We can see that (A) also quickly eliminates our other contrary possibility about the relation of (C) to justifying evidence that Smith has at t2, the possibility that (C) is itself evidence for Smith. At t2 a reasonable objection has been made against (R) to Smith. If (C) were evidence for Smith at that point, then he could simply rely on (C) without further support in order to refute that objection. But (C) is not self-evident. It declares certain evidence misleading. Why should Smith accept it? (C) is unreasonable for Smith to accept in the absence of support, and (R) is Smith's only potential source of such support. As (A) asserts, and manifestly, once (R) has been cast into doubt, wielding (C) against the objection to (R) would not refute the objection. It is better to deny that (C) is evidence and return to the conception of evidence as data, thereby gaining evidentialist grounds for disallowing this violation of (A). Thus, using (A) we can defend the view that (C) is not available to Smith at t2 to argue against the force of (2).[16]

Conclusion

We can be misled out of knowing something. This possibility is not one of the glories of human existence. But it is a fact. It is also a fact that we cannot insulate ourselves from being thus misled by inferring from a known truth that evidence against it is evidence against a truth. That inference, though impeccable, turns out not to do us any good when we need it. This lack of insulation is not a joyous fact, either. But it is another way that things are.

An evidentially based account explains how all of this can be. We are in no position to disregard evidence when it goes strongly against something that we knew, because we have nothing beyond our evidence as our fundamental epistemic basis for disregarding as well as for knowing. When serious doubts arise, there is nothing better that we can do for the sake of knowing than to keep investigating. Fortunately, further inquiry has frequently overcome deception and rarely made things worse. Induction from this fact is encouraging.[17]

[16] This explanation relies on an asymmetric epistemic dependency relation among reasons. To account for the justification to Smith, at t2 of (2), rather than (3), the explanation uses the fact that (3) depends on (R) in a way that (R) does not depend on (3). A coherentist who denies the existence of such asymmetric support relations (or denies their justificatory relevance) seems unable to account for the fact that it is the members of the combination—(1), (2), not-(3), and not-(R)—that are justified for Smith at t2. Ignoring this dependency appears to leave no good reason to deny that Smith would be equally reasonable in believing at t2 the propositions—(R), (1), (3) and not-(2). Since in fact this would be an unreasonable refusal to defer to compelling counter-evidence, an inability to account for this constitutes an objection to any such form of coherentism.

[17] I am grateful to James Cargile, Stewart Cohen, Richard Feldman, epistemology discussion group participants at Brown University, and colloquium participants at UC Davis for helpful comments on previous drafts of this work.

12

Making Sense of Skepticism

Richard Feldman and Earl Conee

This paper compares the way evidentialist and non-evidentialist theories deal with skepticism. We first formulate several familiar arguments for skepticism and then argue that non-evidentialist theories of knowledge and justification fail to engage the considerations that underlie these arguments. In contrast, an evidentialist approach to the justification that is required for knowledge best facilitates an appreciation of skeptical arguments. In the final section of the paper we propose an evidentialist theory of knowledge-level justification that enables us to explain both the credibility and the failure of the skeptical arguments.

I. Arguments for Skepticism[1]

The following four arguments are designed to represent the main classic defenses of external world skepticism.[2]

The Possibility of Error Argument (PE)

1. Any belief about the external world could be mistaken.
2. If a belief could be mistaken, then it is not knowledge.

---

3. Therefore, no belief about the external world is knowledge. (1), (2)

[1] Some parts of this section, the next one, and the final one are heavily revised versions of material originally appearing in chapters 6 and 7 of Richard Feldman, *Epistemology* (Englewood Cliffs, NJ: Prentice-Hall, 2003).

[2] We include as "beliefs about the external world" only contingent propositions implying the existence of an external world. Perhaps a belief such as what one expresses by "I exist" or the belief that either it is the case that there is water or it is not the case that there is water cannot be mistaken. But they do not count for present purposes as beliefs about the external world

(PE) summarizes familiar arguments based on skeptical scenarios. These arguments support premise (1) in the following way. No one has what Descartes termed "certain indications" to distinguish dreams from reality. Anyone could be a victim of an evil demon, a brain in a vat, or the like. So any external world belief could be mistaken. Premise (2) reflects something about the standards for knowledge. We shall return to it later.

