scientific considerations also reflected in Quine. Yolton (1979) presents the widely accepted view that classical empiricists' theory of experience was in large part influenced by their reading of early modern theories of optics and vision. In the *Treatise*, for example, Hume states that "my senses convey to me only the impressions of colour'd points, dispos'd in a certain manner" (1978, 34). This and other passages (in Hume and Berkeley) bear close resemblance to passages from Newton's *Opticks*.[65] Elsewhere Hume makes reference to optical analogies popular in his day, such as those involving "objects seen in mirrors, the eye compared to a lens, the understanding compared to a *camera obscura*," etc. (Yolton 1979, 210). Hume even agrees with the claim of Berkeley's theory of vision that our sight could not give us information about such things as the distance or extension of objects (Hume 1978, 56, 191).

Of course some of this should not be surprising to Quine, who makes reference at least to the issue of the scientific sources of skepticism about three-dimensional awareness (1974, 1–2). But he does not seem to appreciate just how deep the roots of this skepticism run. In response to the contention that we could not be aware of three dimensions if we begin with a two-dimensional retinal image, Quine responds that it is of course by examining three-dimensional eyes that we first discover the two-dimensional retinal image. This response shores up his naturalism, but it concedes that the retinal image—or more generally, mere sensory *stimulation*—is our only source of information about the world. It is this "proximal stimulus" view that Davidson (1990a) identifies as one of the chief sources of Quine's own skepticism. Davidson suggests that a "distal stimulus" view, one which locates the source of meaning and evidence in the "very events and objects that [observation sentences] are naturally and correctly interpreted as being about" (72) would not have the same skeptical tendencies.[66] In spite of this, the "proximal stimulus" view is intensely popular among empiricists, and may even sound uncontroversial at first blush to the modern reader.

Nevertheless, there are serious scientific (and philosophical) grounds for doubting the "proximal stimulus" view, which holds the retinal image as the paradigm of perceptual information. A persuasive defense of a "distal" theory of perception was famously advanced by psychologist J.J.

[65] Newton speaks of rays of light conveyed "in so many Points in the bottom of the Eye, and there to paint the Picture of the Object upon that skin" (Newton 1952, 15)

[66] Davidson, of course, does not see any salvation for foundationalism, even in the "distal stimulus" view. He only sees a source of meaning and evidence, not justification. Justification, for him, must come through coherence. In what remains, I will show how a decent "distal stimulus" theory will actually provide grounds for foundationalism not only about meaning and evidence, but justification as well.

Gibson in the mid-to-late twentieth century (1966; 1986). Gibson insisted that a theory of vision is not exhausted by a theory of optics, and that theories treating isolated stimuli as the only source of perceptual information are impoverished accounts of the biological nature of perception. In his theory of vision, Gibson replaces the retinal image with his concept of the "ambient optic array," the totality of structured light surrounding the organism, which, he argues, uniquely specifies (i.e., isomorphically covaries with) the layout of the organism's environment. Perception of objects in an environment works through "information pickup" from this ambient array, via the organism's *active* interaction—involving all of its sensory modalities, in the context of its total bodily motion—with that light. Through this active interplay, the organism discovers the "invariant" properties of the world, for example the invariant relationship among the different angles of a solid object. The organism is not aware of the light, but through a physiological process, *uses* information in the light to achieve (not infer) awareness of its environment. Gibson originally formulated his theory in the course of his research on the visual perception of WWII fighter pilots, whose distance-perception suffered from anomalies which could not be understood using the classical empiricist ideas of "cues" for depth perception. But he gradually assembled an impressive array of further experimental evidence (1986, 147–202). His theories spawned an entire discipline, ecological psychology, which continues active research to this day. This research is now adding steam to work in the philosophy of perception, where more and more philosophers are beginning to take up the cause of "direct realism" (Pollock 1986; Pollack and Oved 2005; Kelley 1986; Putnam 1994; Huemer 2001; Noë 2002). According to direct realism, perception of the world is "direct" in that involves neither awareness of mental intermediaries nor inferences drawn from them. Of course Gibson and direct realism are not without their critics, but the strength of the research program and increasing popularity of its arguments suggest that the epistemic significance of direct perception is not to be brushed aside.

