

Evidence: Fundamental Concepts and the Phenomenal Conception¹

Thomas Kelly
Princeton University

The concept of evidence is among the central concerns of epistemology broadly construed. As such, it has long engaged the intellectual energies of both philosophers of science and epistemologists of a more traditional variety. Here I briefly survey some of the more important ideas to have emerged from this tradition of reflection. I then look somewhat more closely at an issue that has recently come to the fore, largely as a result of Williamson (2000): that of whether one's evidence supervenes on one's non-factive mental states.

1. Terminology and Fundamental Concepts

Consider some scientific theory that currently enjoys widespread acceptance among members of the relevant scientific community—for example, the theory of continental drift. What makes the theory of continental drift worthy of being believed by contemporary geologists? A plausible and popular answer is that the evidence possessed by contemporary geologists strongly suggests that it is true. Paradigmatically, those theories that are worthy of being believed enjoy such status in virtue of the availability of evidence sufficient to justify belief in their truth.

Evidence that supports or tells in favor of a given theory *confirms* that theory. On the other hand, evidence that tells against a theory *disconfirms* that theory. Of course, a given piece of evidence might confirm or disconfirm a theory to a greater or lesser degree. The term 'verification' has traditionally been used to signify the maximal degree of confirmation: evidence *verifies* a theory in the relevant sense just in case it conclusively establishes that the theory in question is true. At the opposite end of the spectrum, *falsification* signifies the maximal level of disconfirmation: evidence falsifies a theory just in case it conclusively establishes that the theory in question is false.

¹For helpful comments on an earlier version, I am indebted to David Christensen.

Just as different pieces of evidence can confirm a theory to different degrees (i.e. some pieces of evidence which confirm a theory might confirm it more strongly than other pieces of evidence which also confirm it), there are interesting and important differences among the ways in which theories are confirmed by evidence. Thus, a given piece of evidence might lend support to a theory directly (as when the theory provides a particularly good explanation of that piece of evidence), or more indirectly, by, e.g., disconfirming or falsifying some otherwise formidable rival theory. Thus, the observation that certain substances gain weight upon undergoing combustion constituted particularly strong evidence for the oxygen theory of combustion, inasmuch as the phenomenon in question seemed to strongly disconfirm its main rival, the phlogiston theory of combustion. As this example illustrates, there is no simple relationship between how strongly a given piece of evidence confirms a given theory and the directness or indirectness with which it bears on that theory (at least in the intuitive sense of ‘directness’ at issue here): some of the strongest evidence for a theory might bear on the theory rather indirectly, while evidence that bears more directly on the theory might be relatively weak.

In considering questions about how a given body of evidence bears on a theory, it is crucial to distinguish between the *balance* of the evidence and its *weight*. Intuitively, the balance of the evidence concerns how decisively the evidence tells for or against the theory. On the other hand, the weight of the evidence is a matter of how substantial the evidence is. As one acquires more evidence which bears on the theory, the weight of one’s evidence increases; this may or may not make a difference to the balance of one’s evidence, or to what it is reasonable to believe on its basis. Suppose that I select a coin at random and flip it twice; it lands ‘heads’ on the first toss and ‘tails’ on the second. I am then informed that the same coin will be tossed again exactly one year from now. How much credence should I give to the proposition that the coin will land heads on that occasion? Given the evidence available to me, it would seem reasonable to invest credence .5 in the relevant proposition. Suppose that I subsequently flip the coin several thousand more times. I carefully keep track of the outcome of each flip; it turns out that the proportion of ‘heads’ to ‘tails’ is well within the range that one would expect on the assumption that the coin is fair. How much credence should I *now* give to the proposition

that the coin will land heads when it is flipped one year from now? Answer: .5, for the balance of my evidence has not appreciably changed. However, although the balance of my evidence has not changed, its weight has increased significantly.

As this example suggests, while the balance of one's evidence with respect to a given proposition is typically reflected directly in how confident it is reasonable for one to be that that proposition is true, the weight of one's evidence is not. Informed of what someone believes about some question on the basis of her evidence, and assured that she is reasonable in believing as she does, one is typically in a position to draw inferences about the balance of her evidence but not about its weight. (That is, one has no way of knowing how many times I have seen the coin flipped, told only that I reasonably give credence .5 to the proposition that the coin will land heads when it is flipped one year from now.) However, although the weight of one's evidence is not directly reflected in what it is reasonable for one to believe about the target proposition, it often manifests itself in what it is reasonable to believe when additional evidence is acquired. Thus, imagine that the coin is flipped seven consecutive times and lands 'tails' each time. Taken by itself, this sequence suggests that the coin is biased in favor of 'tails'. If prior to observing the sequence my credence for the proposition that *the coin will land heads when it is flipped in one year's time* stood at .5, how much should I reduce my credence in the light of this new information? Not much, if my prior credence was based on having observed several thousand flips during which the coin behaved in the manner of a fair coin. More, if my prior credence was based on having observed a mere two trials during which the coin behaved as though it were fair. The general moral: the weight of one's evidence with respect to some proposition tends to manifest itself, not in how confident it is reasonable for one to be that that proposition is true, but in how it is reasonable to respond when one acquires additional evidence which bears on the question.²

² An excellent, technically sophisticated recent discussion of the distinction between balance and weight, the relationship between each and rational belief, as well as related issues is Joyce (2005). A classic account of the distinction is Keynes (1921). Among other relevant discussions, see especially Skyrms (1980) on the 'resilience' of credences in the face of additional data.

In some cases, the available evidence might *underdetermine* the choice between rival theories in a way that makes suspension of judgment among them the reasonable response. Some cases of underdetermination are uncontroversial. When my evidence suggests that the coin is fair, my evidence underdetermines the choice between the hypothesis that *the coin will land heads the next time it is flipped* and the rival hypothesis that *the coin will land tails the next time it is flipped*. Given that my evidence underdetermines the choice between these two rival hypotheses, the uniquely reasonable response to this state of affairs is for me to suspend judgment on the question of whether the coin will land heads or tails on the next flip. Presumably, much experimentation within the sciences is designed and conducted with an eye towards eliminating underdetermination of this general kind—although typically with respect to questions of greater theoretical interest and importance.

Not all underdetermination claims are uncontroversial. A favorite tactic of skeptics is to claim that our evidence will inevitably underdetermine the choice between our ordinary, common sense views and various skeptical scenarios to which he calls our attention. Thus, the skeptic about our knowledge of the external world maintains that one's evidence (understood, perhaps, as the totality of one's current experiences) does not favor one's ordinary, common sense views about one's surroundings over various skeptical alternatives (e.g., the hypothesis that one is hallucinating in an undetectable way). Other traditional epistemological skepticisms—about, say, induction, or other minds—can also be illuminatingly reconstructed as arguments from underdetermination and are sometimes explicitly presented as such. Controversial and far-reaching underdetermination claims also figured prominently in much 20th century philosophy of science. For example, a recurrent theme in the writings of W.V. Quine is the claim that any possible body of observational data will inevitably underdetermine the choice between rival theories. (See, e.g., Quine (1980)). Sympathy for such claims has contributed significantly to the popularity of various non-realist accounts of science.