(PE) also encapsulates another traditional skeptical argument, the Argument from Illusion. That argument begins with the thought that illusions sometimes fool people about the external world. The argument continues with the claim that illusion is always possible. The conclusion inferred is that no one knows anything about the external world. Like (PE), this reasoning focuses on a possibility of error. In the Argument from Illusion, the error would stem from deceptive perceptual experiences. We will accommodate this feature of the Argument from Illusion when we discuss (PE) below, concentrating on errors that are possible while we have the same basis for belief.

A second argument for skepticism turns on the introspective indistinguishability of any case of true external world belief from some possible cases of false belief. Fallibilist views about knowledge and justification imply that one can have knowledge on the basis of justification that does not entail the known proposition. Presumably, the fundamental basis for external world justification is perceptual experience. There are pairs of possible cases in which someone believes the same external world proposition and things seem the same to the person, yet in one case the proposition is true and in the other it is false. This implies that if one can have external world knowledge on the basis of perceptual experience, then a case in which one has the knowledge can be introspectively the same as a case in which one lacks the knowledge. Our second skeptical argument contends that the consequent of this conditional is false, and so there cannot be such knowledge.

The Introspective Indistinguishability Argument (II)

1. If anyone has knowledge of the external world on the basis of experience, then there can be cases of knowledge that are introspectively indistinguishable from cases of false belief.
2. There cannot be cases of knowledge that are introspectively indistinguishable from cases of false belief.
3. No one has knowledge of the external world on the basis of experience. (1), (2)

4. If anyone has knowledge of the external world, it is on the basis of experience.

---

5. No one has knowledge of the external world. (3), (4)

(II) also summarizes a traditional skeptical argument, the Argument from Appearance. That argument proceeds from the premise that people have direct awareness of how external things appear to be, and not of how they really are. The argument adds that appearances might be misleading. The inference drawn is that no one has knowledge of the external world. This traditional concern about a possibility of misleading appearances is captured by the possibility that (II) describes of false beliefs with no introspectable difference from true beliefs.

The Transmissibility Argument is a third argument for skepticism. The gist of the argument is this: ordinary beliefs about the external world imply that the skeptical scenarios are false. If one had knowledge in the ordinary cases, and realized that this implication holds, then one could deduce, and thus know, that the skeptical scenarios are false. But, according to the argument, no one knows that the skeptical scenarios are false. So it must be that no one knows the ordinary propositions from which the falsity of the skeptical scenarios could be inferred.

To formulate this argument efficiently, we will let (O) be any ordinary external world proposition that we would typically claim to know and let (SK) be any skeptical hypothesis inconsistent with (O). (O) will then imply that (SK) is false. Let S be any ordinary person who knows that (O) implies that (SK) is false.

The Transmissibility Argument (TK)

1. S cannot know that (SK) is false.
2. (O) implies that (SK) is false, and S knows this.
3. If S knows that (O) is true, and that (O) implies that (SK) is false, then S can know that (SK) is false.

---

4. S does not know (O). (1)–(3)

One other skeptical argument deserves separate discussion. It relies on the idea that experiential evidence gives no better reason to accept a commonsense view of the world than to accept alternative hypotheses such as that one is dreaming (DR), that one is a brain in a vat (BIV), or that one is deceived by an evil demon (ED). The lack of a better reason to believe a commonsense view is claimed to imply a lack of justification, and hence a lack of knowledge.

The Alternative Hypotheses Argument (AH)

1. The experiences people have provide no better reason to believe ordinary external world propositions than rival skeptical hypotheses, such as (DR), (BIV), and (ED).
2. If experiences do not provide better reason to believe one external world hypothesis than to believe another, then people are not justified in believing the one.

---

3. People are not justified (and thus do not know) ordinary external world propositions. (1)–(2)

II. Non-Evidentialist Theories

In this section we will briefly describe five non-evidentialist theories of knowledge and justification. Since our goal here is to examine the kinds of resources that these theories make available for dealing with skepticism, we will not present detailed versions.