Direct realism is, of course, a controversial philosophical position. But it is a position which did not exist in its current state of development at the time Quine pondered the Humean predicament and its seemingly unavoidable anti-foundationalist consequences. Many traditional objections to direct realism have since been answered (see Le Morvan (2004)), and I do not wish to restate these here. What is important, I think, is that this *is* a philosophical debate, even if there are also scientific issues it brushes up against. In the next few pages, I will mention some of the scientific and philosophical issues that arise in the debate, with emphasis on the latter. I will not be able to answer every objection to direct realism, and that is not the point. The point is that there is a philosophical (though scientifically informed) theory available that helps to answer the first of the Sellars-Davidson points. Perception surely is a partial cause of our beliefs, but the Gibsonian point is that it is not *merely* a cause. Justification may be a species of causation, provided that the right kind of causation is involved. If Gibson and others are right, perception is a form of cognitive awareness of the world. Insofar as the world is what our beliefs are about, at least part of the ground is cleared for seeing how perception might also help justify beliefs about the world, and how the second of the Sellars-Davidson points might be answered.

Probably the most influential critique of Gibson, on both scientific and philosophical grounds, was advanced by Fodor and Pylyshyn (1981). They claim that even though there are important elements of truth in Gibson's critique of predecessor theories of vision (such as Helmholtz's), his "direct perception" cannot be as direct (non-inferential) as he would like. They argue that the subject must make use of *inference* from the ambient light array in order to determine properties of the environment. It will be useful to briefly review Fodor and Pylyshyn's objections, not only to see how Gibsonians might answer them, but also to see how they depend on philosophical, rather than scientific presuppositions.

Fodor and Pylyshyn argue that Gibson's view of information pickup trivializes perception, unless can find a way to constrain information pickup to the specific properties of the environment he mentions, such as "texture, shape, illumination, reflectance, and resistance to deformation" (144). Gibson characterizes these as "ecological properties," but Fodor and Pylynshyn argue that Gibson has no non-circular way of defining which properties count as "ecological." In particular, they allege that he cannot define the ecological properties, because there are no "ecological laws" in terms of which these properties might be understood as being "projectible." They block another route to understanding property detection in terms of the individuation of perceptual systems which respond to them, arguing that Gibson's attempt to individuate these systems in terms more holistic than individual modalities is lacking an independent criterion.

In fact I think that Gibson could indeed try to define ecological properties in terms of ecological laws, but he would in turn define ecological laws in terms of the ultimate functions of perceptual systems understood in the most holistic sense. He would say there is only *one* perceptual system: the perceptual system of the living, acting, organism, and its function is to aid the organism in its survival. Ecological properties, then, are those aspects of the environment which bear most directly on the survival of the organism. This fairly straightforward biological significance is something that Fodor and Pylyshyn's computationally-oriented theory, like so many others, seems to ignore. Turvey et al. (1981) elaborate upon this error, showing how the "affordances" of objects—their dispositions to be grabbed, climbed, dug into, copulate with, crawled upon, lept over, alighted upon, etc.—are the properties which information pickup responds to, *because* in the evolutionary history of the animal, these are the properties that have been most crucial in the organisms success at grabbing, climbing, digging, copulating, crawling, leaping, alighting, etc (1980, 260–7).

Fodor and Pylyshyn's remaining objections are more philosophical than they are scientific. To begin with, they argue that Gibson cannot account for what *they* take to be paradigm cases of perception, such as perceiving that something is a shoe or that something is a Da Vinci. But these cases are paradigmatic only on the "Establishment" view that treats perception as a form of judgment or "perceptual belief," according to which any justifiable "perceives that" locution counts as a *perception* in the strictest sense. According to Gibson, however, what makes perception "effective" is not necessarily that it results in a belief, but more fundamentally that it involves and guides activity. Gibson is not committed to the idea that perceptual awareness is in propositional form.

The view that perception comes in propositional form is, of course, endemic in philosophy, and Fodor and Pylyshyn support it by insisting that if perception were not propositional, we would have the "problem" of being unable to account for "misperception," i.e. sensory illusions and deception. To the naïve reader, avoiding the possibility of sensory deception would seem to be a strength, not a weakness, of a theory of perception that seeks to account for the possibility of foundational justification. One way of responding to the argument from illusion is that while we do often commit errors in *judgment* regarding the stick in water, this does not mean that perception itself misrepresents the object. Propositional content is not needed to account for what the illusion of a bent stick has in common with an actual bent stick. All that is needed is to say that each of these involves a similar perceptual look (Turvey et al. 1981, 275).[67] An illusory percept has this look in common with "veridical" percepts of actually bent sticks, but is illusory only in that it is an unusual form of perceiving a stick that is actually straight.