Of course, both the plausibility and interest of claims to the effect that evidence inevitably underdetermines theory choice in some domain depend on exactly what is meant by 'underdetermine'. The weakest interpretation of those that sometime seem intended is that our evidence will inevitably *logically* underdetermine the choice between

theories, in the sense that there will always be more than one possible theory that is logically consistent with our evidence. This claim, which in effect amounts to the claim that our evidence does not typically *entail* some particular theory, seems unimpeachable. (In post-positivist philosophy of science, it is universally conceded that the content of our best scientific theories far outstrips the content of the evidence which is taken to support them; indeed, increasing recognition of this fact was a primary driving force in the development of positivism itself.) However, logical underdetermination in this sense should not be confused with underdetermination in the sense which warrants suspension of judgment. For the idea that logical underdetermination warrants suspension of judgment should seem attractive only to those who espouse a particularly crude form of inductive skepticism, according to which the only reasonable inferences are those which mirror relations of logical entailment. (A coin is flipped several thousand times, in a wide variety of circumstances; each time, the coin lands ‘heads’. This evidence logically underdetermines the choice between the hypotheses that

- (i) the coin is biased in favor of heads
- (ii) the coin is biased in favor of tails, and
- (iii) the coin is fair

inasmuch as it does not logically entail any of the three. But inasmuch as agnosticism among the three hypotheses is not the uniquely reasonable response to the evidence, the choice between the hypotheses is not underdetermined by the evidence.)

A different underdetermination claim that is sometimes advanced is the following: our evidence will inevitably underdetermine the choice between theories in the sense that, for any given body of evidence, there will always be more than one *potential explanation* of that evidence—roughly, more than one alternative theory which would, if true, account for our evidence being as it is. This underdetermination claim, although significantly stronger than the one just considered, is also quite plausible. For consider some body of evidence $E_1 \dots E_n$ and some theory T which is a potential explanation of that evidence. Consider next the alternative theory T^* , according to which

T*: A Cartesian evil demon is making everything seem exactly as though theory T is true.

Plausibly, theory T* is a potential explanation of the evidence just in case theory T is: that is, Theory T* is the sort of thing which, if true, would suffice to account for the evidence E1...En, given that E1...En is potentially explained by the original theory T. But once again, no non-skeptic should concede that the sense in which evidence E1...En underdetermines the choice between T and T* is the same sense of underdetermination in which underdetermination requires suspension of judgment. That is, no non-skeptic should concede that a body of evidence equally confirms any hypothesis that potentially explains it.

As is perhaps suggested by what has been said thus far, an underdetermination thesis that would seem to have genuine skeptical implications is the following:

For any body of evidence, there will always be alternative theories that constitute *equally good* explanations of that evidence.

However, no serious attempt to establish anything like this thesis has ever been made. Indeed, it is unobvious how one might go about attempting to establish this thesis, given that we do not yet have anything like a fully adequate descriptive account of the factors that constitute explanatory goodness.³

Evidence which underdetermines the choice between two rival theories is evidence which does not favor either theory over the other. In contrast, *misleading* evidence is evidence which suggests that something which is in fact false is true. Thus, suppose that my goal is to frame the butler: I want others—especially those in positions of power and authority—to believe that the butler is responsible for some particularly heinous crime. Given this goal, my best strategy will typically be to plant misleading evidence that suggests that the butler is guilty. Success in the endeavor would consist, for example, in arranging for his fingerprints to be discovered on the murder weapon, or for his blood to be found at the scene of the crime.

³Recent discussions of underdetermination are vastly more sophisticated than many earlier ones; for a sampling, see Earman (1993), Laudan (1991,1996), Okasha (2002), Hofer and Rosenberg (1994), and the relevant essays in Lipton (1995).

Although misleading evidence suggests what is in fact false, it is nonetheless genuine evidence, in the sense that it satisfies the conditions for being evidence (whatever those conditions are). While counterfeit money is not money, and fool's gold is not gold, misleading evidence is evidence, no less than non-misleading evidence is. The fact that misleading evidence is genuine evidence is why beliefs based on misleading evidence can be reasonable, given that what it is reasonable to believe depends on one's evidence. When the authorities are taken in by my flawless deception and come to believe that the butler is guilty, their evidence-based belief that he is guilty is no less reasonable for being false.

As one's evidence changes over time, what it is reasonable for one to believe on the basis of that evidence undergoes corresponding changes. Suppose that at some later point in time I confess to having framed the butler. If the authorities continue to believe that the butler is guilty in the wake of my confession, then their doing so is unreasonable. An obvious moral: to the extent that what it is reasonable to believe depends on one's evidence, what is relevant is the bearing of one's *total* evidence. It is tempting to think that if a given body of evidence $E_1 \dots E_n$ is sufficient to justify believing p , and one's belief that p is based on $E_1 \dots E_n$, then one's belief that p is justified. However, this is a mistake. For even if $E_1 \dots E_n$ is sufficient to justify believing p when considered in isolation (that is, $E_1 \dots E_n$ would not need to be supplemented by some further piece of evidence or reason for thinking that p is true), $E_1 \dots E_n$ might constitute only a proper subset of one's total evidence. And one's total evidence might include some further piece of evidence E^* such that one is *not* justified in believing p given evidence $E_1 \dots E_n$ and E^* . Thus, perhaps when the only relevant evidence possessed by the authorities consists of

E1 The butler's fingerprints were found on the murder weapon

and

E2 The butler's blood was found at the scene of the crime

it is reasonable for them to believe that the butler committed the crime; no additional evidence of the butler's guilt is required in order for the relevant belief to be a reasonable one. Moreover, notice that, after my later confession, E1 and E2 are still true, and presumably, known to be true by the authorities. Nevertheless, the belief that the butler committed the crime would be unreasonable in the wake of my confession, even if both before and after my confession the relevant belief is based on both E1 and E2. This is because after my confession, the total evidence available to the authorities includes the information that

E3 TK has confessed to framing the butler

and what it is reasonable to believe is always relative to one's total evidence.⁴

When I confess to the authorities, they gain new evidence which is relevant to the question of criminal's identity. This new evidence stands in a particularly interesting relationship to the evidence that they already possessed. Inasmuch as the new evidence E3 undermines the justification for believing that the butler committed the crime that would otherwise be afforded by E1 and E2, evidence E3 constitutes *defeating* evidence: it defeats the justification for holding the relevant belief that would have been afforded by E1 and E2 in its absence.

Significantly, defeating evidence is itself susceptible to being defeated by yet further evidence. Thus, suppose that my confession was made in response to credible threats of physical violence. At some later point in time, the authorities might thus learn that

E4 TK's confession was coerced

Perhaps the addition of E4 to the total evidence possessed by the authorities undermines the probative value of my confession; if so, then E4 constitutes defeating evidence with respect to E3 and thus tends to reinstate the justification for believing that the butler committed the crime afforded by the original evidence of his guilt.