One of the first non-evidentialist theories to gain widespread recent attention is the causal theory. The basic idea is that one has knowledge when one's belief is causally connected to the facts in the right way. An initial formulation is:

C. S knows p iff S's belief in p is caused by the fact p.

There are cases involving deviant causal chains and bizarre causal links between facts and beliefs in those facts. Such cases prompted a leading defender of the causal theory, Alvin Goldman, to replace (C) with:

C*. S knows p if and only if the fact p is causally connected in an appropriate way with S's believing p.[3]

For present purposes it is not necessary to discuss in detail what counts as an appropriate connection. It will suffice to note that ordinary perception, memory, introspection, and reasoning are among the leading candidates.

The second non-evidentialist theory is the tracking theory. It holds that one has knowledge of a proposition provided one's belief in that proposition tracks the truth of that proposition:

TT. S knows p iff (1) p is true, (2) S believes p, and (3) S's attitude toward p tracks the truth-value of p: if p were not true, then S would not believe p; and if p were true, then S would believe p.

[3] "A Causal Theory of Knowing," *Journal of Philosophy*, 64 (1967), 357–72; see p. 369.

Some counterexamples to (TT) involve people who would use a crucially different method to determine what to believe about P, if P were not true. Such examples prompted a leading defender of this theory, Robert Nozick, to propose the following modified version:

TT*. S knows p iff (1) p is true, (2) S believes p, (3) there is a belief forming method M that S used to form the belief in p and (4) when S uses method M to form beliefs about p, S's beliefs about p track the truth of p.[4]

A third non-evidentialist theory is reliabilism. Reliabilism has been developed both as a theory of justification and as a theory of knowledge that omits any justification condition. These differences will not affect the points to be discussed below, and it is slightly simpler to proceed in terms of reliabilist theories of knowledge. In its simplest form, reliabilism is the view that:

R. S knows p iff S's true belief in p results from a reliable belief-forming process.[5]

Reliable belief-forming processes are processes that generally lead to true beliefs.[6] Details about how reliability is measured will not affect the discussion that follows.

A fourth non-evidentialist theory is the proper function theory, championed by Alvin Plantinga.[7] The central idea is:

PF. S knows p iff S's true belief in p results from the proper functioning of S's cognitive system.

A variety of problem cases have led Plantinga to propose that knowledge requires warranted belief, where warrant is something closer to the following:

PF*. A belief is warranted iff (1) it results from the proper functioning of the believer's cognitive system in a suitable environment, (2) the segment of the system that produced the belief is aimed at the truth, and (3) the overall system usually produces true beliefs when it is in a suitable environment.[8]

[4] *Philosophical Explanations* (Cambridge, Mass.: Harvard University Press, 1981), 179.

[5] There are a number of bells and whistles that can be introduced into reliabilism. One can add what amounts to a "no-defeaters" clause, holding that one has knowledge only if there is no reliable process that would lead S not to believe p. One could also introduce the distinction between reliable and conditionally reliable belief-forming processes. (See Alvin Goldman, "What is Justified Belief?", in George S. Pappas (ed.), *Justification and Knowledge* (Dordrecht: D. Reidel, 1979), 1–24.) These modifications will not matter in what follows, and we will ignore them here.

[6] See "The Generality Problem for Reliabilism" (Ch. 6) for problems about the bearers of reliability.

[7] See *Warrant and Proper Function* (Oxford: Oxford University Press, 1993).

[8] Ibid. 46–7.