There is, of course, a philosophical dispute about whether illusions can even be described as illusions without implying that they involve misrepresentation. The usual locution is that an illusion involves seeing a straight stick *as* a bent stick, and since the straight stick is not bent, this implies a representation of something as something that it is not. Of course it is important to distinguish two senses of "seeing a straight stick as a bent stick." One sense involves applying the concept "straight" to the stick, which is the propositional form of awareness we are not at the moment considering. The

[67] It is true that certain forms of perceptual awareness will deliver more or less information, which may, in a given situation, have more or less biological utility. But this does not mean that the senses are "mistaken." First of all, biological disvalue is not the same as epistemic disvalue. Simply because some forms of awareness are out of the ordinary does not mean they fail to be forms of awareness of the world. It is possible to grasp the same objects in different forms. Second, what appears to be a biological disadvantage in a narrow circumstance can be understood as a consequence of the possession of a biological capacity that is itself of great utility. The ability to detect refraction differentials, for example, while misleading to a subject's judgment in certain cases, is nonetheless essential to the organism's ability to detect the difference between solid, liquid, and air—a crucial ability from the perspective of survival.

remaining sense, then, is simply a comparison of the two perceptual appearances. How we are to understand perceptual appearance, of course, is a difficult philosophical question. If appearances are *qualia*, which many philosophers have difficulty reducing to the physical, it may be difficult for external observers even to compare appearances to objects in order to say that the two are different and that *qualia* misrepresent. And as Berkeley argued long ago, it is impossible for the observer himself to get outside of his consciousness and compare his ideas with things outside his ideas. (Of course we can agree with Berkeley about this, without following him to his idealist conclusions, if we remember that *qualia* are our forms of awareness, not the *objects* of awareness.) Now of course there may be physical correlates on which *qualia* supervene, and perhaps these may be compared to the objects themselves. The big question, however, is whether there is a naturalistic concept of representation that could sustain a fruitful comparison. I have discussed problems with this project in some detail in chapter 3. The most viable theory is that representations are somehow isomorphic with what they represent in the world. One problem with applying this view to the question of illusion is that arguably an illusory representation of a straight stick as bent is still isomorphic with the straight stick on some way, because isomorphism is cheap. The causal connection between the two virtually guarantees this. Surely the illusory representation does not have the same isomorphism with a bent stick as a non-illusory representation, but then the question becomes, why privilege one isomorphism over the other?[68]

A final objection to Gibson from Fodor and Pylyshyn is that even if Gibson can specify the detected ecological properties that bring about perception, it is only through a process of *inference* that one can proceed from these detected properties to perceiving the environment. If perception requires inference, it is clearly not direct, and for that matter, not clearly distinguishable from judgment, and

[68] Turvey et al. (1981) make a useful response to the objection that direct realism is somehow incompatible with the possibility of illusion. They point out that to say that the senses are "in error" suggests that they should represent the bent stick in water (for example) differently than they actually do. But *this* ignores the fact that the senses' response to the structure of light in a way that displays the "differential in refractive indices between the media of air and water" actually delivers important information to the organism, and that there is no standard by which to judge this report as "in error."

therefore subject to error. This inferentialist view of perception is not unique to Fodor and Pylyshyn's particular brand of computationalism, but extends back as far as the perceptual theories of Helmholtz.

It is undeniably true that a great deal of physiological processing intervenes between the reception of stimuli (globally understood) and the production of percepts. The question is whether there is good reason to characterize this physiological process as a process of *inference*. This question is parallel to the question of whether or not to consider physical correlates of *qualia* to be representations. Arguably a naturalistic account of perceptual inference depends on a naturalistic account of representations, for surely naturalists would admit that inference is not inference unless it involves the manipulation of representations. So all of the problems raised above and in chapter 3 count against this position, as well.

There are a number of other problems with characterizing mere physical processing as "inference." Our usual conception of inference would seem to commit the inferentialist to the idea that pre-conceptual inference would require innate knowledge, and also the possibility of reversing illusions. I will not elaborate on these at present.[69] There is one additional objection from the commonsense concept of inference that is worth focusing on, for the sake of bringing out some of the purely philosophic points at issue here. An inference is usually understood as a kind of argument drawn from known premises. The "premises" of inferential perception would, presumably, be either