⁴ A classic discussion of the requirement of total evidence is Hempel (1960).

Clearly, different individuals will typically differ significantly in the total evidence that they possess, and the total evidence possessed by any one individual will change significantly over time. Given that what it is reasonable for one to believe at any point in time depends on the total evidence which one possesses at that time, what it is reasonable to believe will be a highly relativized matter. Inasmuch as even contemporaries who inhabit the same cultural circle will differ significantly in their particular life histories and the more specific ways in which they are embedded in the world, such individuals will typically differ significantly in the evidence that they possess, and therefore, with respect to what it is reasonable to believe. We would thus expect to find even greater differences between individuals who inhabit radically different cultures or who are far removed from one another in space and time. The typical human being alive at the turn of the 21st century possessed vast troves of evidence that were simply unavailable to human beings who lived in earlier epochs. By the same token, much evidence that was readily available to those who lived centuries earlier has no doubt been forever lost to us. The fact that individuals who inhabit different epochs or cultures will differ significantly—indeed, in many cases, differ radically—in what it is reasonable for them to believe has on occasion encouraged the view that what is true is itself relative to time or culture. But this is a vulgar confusion. The uncontroversial kind of epistemic relativism discussed here, while undoubtedly a fact of great significance in many ways, lends no support to relativism about truth. The fact that it might be reasonable for the authorities to believe that the butler committed the crime given the total evidence available to them, while it is reasonable for me to believe that *I* committed the crime given the total evidence available to me, does nothing to show that the truth about who committed the crime might differ from individual to individual, or be relative to one's perspective on the world.

What has been said thus far I take to be relatively uncontroversial. Unfortunately, here as elsewhere in philosophy, agreement runs out quickly: some of what would seem to be among the most fundamental questions about evidence are matters of ongoing, substantive dispute. Here are three such questions.

(i) The first question concerns what we might call the *ontology* of evidence. What kinds of things are eligible to count as evidence? Prominent candidates within the philosophical tradition include sense data, experiences with representational content,

token beliefs, facts, propositions, and, in Quine's distinctive brand of empiricism, 'the stimulation of one's sensory receptors' (1969: 75). Here, a key division is between those who think that all evidence is *propositional* and those who think that at least some evidence is more *object-like* than proposition-like. According to one prominent version of classical foundationalism, our evidence ultimately consists of sense data, mental particulars with which we are directly acquainted. On this traditional view, all of the evidence that we ultimately have to go on is non-propositional. More recently, however, the debate has been between those who hold that all evidence is propositional (e.g., partisans of facts, token beliefs, and experiences with representational content) and those who hold, more liberally, that some of our evidence is propositional and some non-propositional.⁵

A second important division under this general heading is between theorists who hold that what evidence one has is completely fixed by one's non-factive mental states and theorists who deny this. This division between 'evidential internalists' and 'evidential externalists'⁶ cuts across the propositionalist/non-propositionalist distinction. Thus, a proponent of the classic sense data theory mentioned above will hold that one's evidence is non-propositional and entirely a matter of what mental states one is in; this second, internalist commitment will be shared by various paradigmatic propositionalists, e.g., a coherentist who holds that one's evidence consists of token beliefs that are sufficiently well-integrated with one another. On the other hand, a propositionalist who thinks that one's evidence consists of all and only those propositions that one knows (e.g., Williamson 2000) will think that what evidence one has depends not only on what non-

⁵ Williamson (2000:194-200) argues that all evidence is propositional; Neta (forthcoming) offers some trenchant criticisms of those arguments and attempts to supply stronger ones for the same conclusion. Plantinga (1993) defends the view that some evidence is non-propositional; that is also the view of Conee and Feldman (2004).

Given the synonymy or near synonymy of 'evidence' and 'reasons for belief', the question is essentially the familiar one of whether only propositions can be reasons for belief. See Unger (1975: 204-206) and Davidson (1986) for affirmative answers to this question; Moser (1989: 47-125) and Millar (1991) take the contrary view.

⁶ The terminology is due to Silins (2005). Representative evidential internalists would include Bonjour (1999) and Audi (2001); representative externalists would include McDowell (1982, 1995) and Williamson (2000).

factive mental states one is in but also on how things stand in the external world. This ‘externalism’ about evidence will also be endorsed by a non-propositionalist who takes at face value the many ordinary assertions which suggest that physical objects can themselves count as evidence in certain contexts.⁷ After some sharpening, this issue will be our primary focus in Section 2 below.

(ii) Another question concerns the relation that must obtain between an individual and a piece of evidence E in order for E to count as part of that individual’s total evidence. Thus far, we have spoken of individuals *possessing* evidence. But it is far from clear what it is to possess evidence in the relevant sense. Presumably, the way in which one possesses one’s evidence is not a matter of *physical* possession, in the way that one might be in physical possession of one’s car keys at a given moment in time. Nor does legal possession seem to be a good model here: the way in which one possesses one’s evidence does not seem much like the way in which one possesses one’s property. What relation does one bear to a piece of evidence, when one possesses that evidence in the relevant sense? Of course, the answer that a theorist gives to this question will depend a great deal on how she answers the first question, i.e., on her view as to what sorts of things are eligible to serve as evidence. For example, according to classic sense data theory, one possessed one’s evidence in virtue of being *acquainted* with it, where acquaintance was a *sui generis*, unanalyzable relation of perfect psychological immediacy that one bears to the sensations that one is having at a given moment in time. On the other hand, suppose that a theorist holds, in answer to the first question, that an individual’s evidence consists of some body of propositions. The question will then arise: which set of propositions is it? The set of propositions that the individual believes? The set of propositions that she believes with certainty? Or perhaps: the set of propositions that she ought to believe or would be justified in believing? Or—yet another possibility--the set of propositions that she knows? Each of these answers has found its defenders.

(iii) A third question concerns the relation of confirmation itself: what relation must E stand in to H in order for it to be true that E confirms H? With respect to this question, one answer has emerged as more popular than any competitor. According to this view,

⁷ Consider, for example, the way in which the murder weapon itself is routinely referred to as evidence during a criminal trial.

confirmation consists in increase in probability: evidence E confirms hypothesis H just in case the probability of H conditional on E is greater than the unconditional probability of H. (Similarly, E disconfirms H just in the case the conditional probability of H on E is lower than the unconditional probability of H.) This probabilistic explication of confirmation, which is central to Bayesian epistemology, enjoys fairly widespread support. Even here, however, the extent of the agreement should not be overestimated. In addition to those who deny that the relation of confirmation should be explicated in probabilistic terms, those who accept the underlying idea often differ among themselves regarding the key issue of what kind of probability is relevant. Should confirmation be explicated in terms of *subjective* probability, as the orthodox Bayesian suggests? Or is some more objective variety of probability called for?⁸

As in the case of question (ii), how a theorist answers question (iii) typically constrains, and is constrained by, her views about the ontology of evidence. For whatever evidence is, it must be the kind of thing which is eligible to serve as the first *relatum* in instances of the confirmation relation. So, for example, theorists who are committed to a probabilistic explication of the confirmation relation sometimes argue that evidence must be proposition-like on the grounds that only entities with propositional structure can stand in probabilistic relations.