[69] Kelley (1986) points out two further problems with the inferentialist view of perception. First, inference is also usually understood as involving the application of background knowledge to known premises. However the kind of background knowledge presupposed by the sorts of inferences allegedly needed for the transition from detected properties to perception would be of an extremely advanced kind. In order to calculate three-dimensional perception based on ocular disparity, for example, the pre-conceptual child would need knowledge of the geometry of parallax. Since newborn children can perceive three-dimensional objects—but have not yet studied geometry—this knowledge would need to be innate. Even if this kind of innate knowledge is not impossible, the burden of explaining its possibility is nonetheless a heavy one. Direct realists, by contrast, can accept any physiological findings relevant to the mechanism of perception, while simply insisting that whatever the mechanism is, it is not inferential. See Kelley (1986, 61). Second, inference is usually understood as a conscious mental process. The kind of inference described so far—involving premises and background knowledge of which the perceiver is not aware—would have to be a special kind of unconscious inference. The discovery that the conclusion of a conscious inference is false often leads us to reassess its premises and reject the same inference in the future. Yet if perception is inferential, this means that perceptual illusions are as well. Only we cannot "reverse" perceptual illusions as we might reverse other "erroneous" inferences. We do not stop seeing the stick as bent when we learn that it really isn't. See Kelley (1986, 67).

the properties of the retinal stimuli (in the proximal viewpoint) or the detected ecological properties (in the distal viewpoint). Surely we are aware of objects *by means* of retinal stimuli or surface properties, but the question is whether this implies *awareness* of these properties *as* properties. The proper direct realist response would be to suggest that the means of awareness and object of awareness should not be confused, so awareness of objects by means of retinal stimuli and surface properties does not imply awareness *of* these stimuli or properties, and therefore does not imply the use of that knowledge as premises in an inference. Now there is a long tradition, most pronounced in the Kantian tradition, which holds that only *absolutely* unmediated awareness counts as awareness. This is what prompts philosophers to look for the objects of awareness in the most "internal" states possible, rather than in external objects. But there is simply no good philosophical reason to think that awareness counts as awareness only if it is unmediated. Hence there is every reason to take phenomenological facts at face value: we are aware of objects, not images or stimuli or even surface properties.[70]

Even if Gibsonian direct realism is plausible, another part of that ground remains to be cleared. As BonJour has noted in his recent critique of direct realism (2004), even if we accept that perception is a form of awareness unmediated by inference from internal mental objects, this does not yet show how it is related to the justification of beliefs. Particularly if we are to understand direct perception as a non-conceptual, *non-propositional* form of awareness, there is then a serious problem about how such a form of awareness could come to justify conceptual, propositional beliefs.[71] BonJour considers Reynolds' (1991) version of this view, according to which non-propositional experience justifies by being the object of recognitional skills that can be captured in rules but that are not articulated as such by the recognizers. BonJour is concerned, however, that Reynolds' view of

[70] See Kelley (1986, 66-8).

[71] Considering this version of direct realism is clearly advantageous, as it avoids the objections BonJour raises to the other versions which consider perception to be either propositional but non-conceptual Huemer (2001) or both propositional and conceptual Brewer (1999). The objection to both such views is, essentially, that if perception is propositional, it could be true or false, and is then in need of justification itself. This means it cannot serve the role of a foundation that is usually demanded of it. Of course I think the answer here is that perception is non-propositional, and the automatic deliverances of the senses are neither true nor false. See paragraph 4 of footnote 16 for more.

justification is not sufficiently internalist: according to it, we follow these recognition rules only because our "epistemic community" has adopted them for their efficacy in producing true beliefs—it is not clear if we adopt them because of our own understanding of their connection to truth (Bonjour 2004, 365). As before, this skepticism about the efficacy perception stems from a core of truth: there does indeed seem to be a difference between merely learning to replicate a social skill and forming a justified belief. If we are to defend the epistemic significance of direct realism, more is needed.

One possibility is that we can understand perception's epistemic relevance to propositional justification via *concepts*. Justified judgment essentially means the justified application of concepts to a new situation. Of course there is a longstanding philosophic tradition, deriving from Kantian-Hegelian premises, that takes a dim view of subject-predicate judgment as a source of genuine cognition. This view finds its modern equivalent in the doctrines of Sellars and Brandom, who allege that an organism's ability to categorize its environment makes it sentient, but not sapient. To cognize is to do more than to have a reliable differential response to the environment. Cognition is instead a product of the inferential commitments of one's judgments, commitments entered into through social interaction. It is of course now obvious that we are dealing with a *paradigmatic* non-naturalistic philosophic objection to the possibility of foundationalism. If naturalized epistemologists have to rely on "myth of the given" objections to counter foundationalism, they are in a great deal of trouble, because these skeptical doubts are not prompted by any obviously scientific problems—unless it makes sense to say that Kant's and Hegel's theories were so prompted.

Nevertheless I want to say a few things about how to answer BonJour's concern in a way that addresses the problems raised by Sellars and Brandom. Of course one issue is the worry about the difference between sentience and sapience. I hope that I have already answered that objection to a degree by showing that perception is not a mere response to the environment, but an active, cognitive probing of it. This doesn't do everything to explain the possibility of *human* sapience, of course, because even animals have perceptual capacities. A question remains as to how perception gives any authority to subject-predicate judgments.

John McDowell, a partisan for the Sellars-Brandom critique of the authority of the empirically given, considers one way in which perception might be taken as authoritative. We may be able to understand the conditions of justified concept-application in relation to the conditions in which the same concepts were originally formed: we are justified in applying a concept whenever we observe aspects of our environment to be similar to the aspects of the environment from which we *formed* the concept. McDowell considers just this possibility in order to dismiss it:

> We could not begin to suppose that we understand how pointing to a bit of the Given could justify the use of a concept in judgment—unless we took this possibility of warrant to be constitutive of the concept's being what it is, and hence constitutive of its contribution to any thinkable content it figures in. . . . The supposed requirement is reflected in a familiar picture of the formation of such concepts, a picture that is a natural counterpart to the idea of the Given . . . [that] the associated conceptual capacities must be acquired from confrontations with suitable bits of the Given. . . . A subject would have to abstract out the right element in the presented multiplicity (1996, 6–7).

So McDowell acknowledges that an abstractionist view of concepts could help explain how "the Given" justifies judgments. But invoking the work of Geach (1957), he simply denies that observational concepts *could* be abstracted. Now Geach offers compelling reasons for why *color* concepts could not be the first ones we abstract from perceptual experience, contrary to traditional empiricists. Indeed he is right that young children form concepts of middle-sized objects well before they form color concepts (or any other attribute concepts, for that matter) (Anglin 1977, 11). But at no point do Geach or McDowell offer reason to doubt that *these* concepts of middle-sized objects might be formed first by abstracting from perceived similarities. It is worth considering if that is possible.

John Pollock and Iris Oved (2005) understand that we recognize middle-sized objects, like cats, even when we cannot articulate a description of cats in terms of shape- or color-concepts. They speculate about how this recognition might result from a "cat detector" cognitive "module," that is acquired after seeing how the various parts of cats fit and move together in a certain uniform way, in just the same way that chicken-sexers learn their skill on the basis of an implicit learning of chicken-parts. Perhaps this results from the formation of a connectionist "neural network" that permits a sophisticated form of pattern recognition (333–8). What Pollock and Oved are gesturing toward is the recognition that there is a *perceptual similarity* relation available for three-dimensional shapes. Even though young children do not yet understand shape concepts—especially not those of the three-dimensional kind—it does seem that the most salient similarities among the first kinds of objects they do conceptualize (types of animal, people, food, artifacts, etc.) are in regard to shape (Anglin 1977, 71). It is important that this similarity is perceptual—meaning non-propositional and non-inferential—for otherwise we fall prey to Sellars' objection that a *judgment* of similarity would presuppose further concepts (such as that of "similarity," and that of the respect in which things are said to be similar) (Sellars 1963, 160).[72]

There is, however, still something of a puzzle about how such three-dimensional similarities could be grasped without the possession of any prior concepts. This is where Gibson's view of perception begins to pay dividends even for our understanding of concept-formation, where our answer to the first Sellars-Davidson point informs our answer to the second. Recall that central to Gibson's theory is the distinction between information that is "picked up" from the ambient optic array on the one hand, and the things or properties of the world that are perceived on the other. Information about properties such as shape and resistance to deformation are picked up, and "invariant" relationships among these are extracted in a way that produces an awareness of the object as a whole. So something like the grasp of three-dimensional shape is already involved in perception of a given object in the first place: it comes as little surprise, then, that this ability can be marshaled for the purpose of grasping similarity relationships.

[72] These similarities would need to grasped relationally, by seeing how two objects differ less from each other than they do from a third. They could not result from simply reading off common features from any two objects, in the Lockean fashion, for this would presuppose a conceptualization of the features in question. For more, see Kelley (1984, 336-342).