Each of these questions is a large one about which much could be said. Indeed, a thorough discussion of any one of the three would be well beyond the scope of the present essay. In what follows, we will look at one of the issues that arises with respect to the ontology of evidence: that of whether one's evidence supervenes on one's non-factive mental states, or whether what evidence one possesses typically depends on how things stand in the external world. In addition to its intrinsic interest, this is an issue on which the dominant philosophical tradition differs sharply from the view of evidence that seems implicit in both scientific practice and in everyday life. Moreover, as we will see, certain recent developments in philosophy have conspired to push this issue towards center stage.

⁸ Representative examples of the former approach include Jeffrey (1992, 2004) and Howson and Urbach (1993); representative examples of the latter include Horwich (1982) and Williamson (2000). The papers collected in Achinstein (1983) are largely concerned with questions about the relation of confirmation.

2. The Phenomenal Conception of Evidence

Intuitively, one's evidence is *what one has to go on* in arriving at a view. Evidence is what Sherlock Holmes carefully collects and surveys, and that from which he ultimately infers the identity of the person who committed the crime. As such examples suggest, talk of evidence seems most at home in contexts in which the truth would otherwise be unobvious. Plausibly, there are some propositions whose truth or falsity we grasp in an utterly direct, unmediated way. Consider, for example, simple arithmetical truths such as the proposition that $2+2=4$. Traditionally, such truths have been claimed to be 'self-evident': allegedly, they need only be understood in order to be known. If the truth value of every proposition were transparent in this way, perhaps we would have little or no use for evidence. In contrast, a central function of evidence is to make evident that which would not be so in its absence.

In general then, evidence seems to play a mediating role vis-à-vis our efforts to arrive at an accurate picture of the world: in cases in which the truth is not transparent, we seek to believe what is true by way of holding beliefs that are well-supported by the evidence, and we seek to avoid believing what is false by way of not believing that which is not well supported by the evidence. Thus, one's recognition that *the earth is roughly spherical in shape* seems to depend on one's evidence in a way that one's recognition that $2+2=4$ does not. Of course, it can be a contested matter whether one's access to truth in some domain is problematic—and thus, whether one is dependent upon evidence for grasping truths about that domain. For example, common sense holds that we often have unproblematic access to facts about our immediate physical environment *via* sense perception. In contrast, much traditional epistemology holds that one's access to such truths is always deeply problematic; what is unproblematic, rather, is one's recognition that one's experiences represent the world as being a certain way. Hence, much traditional epistemology construes the relationship between one's experiences and one's beliefs about the physical world on the model of the relationship between evidence and a scientific hypothesis.

Suppose that I lack the capacity to detect the presence of Xs directly: to the extent that I am able to detect their presence at all, I must rely on evidence in order to do so. Perhaps I treat the presence of Ys as evidence, increasing my confidence that Xs are present whenever I judge that Ys are. What must be true of the Ys, in order for them to effectively play this role? Here are two plausible requirements. First, there must be a positive correlation between the presence of Xs and the presence of Ys. Second, it must be, at least often enough, *easier to recognize* that Ys are present than that Xs are present. For even if I know that the presence of Ys guarantees that Xs are also present, this will avail me little unless I can detect the Ys more readily than the Xs.

These two requirements—that evidence should be the kind of thing which is, in general, a reliable indication of that for which it is evidence, and that it should be relatively easy to recognize—can pull the theorist of evidence in opposite directions. When one is dependent on evidence for recognizing truths in some domain, there is a non-trivial gap between the truths in question and what one would be in a position to know in the absence of evidence; the function of evidence is to close or narrow the relevant epistemic gap. If evidence is conceived of in such a way so that it is unproblematically and immediately given to the subject, then it will be easy enough for the subject to recognize her evidence, but doubts may arise about whether evidence so understood is the kind of thing which can adequately close the original gap. Thus, consider views on which one's evidence is exhausted by how things currently seem or appear to one. No doubt, one is almost always in a position to appreciate how things currently seem or appear. But among the facts which we ordinarily suppose can be known on the basis of evidence are various facts about the distant past, e.g., that *the Battle of Hastings occurred in the year 1066*, or, more distantly still, that *dinosaurs once roamed the surface of the earth*. Offhand, the gap between facts such as these and facts about how things seem or appear to one at the current moment would seem to be quite wide. Indeed, one might reasonably worry that no amount of evidence would put one in a position to know facts about the distant past, so long as evidence is understood to consist exclusively of facts about how things currently seem or appear. On the other hand, if evidence is conceived of in a more liberal manner, so that it is more closely tied to the target subject matter, concerns might naturally arise about whether one possesses the

requisite access to evidence so understood; the risk is that a new gap will be introduced, this one between the person and the putative evidence itself. For example, perhaps certain facts about the slightly less distant past strongly reflect facts about the slightly more distant past. But are facts about the slightly less distant past the sort of thing that are given to one in such a way that they can credibly be claimed to be among one's evidence? That is, are such facts really part of what one has to go on?

On the whole, the epistemological tradition since Descartes has opted for conceptions of evidence on which its availability to the person whose evidence it is is maximally unproblematic; the pressing challenge has been to show that evidence so understood is still the kind of thing which, when amassed in sufficient quantities, puts one in a position to know truths about the target subject matter. In particular, the dominant tradition has tended to identify one's evidence with some proper subset of one's non-factive mental states—for example, one's experiential states, or perhaps, one's experiential states together with one's beliefs. The challenge has been to explain how such an evidential base is adequate to underwrite our knowledge of the external world of tables, chairs, and other people—to say nothing of, e.g., our knowledge of the objects of theoretical science or of past historical events, both of which would seem to stand at yet more daunting epistemic remove. Famously, for Descartes that which is inner is epistemologically prior to that which is outer. In this respect if not in others, he was closely followed by the classical empiricists: for Locke, Berkeley, and Hume, what one ultimately has to go one in forming a view of the world are mental entities, impressions and ideas. Essentially this picture was bequeathed by the classical empiricists to much twentieth century analytic philosophy, although those working in the latter tradition tended to couch their theories in terms of 'sense data' or 'experiences'.

On this traditional picture of evidence, what evidence one has supervenes on one's non-factive mental states. Let us call this view the **phenomenal conception of evidence**.⁹ According to the phenomenal conception of evidence then, the evidence that Holmes ultimately has to go on in solving the crime does not consist in facts about the external world, such as the fact that the suspect's fingerprints were found on the murder

⁹ The name is from Williamson (2000).

weapon. Rather, it consists in such things as Holmes' visual experiences as of the suspect's fingerprints being on the murder weapon, or (when such visual experiences are no longer occurrent) Holmes' apparent memories of such experiences.¹⁰ But why suppose that Holmes' evidence is limited in this way, and does not include seemingly relevant facts about the external world which he has observed?