Jesse Prinz (2002) has formulated a new theory of "concept empiricism" that explains how concepts of middle-sized objects are abstracted from perceptual experience, which exploits mental devices he calls "proxytypes," products of "long-term memory networks" which permit the grouping together of quantitatively-varied percepts. Simple examples of these networks include "hierarchical" representations, which encode and integrate changes in experience as one "zooms" in or out on an object, "transformational" representations, which do the same for moving objects, and most importantly, "*predicative"* representations, which exploit the other types to group objects on the basis of similarity (as if a similarity range could be grasped as a transformation from one similar to another) (141–4).[73] Here, in particular, we see how proxytype theory predicts just what we would expect about similarity in light of Gibson's findings about perception:

> Proxytype theory offers an attractive account of basic-level categorization. The basic level generally lies at an intermediate [middle-sized] level of abstraction. . . . This suggests that shape similarity confers an advantage in learning and categorization. Proxytype theory predicts this because proxytypes are perceptually derived and shape plays a very dominant role in object perception. (163) [74]

Now Prinz's theory is somewhat controversial, but it does have the advantage of cohering with a vast amount of evidence from developmental and cognitive psychologists about concept-formation and use, which previous theories seemed to be able to deal with only piecemeal. His theory is also incomplete: it lacks, for example, a better account of how awareness of similarity makes possible a

---

[73] See also Kelley (1984), 345.

[74] It is also worth noting how similarity understood in light of Gibson and Prinz helps answer long-standing objections to similarity as the basis for an awareness of natural kinds, as expressed by Quine (1969b). Quine says there is something "logically repugnant" about similarity, understood in the comparative fashion mentioned in footnote 18, because of the inability to identify the properties in terms of which this similarity should be formulated. Quine repeats the familiar point that two objects are similar in any number of regards, and similarity threatens a promiscuity of natural kinds unless the relevant respects of similarity are constrained. Quine speculates that standards of similarity may be innate and determined by evolution. Here we can agree with him, but the types of similarity standards he assumes to be unique are primarily color-similarities, which is an age-old empiricist bias that has long been contradicted by developmental evidence. Children are far more impressed by similarities in shape than they are by color, and we can now understand *this* in evolutionary terms by reference to Gibson's understanding of ecological properties as forming the basis for the information pickup of perception. And when similarity is understood first in terms of shape, similarity also becomes much more respectable, scientifically. The shapes of objects have far more consequences for their causal interactions, which makes concepts formed on the basis of perceptual similarity "natural kind" concepts from the beginning, not just superficially-held "nominal essences." Science itself, then, to the extent that it begins, developmentally, with similarity concepts, is not, as Quine says "rotten to the core."

unitary grasp of seen and unseen referents of a concept. In fact I critiqued Prinz in chapter 3 for offering his theory as a purely naturalistic account of concepts. Perhaps the most overwhelming reason to consider mine to be a non-naturalistic approach is because I do not attempt to offer a naturalistic solution to the "*qua*" problem: the problem of which of many classes of things a categorized object is represented as belonging to. I raised that problem for higher-level concepts which proxytype theory could not address by itself, but it is worth mentioning that even proxytypes face their own version of the problem. Even though the mind *can* transform images in long-term memory networks in the ways described above, there is a question about whether it ever *does* transform them into all of the possible objects we take our concepts to refer to. Perhaps a naturalistic theory could be formulated in terms of dispositions to transform proxytypes, but disposition talk is often modally loaded and difficult for the naturalist to reconciling. It is here that I think that the first elements of irreducible intentionality may enter the picture. Perhaps part of what permits us to refer to an infinite range of objects is little more than our intention to regard the range of transformations as possible, and our intention to classify potential objects under that range.[75] This irreducible intentionality would also help erase the Sellars-Brandom concern that subject-predicate judgment involves mere differential response to one's environment. Since I am presenting a non-naturalistic theory of reference, it should be clear that we are not talking about merely differential response.

We are now approaching our dénouement. We have seen how scientific evidence, and some philosophic theorizing, lends credence to a theory of perceptual similarity that helps explain the abstraction of basic observational concepts from perception. We needed such a view of abstraction as an aid for understanding how perception could help justify judgments, in the manner McDowell thought to be impossible. We needed such a clarification about judgments to answer BonJour's concern that, even if perception did afford direct awareness of the world, it would not be such as to justify our judgments. To answer BonJour's concern, and illustrate how all concepts used in judgment

[75] See Rand (1990) for a developed version of this account.

might be justifiably applied, we need an account of concepts beyond simply the middle-sized object-concepts discussed so far. Particularly if we want an account of judgment that allows for inductive generalization, we need some account of action and attribute-concepts, one showing how these might be derived from perception.