There are at least two closely-related reasons for the historical popularity of the phenomenal conception. First, a substantial portion of the history of epistemology since Descartes has been devoted to the consideration of radical kinds of skepticism, and in particular, skepticism about our knowledge of the external world. When such skepticism is salient, our access to what would ordinarily seem to be unproblematically available, mundane facts about the external world naturally comes to seem less and less straightforward. Of course, if skepticism is true, then we do not have cognitive access to facts about how things stand in the external world, and so such facts cannot play the role of evidence. But even if such skepticism is ultimately rejected, and the common sense view that we do have ample knowledge of the external world upheld, the skeptical challenge might still make our access to such facts seem to be too much of a cognitive *achievement* for such facts to be theoretically attractive candidates for being among the things which we have to go on. If, for example, it turned out that our only access to facts about the external world was *via* some kind of theoretical or inductive inference, then such facts would seem to be less natural candidates for playing the evidence role. For in that case, it would be more natural to identify one's evidence with whatever it is from which facts about the external world are inferred, rather than with the inferred facts themselves.

This brings us to a second, closely-related reason for the historical popularity of the phenomenal conception, a reason which concerns the picture of perception that dominates

¹⁰ The idea that paradigmatic evidence consists of token experiences is one which will find favor among traditional foundationalists; coherentists, on the other hand, will tend to identify Holmes' evidence with (e.g.) his *belief* that the suspect's fingerprints are on the murder weapon (provided that that belief is sufficiently well-integrated with Holmes' other beliefs). On either view, the supervenience of Holmes' evidence with his non-factive mental states is upheld. Thus, the phenomenal conception of evidence is a common commitment of both coherentism and the historically most popular version of foundationalism.

the broadly empiricist tradition. On the picture in question, which is common to Locke, Berkeley, and the rival sub-traditions which followed each philosopher, the immediate objects of perception are mental entities. Inasmuch as it is quite natural to think of evidence, at least in paradigm cases, as that which one directly observes as opposed to something which one infers from what one directly observes, such a picture of perception naturally encourages the phenomenal conception of evidence. In short, as long as a common sense view of perception is in place, according to which Holmes can directly observe various facts about how things stand in the external world, it is natural to think of such facts as part of Holmes' evidence. However, once this common sense view of perception is displaced on the grounds that it amounts to an untenable 'naïve realism', it becomes natural to think of Holmes' evidence in a different way as well.

Of course, the fact that there is an historical explanation for the dominance of the phenomenal conception of evidence does not mean that we have compelling reason to accept it. Indeed, I believe that we should reject the phenomenal conception and adopt in its stead a way of thinking about evidence that bears a greater resemblance to the notion as it is employed outside of philosophy, e.g., in scientific inquiry and in ordinary life. While some of the reasons for doubting the phenomenal conception are venerable ones, others have emerged only lately, as a result of recent developments in the theory of perception and elsewhere. In the remainder of this essay, I sketch what I take to be the current dialectical situation with respect to this issue.

One reason for suspicion of the phenomenal conception has already been noted: if our evidence consists exclusively of non-factive mental states, it is far from clear that it provides us with an epistemic foothold in the world sufficient to underwrite the knowledge that we ordinarily take ourselves to have. For the broadly Lockean tradition of 'representative' or 'indirect' realism, what one is immediately given in perception is not the chair itself but rather a mental representation of the chair; whatever knowledge one has of the chair itself is in effect inferential knowledge, knowledge which is inferentially based on what one ultimately has to go on. Again, on this view, our cognitive relationship to ordinary objects such as tables and chairs resembles the scientist's relationship to theoretical entities such as electrons and quarks on a standard scientific realist account of the latter.

From the beginning, this picture has consistently aroused suspicions that it ultimately leads to skepticism. Indeed, the task of accounting for our knowledge of the external world given such a picture of our evidence has often seemed so daunting that at various historical junctures its proponents have found themselves on the defensive against idealists and phenomenologists, for whom the relationship between our experience of the world and the world itself is more akin to a relationship of constitution than an evidential one. Consider, for example, the following argument that representative realism leads to skepticism, an argument which was frequently rehearsed during the heyday of twentieth century phenomenology.¹¹ According to the representative realist, our perceptual knowledge of tables and chairs is in effect inductive knowledge, knowledge which rests upon an evidential base consisting of our experiences. But inductive inference is at bottom a matter of projecting previously observed correlations to new cases. One sees the lightning and subsequently hears the thunder; after observing both lightning and thunder together a sufficient number of times, one learns to treat the occurrence of lightning as evidence for the occurrence of thunder, and one justifiably increases one's credence that one will hear thunder in response to seeing lightning. However, on the representative realist's own picture, we are never in a position to draw conclusions about the external world on the basis of our experiences in this way. For on the picture in question, we never have the opportunity to directly observe tables and chairs themselves; rather, what we are given are certain table- and chair-like experiences. Hence, we are never in a position to observe that, e.g., our chair-like experiences generally occur when we are in the presence of chairs, for we lack any independent access to the latter. It seems then that we have no way of arriving at the kind of inductive base that would license us in justifiably concluding that we are in the presence of a chair on the basis of undergoing chair experiences. Perhaps God would be in a position to observe that one's chair experiences typically occur when one is in the presence of chairs; if so, then He would be in a position to justifiably conclude that one is in the presence of a chair from the fact that one is currently undergoing chair experiences of a suitable sort. But of course, we do not occupy the God's eye point of view. Hence, the proponent of the

¹¹ A classic statement is Ayer (1936), who uses it to motivate the semantic thesis that sentences about tables and chairs are equivalent to sentences about sense data.

argument concludes, the representative realist picture of our relationship to the external world leads inevitably to skepticism.

As it stands, this phenomenalist argument is not compelling. After all, it is admitted on all sides that no one has ever directly observed subatomic particles such as electrons and quarks. On the assumption that scientists know that such particles exist on the basis of their observed effects, it must be possible to gain knowledge of some kind of entity *via* induction, even if one has never actually observed entities of the relevant kind correlated with that which one treats as evidence of their presence. Thus, the phenomenalist argument rehearsed above founders on the fact that the picture of induction which it presupposes—a picture presupposed by a good part of the broadly empiricist tradition—is an impoverished one. On a standard scientific realist picture, the scientist's inductively-arrived at knowledge of electrons and quarks is a matter of *inference to the best explanation*: such entities are posited by the theory which best explains the evidence which the scientist observes. Justification for believing that such entities exist and possess such-and-such properties is thus of a piece with the justification that one possesses for believing the relevant theory. Unsurprisingly given the parallelism, inference to the best explanation has been embraced enthusiastically by many realists who accept the phenomenal conception of evidence.¹² On the envisaged picture, what we ultimately have to go on in arriving at a view of the external world are our non-factive experiential states; our knowledge of tables and chairs is underwritten by the fact that the existence of such entities is part of the best explanation of why we have the experiences that we do, or why our experiences tend to have some feature (e.g., a relatively high degree of orderliness or coherence).