In Gibsonian theory, not only are middle-sized objects or entities perceived, but some particular attributes, actions, and relationships are as well. (Other recent psychological work influenced by Michotte (1963) even suggests the possibility of perceiving *causal* relationships.[76] The perceiver perceives a total environment, and to the extent that any types of entity, attribute, action, or relationship have played an ecologically significant role in the organism's evolutionary history, it makes sense that the organism should come to develop a capacity to perceive them, i.e. to be able to *discriminate* one particular entity from another entity, one particular attribute from another, etc. To perceive an object is just to be able to see it as distinct from other objects against a shared spatial background. Perception of other properties is also relational in this way (contrary to classical empiricists): one does not perceive the particular shape of an object, for example, by discriminating its shape from its color, but instead by discriminating its shape from another object's shape. Perceptual similarity can be understood in terms of perceptual discrimination: two objects, for example, are perceived as similar if the respect in which they are perceived as differing (three-dimensional shape) is less than the respect in which they each differ from some third outlier. By a parallel act of double discrimination, we could also come to perceive similarities in attribute, action, or relationship. We can perceive similarities in shape if, for example, we perceive two objects of different types each of which is spherical (e.g., an orange and a baseball), as against everything else. Or we can even perceive similarities in action if, for example, we perceive two objects of different types each of which is rolling (again the orange and the baseball), as against other objects at rest.

[76] See Leslie and Keeble (1987), Scholl and Tremoulet (2000), and Prinz (2002, 173-177).

Once we appreciate the way perceptual similarity can be extended even to categories apart from entities, forming concepts from these similarities could proceed using at minimum something like a Prinzean proxytype. Applying these concepts in the act of judgment would then yield generalizations ("balls are round," "balls roll," etc.), seemingly capable of justification by reference to perceiving the same similarities that first permitted the formation of these concepts.[77]

We were first led down the path of looking for the foundations of induction by John Norton, who argued that as long as one inductive generalization could be understood as justified by a material fact induced by another, there need be no infinite regress of inductive justification as long as we appealed to more and more basic facts, eventually terminating in "brute" facts of experience. We have now gone some distance in explaining just what these "brute facts" might look like (literally). But we are now also in a position to appreciate and enrich a tantalizing suggestion of Norton's (2003, 668) in a final footnote:

> I reject the simple argument that such brute facts are always singular and that no collection of singular facts can license a universal. The problem is that canonical singular facts—"the ball is red"—already presuppose universal knowledge. In this case, it resides in the recognition that the thing is a ball and that its surface is red, thereby admitting recognition of commonality with potentially infinite classes of objects and colors.

We can now see just how true this is. The very formation of concepts is, in a way, a quasi-inductive process, inasmuch as we must abstract from similarities to refer to an open-ended class. The application of concepts in judgment therefore presupposes this universality.[78] This also enriches a suggestion from Lange (2004, 218):

---

[77] On non-Humean theories of causality, being able to generalize about the actions of entities (e.g., "balls roll," "knives cut") already constitutes causal knowledge (Harré and Madden 1975). On this view, causal connections are not primarily regularity relations between types of event, but relationships between types of actions and types of entities. It is worth noting that an additional source of Humean problems is the view that causal relations are only among events—where the events are interpreted psychologically as sensory atoms. With the rejection of the retinal image view of visual perception, however, the Humean metaphysics of events loses its psychological grounding. See Harré and Madden (1975, 49-67).

[78] It may be objected—following Hume (1975, 33)—that even if we possess general knowledge connecting subject and causal predicate, we do not know when to apply it to particulars, and hence we can never make predictions about singular concretes in the future or in unobserved distant locations. Hume's idea is that there are no known connections between "sensible qualities" and "secret powers," and so merely recognizing objects

> If we cannot make observations without undertaking commitments that go beyond them, then apparently the inductive skeptic must abstain from making *any observations at all*. Thus the problem of induction is circumvented, since it presupposes that we could be entitled to make observations without having any reason to infer inductively beyond them.

If this is right, we have succeeded not only in stopping the inductive skeptic's *reductio*, but in the course of doing so, we have shown that quite possibly the Humean skeptic about induction does not even have the logical right to assume the premises he did in the first place. He cannot assume the premise that we can formulate certain observation reports about singular facts, which are assumed not to imply generalizations, as there would be an outright contradiction involved in the idea that we might make observation reports but not imply any generalizations. This gives reason to wonder whether many of the other premises assumed for reductio might not involve similar contradictions with the skeptic's conclusion. It also gives further reason to wonder about how many of these doubts are prompted by purely scientific discoveries, rather than by philosophic presuppositions.