For a realist who accepts the phenomenal conception of evidence, perhaps this is indeed the most promising route to pursue.¹³ However, it would be unwise to assume

¹² Prominent contemporary examples would include Laurence Bonjour (2002, chapter 7; Bonjour and Sosa 2003), Richard Feldman (2003, ch.6) and Jonathan Vogel (1990, 2005). For a brief survey of different ways in which the thought might be developed, see Kelly (2006b). The general idea is arguably present in Locke himself.

¹³ Alternatively, such a realist might simply insist that it is a brute fact that having the experience as of p (where p is a proposition about the external world) justifies one in believing p in the absence of defeating evidence, a fact which is not in any way

that inference to the best explanation amounts to a panacea for the skeptical difficulties which she faces. Indeed, even those of us who count ourselves as friends of inference to the best explanation, and think that its employment makes a significant contribution to our knowledge, have good reason to be suspicious of the idea that it can bear all of the weight that it is being asked to bear in this context. When a scientist adopts one hypothesis over another on the grounds of its, e.g., superior simplicity, there are typically indefinitely many other potential explanations of the relevant phenomenon which are never given serious consideration simply because they are inconsistent with, or vastly improbable on, the scientist's background knowledge of how the world works.¹⁴ Of course, this background knowledge is not primarily knowledge of the scientist's own experiences or mental states; it is, rather, knowledge of the external world itself. In short, it is one thing to claim that explanatory considerations can legitimately guide inductive inference to the theoretical entities of science given a rich evidential base consisting of knowledge of the external world; it is quite another to claim that explanatory considerations suffice to ground the entirety of one's knowledge of the external world, given only a relatively meager evidential base made up of one's own non-factive mental states. Could our knowledge of the external world really be a matter of reasoning to the best explanation *all the way down*?

Perhaps. In any case, no decisive reason will be given here for thinking that the project of reconstructing all of our knowledge of the external world as the deliverance of explanatory inferences from a mentalistic evidential base must inevitably end in failure. But the burden of proof would seem to belong to those who think that such a reconstruction is to be had, and the burden will not be an easy one to discharge. In short, the claim that inference to the best explanation can deliver a realist who embraces the

underwritten by explanatory considerations. Other options are also possible; I will not attempt a survey here.

¹⁴ This has suggested to some that reasoning to the best explanation might be illuminatingly reconstructed as a two stage process: in the first stage, all but a small number of potential explanations are 'filtered out' in virtue of their lack of fit with background theory; in the second stage, the remaining competitor which best exemplifies the criteria of explanatory goodness gets inferred as the true explanation of the evidence. For this suggestion, see Lipton (1991).

phenomenal conception of evidence from skepticism is best regarded as a longstanding promissory note, and one which remains quite far from being redeemed, despite the expenditure of considerable energy on its behalf to date.

Of course, even if it turned out that the phenomenal conception *does* lead to skepticism, acceptance of that conception might yet prove unavoidable if it were forced upon us by considerations drawn from the theory of perception. Above, we noted that it is quite natural for a theorist to think of evidence, at least in paradigmatic cases, as consisting of the kinds of things which he takes us to directly observe in standard cases of perception. Again, both the representative realist tradition which descends from Locke as well as the rival phenomenalist tradition which descends from Berkeley hold that that which is immediately given in perception are mental entities. As long as such a picture is in place, the phenomenal conception of evidence would seem difficult to avoid. Of course, phenomenism is currently moribund: never particularly appetizing for the sober-minded in the first place, the phenomenalist program collapsed completely in the second half of the twentieth century, discredited by the repeated failure to actually carry out any successful phenomenalist analyses, along with the provision of principled reasons by philosophers such as Chisholm and Sellars for thinking that such failures were inevitable. Thus, whatever its challenges, the representative realist alternative would seem to be the more promising of the two. Again, given the extreme ambitiousness of the representative realist project of showing that all of our knowledge of the external world ultimately rests on an evidential basis of mental states, one might harbor reasonable doubts about whether such a project can be successfully executed. Still, even in the event that such doubts prove well-founded, if we are committed to the view that the deliverances of perception are mental, it seems that the correct response would be to retain the phenomenal conception and simply conclude that the extent of our knowledge is significantly more limited than is ordinarily supposed.

Given this dialectical situation, an event of crucial importance for the theory of evidence is the relatively recent resurgence of direct realism within the theory of perception. Although a number of different views travel under the banner of 'direct realism', the idea which unites them is this: in standard cases of perception, one directly perceives that things are thus and so in the external world. That is, in paradigmatic cases,

the primary objects of perception are not mental entities, but rather external objects such as tables and chairs and the states of affairs in which they participate. When one knows that the cup is on the table on the basis of visual perception, there is no sense in which one's knowledge is inferentially based on something which is more epistemically immediate, e.g., one's experience as of the cup's being on the table, or the fact that it looks or appears to one as though the cup is on the table. Of course, such a view is hardly an innovation of recent philosophy. Indeed, it has a strong claim to being the pre-philosophical, 'common sense' view of perception. But for much of the twentieth century, direct realism was simply not considered a live option. Rather, it was the one view about perception that could be safely dismissed as utterly untenable by anyone acquainted with the relevant scientific and philosophical considerations.¹⁵ Against this background, the strong resurgence of direct realism in the later part of the twentieth century is an event as unexpected as it is potentially significant. Here is not the place to review the details of that resurgence¹⁶; suffice it to say that in the hands of its most recent advocates, direct realism has been elaborated with a previously unparalleled sophistication, and that it is once again generally regarded as very much a live option, as viable and defensible as any alternative account of perception.¹⁷ If in fact it is defensible to hold that we can have immediate, non-inferential knowledge that things are arranged thus and so in the external world, then this would seem to undermine much of the traditional motivation for denying that such facts are among one's fundamental evidence, or what one has to go on. Given the link between evidence and perception that we have stressed here, it is unsurprising that the perceived viability of a common sense view of

¹⁵ The low regard in which the view was held during this period is perhaps best reflected in the derisive moniker by which it was known, 'naïve realism'.

¹⁶ In retrospect, a seminal text in the resurgence was Austin (1962). Also of significance here was the work of Armstrong (1961, 1968), Dretske (1969), and Pitcher (1971), as well as the rise of 'disjunctivist' conceptions of experience often associated with Hinton (1967), Snowdon (1981) and McDowell (1982).

¹⁷ For something like the state of the art in contemporary debates concerning the nature of perception, see the essays collected in Gendler and Hawthorne (2006).

evidence waxes and wanes along with the perceived viability of a common sense view of perception.

We have briefly canvassed one line of thought that is sometimes offered against the phenomenal conception of evidence (viz. that it makes skepticism unavoidable) and another that is sometimes offered in its favor (viz. that it is the view of evidence which fits best with a philosophically sophisticated picture of perception). While I expressed considerable sympathy for the first line of thought, I do not think that either is decisive. If that assessment is accurate, where would that leave us? On what I take to be the most natural understanding of the dialectic, the burden of proof belongs to the proponent of the phenomenal conception. It is incumbent upon him to provide some compelling reason for thinking that our evidence is restricted in the way that he supposes, for surely the default view is that it is not so restricted. In ordinary, everyday contexts, we often seem to think and speak about evidence in ways that are inconsistent with the phenomenal conception. Asked what evidence I have for my belief that a particular coin is biased, I might make reference to its having behaved in the manner characteristic of a biased coin in the past. If it turns out that my beliefs about its past behavior are false, it is natural to say that my evidence for the claim that the coin is biased was not as strong as I took it to be, or that I did not have the evidence that I thought I had. Similarly, the phenomenal conception of evidence would seem to fit poorly with the way in which the concept is employed in scientific and legal contexts. Notably, in such contexts, a large value is placed on the *publicity* of evidence, i.e., on the fact that paradigmatic evidence is something that can be shared by multiple individuals. Indeed, it is this public character of evidence which is often taken to underwrite the possibility of an inquiry that is genuinely objective.¹⁸ If evidence is taken to include (e.g.) facts about the external world, then, inasmuch as multiple individuals can be aware of the same facts, then one and the same piece of evidence can literally be shared by those individuals. In contrast, a view according to which my evidence is limited to my own non-factive mental states does not seem to allow for this, inasmuch as I do not literally share my own token mental states with anyone else. At best, the sense in which evidence can be shared by multiple

¹⁸ For more on this aspect of evidence, see Kelly (2006a: Section 4) and the references provided there.

individuals on such a picture is a matter of their being in distinct but similar mental states (perhaps: token experiential states of the same type). It is far from obvious that such a surrogate would provide the kind of objectivity which literally shared public evidence is often taken to afford.

None of this is to say that the phenomenal conception must be mistaken; it is to say that we should not accept it in the absence of very strong reasons for doing so. But if acceptance of the phenomenal conception of evidence is not dictated by considerations drawn from the philosophy of perception, what reasons remain for thinking that it is true? Here is one line of thought that might lead one to accept it. In section 1, we noted that in a given case one might be led astray by following one's evidence: even if *p* is true, one's evidence might misleadingly suggest that *p* is not true. When one's evidence is misleading, one typically arrives at a false belief by believing in accordance with it. We ordinarily assume that such cases are exceptional. Are there possible worlds in which such cases are the norm? Consider a careful and judicious thinker who consistently and scrupulously attends to his evidence in arriving at his beliefs. In our world, these habits lead to cognitive prosperity—the individual holds a relatively large number of true beliefs and relatively few false beliefs. (Or at least, he fares significantly better with respect to truth and falsity than those who fail to attend to their evidence and instead form their beliefs in a hasty or haphazard manner.) Consider next how the same individual fares in a world that is subject to the machinations of a Cartesian evil demon, a being bent on deceiving the world's inhabitants as to its true character. Although the true character of the world in question differs radically from our own, it is, from the point of view of its inhabitants, utterly indistinguishable, for the Demon takes care to ensure that the courses of experiences that the inhabitants undergo are qualitatively identical to the courses of experiences that they undergo in our non-delusory world. In the world run by the Cartesian Demon, our thinker is no less judicious and no less scrupulous in attending to (what he blamelessly takes to be) relevant considerations than he is in our world. Because of his unfortunate circumstances, however, his beliefs embody a radically false picture of his environment. Granted that the thinker's beliefs about his environment are false, are they any less justified than in our world? Is the thinker himself any less rational? Many philosophers maintain that the thinker's beliefs are equally well-justified

and that the thinker himself is equally rational in the two worlds (See e.g., Cohen 1984 and Pryor 2001). Apparently, there is strong intuitive resistance to the idea that a thinker whose underlying dispositions and habits of thought remain unchanged might become less rational simply in virtue of being located in less fortuitous circumstances. However, the judgement that the thinker is equally rational in ‘the good case’ and ‘the bad case’ tends to push one inexorably towards a conception of evidence according to which one’s evidence is exhausted by one’s subjective, non-factive mental states. For if rationality is a matter of responding correctly to one’s evidence, then the judgement that the thinker is equally rational in the good case and the bad case would seem to require that the thinker has the same evidence for his beliefs in both cases. But *ex hypothesi*, the only thing common to the good case and the bad case that is a plausible candidate for being the thinker’s evidence are his non-factive mental states. Thus, the judgement that the thinker is equally rational in both cases, when conjoined with the view that rationality is a matter of responding to one’s evidence in the appropriate way, seems to force the conclusion that the thinker’s evidence is limited to his non-factive mental states *even in the good case*. In this way, the requirement that the thinker has the same evidence in the good case and the bad case leads to the phenomenal conception of evidence.¹⁹

Williamson (2000) argues at length that we should not accept the idea that one has the same evidence in the good case and the bad case. Central to his argument is the contention that, even if one were to adopt the phenomenal conception of evidence, this would not allow one to vindicate the underlying intuitions that seemed to make its adoption attractive in the first place; hence, the phenomenal conception of evidence is ultimately not well-motivated. As we have seen, it is the desire to preserve the intuition that a sufficiently scrupulous thinker in the bad case is no less reasonable than a similarly scrupulous thinker in the good case which seems to rule out any conception of evidence

¹⁹ Silins (2005) is perhaps the most extended defense of the view that the thinker has the same evidence in the good case and the bad case. He argues that, if the sameness of evidence requirement is jettisoned, there will be circumstances in which the thinker in the good case ends up being *less* reasonable than the thinker in the bad case. He suggests, plausibly, that this is a potentially more embarrassing consequence than the opposite possibility: ‘Even if one is willing to accept the result that one is sometimes more rational in the good case than in the bad case, it is harder to live with the claim that one is sometimes more rational in the bad case than the good case’ (p.390).

according to which one's evidence might consist of (say) true propositions or facts about the external world. For a thinker in the bad case is not in a position to recognize facts about the external world; he is, however, in a position to recognize facts about his own experiences. The view that one's evidence is limited to one's non-factive mental states thus seems to be motivated by the idea that one's evidence, no matter what else is true of it, must be the kind of thing that one is always in a position to correctly take into account, at least in principle. But (it is claimed) one's experiences are the things that one is always in a position to correctly take into account. Williamson contends that this last thought is a mistake: in fact, one is not always in a position to correctly take into account one's experiences, even in principle. Indeed, Williamson argues that there is *no* non-trivial condition which is such that one is always in a position to know that it obtains. Thus, the thought that evidence might be such that one is always in a position to know what one's evidence is is a chimerical one. To insist that

in order for x to be among one's evidence, x must be such that one is always in a position to know whether one's evidence includes x

is thus to impose a misguided and unrealizable *desideratum* on the theory of evidence. In short: "*Whatever* evidence is, one is not always in a position to know what one has of it" (2000: 178, emphasis added).

Still, even if one were to give up on the strong accessibility requirement which Williamson criticizes, one might nevertheless be left with the brute intuition that a thinker in the bad case is equally well-justified in his beliefs, and hence, has the same evidence as a thinker in the good case. In deciding whether to accept or reject the phenomenal conception of evidence, how much weight should one give to this intuition, relative to considerations which seem to count against the phenomenal conception, e.g., considerations having to do with publicity?