---

similar to acting objects in the past does not guarantee that they will act the same way or possess the same powers. There is some truth to this point: grasping a similarity is certainly only a defeasible source of justification, and we can make mistakes. But at this point, Hume is at the end of his rope. The objection that we might be mistaken about the application of a generalization to a particular is no longer a special problem for induction. It is really no different than the simple argument from illusion concerning the bent stick in water. Just as we might be subject to shape illusions, so we might also be subject to causal illusions. But this is no special problem for our knowledge, provided that we understand that the error is one of judgment, not the senses, and that it is the senses which provide the ultimate source of understanding and correcting our errors. Just as we can remove the stick from water to see if it is really bent, we can place our bread under a microscope to discover if it has the ordinary chemical sub-structure that would nourish us. Of course we are subject to error when we make judgments about these observations, too, but the general possibility of error does not imply that we might be wrong in any particular case. ("It is possible for me to err" does not imply "It is possible that I am in error now.") Here we are really treading on the grounds of Cartesian skepticism, rather than Humean skepticism. The most I can say is that the manner in which the Cartesian skeptic generates possible sources of error is akin in many important ways to Quine's underdetermination problem. Just as multiple rival theories "possibly" might account for the same observational evidence according to Quine, there might also be different "possible" possible sources of error regarding evil demons, etc., might account for what we perceive. I think each of these arguments is flawed for many of the reasons I have already identified in my section on the underdetermination thesis.

Conclusion

This essay began by considering a strategy, sometimes attributed to Quine, for refuting skepticism by pointing to the scientific assumptions made by skeptics for the sake of *reductio*: If the skeptic has the right to assume science, then the anti-skeptic also has the right to assume science to defuse the absurd consequences of the skeptic's more scientifically isolated assumptions. This strategy does not itself attempt to prove any propositions of philosophy, but simply to dismiss motivations for various kinds of skepticism about philosophical propositions. I have sought to show how this strategy may be trained upon more an even radical skepticism than has been imagined, including Quine's own form of skepticism, the underdetermination thesis. But the underdetermination thesis was only a default position accepted as a result of previous skepticisms about induction and empirical foundations, so I sought to show how a similar anti-skeptical strategy can be applied against these positions, as well.

I have argued above that the Quinean anti-skeptical strategy of appealing to more science does have some limited utility in dissolving Humean doubts about induction, and Sellars-Davidson doubts about the perceptual basis of induction. But my emphasis, particularly in this last section, has been that scientific evidence cannot do the whole job, because not all of the doubts are themselves scientific. Even if we cannot resolve these doubts, my concern then is that naturalized epistemology cannot rely on these doubts to motivate its program. If skepticism seems unavoidable only because philosophers have idiosyncratic presuppositions about mind and representation, then insofar as skepticism prompts naturalism, naturalism is prompted for non-naturalistic reasons. This works to undermine the coherence of the naturalistic view that philosophy works in a manner that is purely continuous with the scientific. If purely philosophical problems got us into skeptical trouble, it may take philosophical solutions to get us out.

But there is some legitimate concern over what a "purely philosophic" solution would be. In the course of proposing some of these solutions, I have made appeals to our understanding of concepts like “perception,” “inference,” “concept” and “intentionality.” I have argued that skeptical arguments fail because of their misapplication of these concepts. Nevertheless I want to emphasize that this methodology needn’t be *a priorist*, even if it is “purely philosophical.” Pure philosophy is only *a priori* if what separates philosophy from science is that only the second involves the empirical method. Other methods of demarcation are possible. It is possible for philosophic knowledge to be empirical, if only more fundamental and universal than scientific empirical knowledge. Philosophy may examine facts—and the concepts that designate them—that are of fundamental important to scientific and all other kinds of knowledge. If we gain better understanding of our concepts of these facts by understanding how they arise from our experience (but extrospective and introspective), rather than through “intuition,” all the better.

Understanding philosophic method this way, while empirical, is also not “naturalistic.” Naturalism takes scientific empirical knowledge as the *exclusive* form of knowledge. My approach certainly takes science as the highest form of empirical knowledge, but not as the exclusive form. Central to the doctrine of empirical foundationalism I have explored, after all, is that there might be pre-scientific commonsense empirical knowledge, on which science itself is based. Philosophic knowledge itself may be of this pedigree, though of a higher order than most pre-scientific knowledge, and certainly more self-conscious. A full articulation of this approach to philosophic knowledge, in particular its empirical approach to “conceptual analysis” is far beyond the scope of this dissertation. Its development will have to wait for another day. But acknowledge its possibility is an important antidote to the view, so often endorsed by naturalists, that the choice is between *a priori* conceptual analysis and scientistic naturalism. This is a false choice, and I hope that much of the work of this chapter—and dissertation—will help to show why.