Here we see that questions about the ontology of evidence, when pursued far enough, ultimately raise fundamental questions about philosophical method. Historically, preoccupation with the traditional skeptical problematic has played a significant role in the widespread acceptance of the phenomenal conception, and what popularity it retains is perhaps due in large measure to the conviction that one's beliefs about the external

world would be no less justified even if one were a recently envatted brain or the plaything of a Cartesian evil demon. As is often noted, such skeptical scenarios typically receive little attention from the scientist, the lawyer, or the ordinary person. It is thus perhaps not surprising that the conception of evidence employed in the laboratory, the courtroom or on the street bears little resemblance to the phenomenal conception. The methodological question for the theorist then, is this: in theorizing about evidence, to what extent should we be guided by judgments relating to the traditional skeptical problematic, and to what extent should such theorizing be responsible to central features of the concept as it is employed in non-philosophical contexts?

References

- Achinstein, Peter (ed.) (1983). *The Concept of Evidence* (Oxford: Oxford University Press).
- Armstrong, David (1961). *Perception and the Physical World* (London: Routledge and Kegan Paul).
- Armstrong, David (1968). *A Materialist Theory of the Mind* (New York: Humanities Press).
- Audi, Robert (2001). "An Internalist Theory of Normative Grounds", in *Philosophical Topics* (29):31-45.
- Austin, J.L. (1962). *Sense and Sensibilia* (Oxford: Oxford University Press).
- Ayer, A.J. (1936). *Language, Truth and Logic*. (New York: Dover)
- BonJour, Laurence (1999). "Foundationalism and the External World", in James Tomberlin (ed.) *Philosophical Perspectives* 13: Epistemology (Blackwell).
- BonJour, Laurence (2002). *Epistemology*. (New York: Rowman and Littlefield).
- BonJour, Laurence and Sosa, Ernest (2003). *Epistemic Justification* (Blackwell Publishers).
- Cohen, Stewart (1984). "Justification and Truth", *Philosophical Studies* 46: 279-295.
- Conee, Earl and Feldman, Richard (2004). *Evidentialism: Essays in Epistemology* (Oxford: Oxford University Press).
- Davidson, Donald (1986). "A Coherence Theory of Truth and Knowledge". In E. LePore (ed.), *Truth and Interpretation*. (Oxford: Blackwell).
- Dretske, Fred (1969) *Seeing and Knowing* (London: Routledge and Kegan Paul).
- Earman, John (1993). "Underdetermination, Realism and Reason". *Midwest Studies in Philosophy*, 18: 19-38.
- Feldman, Richard. *Epistemology*. (New Jersey: Prentice Hall).
- Gendler, Tamar Szabo and Hawthorne, John (2006). *Perceptual Experience* (Oxford: Oxford University Press).

- Hempel, Carl (1960). "Inductive Inconsistencies", *Synthese* 12: 439-469. Reprinted in Hempel (1965).
- Hempel, Carl (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science* (New York: The Free Press).
- Hinton, J.M. (1967). "Visual Experiences". *Mind*, New Series, Vol.76, No.302: 217-227.
- Hofer, Carl and Rosenberg, Alexander (1994). "Empirical Equivalence, Underdetermination, and Systems of the World", in *Philosophy of Science* 61(4): 592-607.
- Horwich, Paul (1982). *Probability and Evidence* (Cambridge: Cambridge University Press).
- Howson, Colin and Urbach, Peter (1993). *Scientific Reasoning: The Bayesian Approach*, 2nd edition. (Chicago: Open Court).
- Jeffrey, Richard (1992). *Probability and the Art of Judgment* (Cambridge: Cambridge University Press).
- Jeffrey, Richard (2004). *Subjective Probability* (Cambridge: Cambridge University Press).
- Joyce, James (2005). "How Probabilities Reflect Evidence", in John Hawthorne (ed.) *Philosophical Perspectives, vol.19: Epistemology* (Blackwell: Oxford).
- Kelly, Thomas (2006a). "Evidence" in *The Stanford Encyclopedia of Philosophy* (Fall 2006 Edition), Edward N. Zalta (ed.)
<http://plato.stanford.edu/archives/fall2006/entries/evidence/>.
- Kelly, Thomas (2006b). "The Cost of Skepticism: Who Pays?" *Philosophical Studies*, vol.131, no.3, December: 695-712.
- Keynes, John Maynard (1921). *A Treatise on Probability*. (London: Macmillan).
- Laudan, Larry (1991). "Empirical Equivalence and Underdetermination" in *The Journal of Philosophy* 88(9): 449-472.
- Laudan, Larry (1996). *Beyond Positivism and Relativism: Theory, Method, and Evidence*. (Boulder, CO: Westview Press).
- Lipton, Peter (1991). *Inference to the Best Explanation* (London: Routledge).
- Lipton, Peter (ed.) (1995). *Theory, Evidence and Explanation*. (Dartmouth Publishing Group).

- McDowell, John (1982). "Criteria, Defeasibility, and Knowledge". *Proceedings of the British Academy* 68: 455-479.
- McDowell, John (1995). "Knowledge and the Internal", in *Philosophy and Phenomenological Research* (55): 877-893.
- Millar, Alan (1991). *Reasons and Experience*. (Oxford: Clarendon Press).
- Moser, Paul (1989). *Knowledge and Evidence*. (Cambridge: Cambridge University Press).
- Neta, Ram (forthcoming). "What Evidence Do You Have?" *British Journal for Philosophy of Science*.
- Okasha, Samir (2002). "Underdetermination, Holism and the Theory/Data Distinction" in *Philosophical Quarterly* 52 (208): 302-319.
- Pitcher, George (1971). *A Theory of Perception* (Princeton, NJ: Princeton University Press).
- Plantinga, Alvin (1993). *Warrant and Proper Function* (Oxford: Oxford University Press).
- Pryor, James (2001). "Highlights of Recent Epistemology". *British Journal for the Philosophy of Science* 52: 95-124.
- Quine, W.V. (1969) "Epistemology Naturalized". In his *Ontological Relativity and Other Essays*. (Columbia University Press: New York, NY).
- Quine, W.V. (1980) *From a Logical Point of View*. (Cambridge, MA: Harvard University).
- Silins, Nico (2005). "Deception and Evidence", in John Hawthorne (ed.) *Philosophical Perspectives, vol.19: Epistemology* (Blackwell: Oxford).
- Skyrms, Brian (1980). *Causal Necessity* (London, Yale University Press).
- Snowdon, P. (1981). Perception, Vision, and Causation. *Proceedings of the Aristotelian Society* 81: 175-192.
- Unger, Peter (1975). *Ignorance: A Case for Skepticism*. (Oxford: Oxford University Press).
- Vogel, Jonathan (1990). "Cartesian Skepticism and Inference to the Best Explanation", *The Journal of Philosophy* 87:658-666.

Vogel, Jonathan (2005). "The Refutation of Skepticism", in Mathias Steup and Ernest Sosa (eds.) *Contemporary Debates in Epistemology* (Blackwell Publishers).

Williamson, Timothy (2000). *Knowledge and Its Limits* (Oxford: Oxford University Press).