Finally, there is the idea that it is the "safety" of a true belief that makes it knowledge. Ernest Sosa is the leading proponent of this sort of theory.[9] Very roughly, a true belief is safe provided that it is unlikely to be falsely held. Consider, for example, a typical person's belief that he sees a tree. It is quite plausible that this belief is safe, since ordinarily, it would take the sight of a tree for an ordinary person to believe that he sees a tree. Since beliefs in metaphysical and logical truths, as well as laws of nature, may satisfy this condition without being cases of knowledge, Sosa adds an additional condition for knowledge. In one place he suggests:

S. S knows p iff (1) p is true, (2) S's belief in p is safe, and (3) S's belief in p results from a virtuous intellectual method.[10]

III. Contextualism

Non-evidentialist theories of knowledge are often combined with contextualism in an effort to respond to skepticism. Contemporary contextualism in epistemology is primarily a metalinguistic thesis about "knowledge" and its cognates. The contextualist view is that the truth conditions for attributions of "knowledge" vary with the context of the attributer.[11] The capacity of contextualism to deal with skeptical arguments is often cited as one of its principal assets. Typically, contextualists hold that something about contexts in which skeptical arguments are prominent makes true the skeptic's denials of external world "knowledge" attributions, while something about ordinary contexts makes true many ordinary external world "knowledge" attributions.[12]

In what follows we will focus on non-contextualist versions of the non-evidentialist theories discussed in section II. We will argue that these theories do not make sense of skepticism, since they do not make clear why anyone would even be tempted by the arguments for skepticism. Yet people routinely do

[9] See "Tracking, Competence, and Knowledge", in Paul K. Moser (ed.), *The Oxford Handbook of Epistemology* (Oxford: Oxford University Press, 2002), 264–86.

[10] Ibid. 275.

[11] Three influential sources are Stewart Cohen, "Contextualism, Skepticism, and the Structure of Reasons", *Philosophical Perspectives*, 13 (1999), 57–89; Keith DeRose, "Solving the Skeptical Problem", *Philosophical Review*, 104 (1995), 1–52; and David Lewis, "Elusive Knowledge", *Australasian Journal of Philosophy*, 74 (1996), 549–67. Cohen defends an evidentialist view, but DeRose and Lewis do not. More recently, Ram Neta has defended a contextualist view about what evidence counts toward justification, in "Contextualism and the Problem of the External World", *Philosophy and Phenomenological Research*, 66/1 (2003), 1–31.

[12] Contextualists can defend an evidentialist theory of knowledge. Those who do can largely agree with the explanation we offer in section VI of why we often satisfy the conditions for knowledge in place in ordinary contexts.

find skeptical arguments tempting. Contextualist versions of these theories offer an explanation. Again, the idea is that in contexts where skepticism is prominent, the truth conditions for "knowledge" attributions are so demanding that ordinary external world beliefs do not meet them. This goes too far. As one of us has argued elsewhere,[13] this is an unwarranted concession to skepticism. The reasons for thinking that skeptics are right when someone is paying attention either to skepticism or to an argument for it are no better than the reasons for thinking that skeptics are right when no one is paying attention to any such thing. Furthermore, external world skepticism can be refuted. As with all enduring disputed positions in philosophy, the skeptical side has been defended with some plausibility. As with the other enduring disputes, there is no need to endorse the conclusions of both sides to make sense of the debate. When all of the arguments for one side can be given their due while denying that any of them is sound, it is most reasonable not to accept that side's conclusion. An evidentialist theory of knowledge can do this for skepticism. That done, there is no good reason also to interpret the skeptic's conclusion so that it says something true, as contextualists do.

Invariantists deny that "knowledge" attributions have contextually varying truth conditions. Our goal here is to argue that invariantist evidentialist theories of knowledge are much better equipped than are the recent non-evidential theories to appreciate the challenge of external world skepticism, and to reply to that challenge. Invariantist evidentialist theories are not only better equipped, they are so well equipped that no explanatory work about skepticism remains for contextualism to accomplish. Invariantist evidentialism provides a fully reasonable assessment of skepticism.[14] By defending all of this, we are not establishing that contextualism is incorrect. We are showing something about what can be done without it.

IV. Non-Evidentialist Theories and Skepticism

Each of the non-evidentialist theories described in section II provides the basis for simple responses to the skeptical arguments.[15] We will briefly describe each of these responses, and then consider some of the issues that arise in assessing

[13] Richard Feldman, "Skeptical Problems, Contextualist Solutions", *Philosophical Studies*, 103 (2001), 61–85.

[14] This is not to say that our version of invariantist evidentialism has all the answers. We leave open important issues. But they are not ones where contextualism holds promise.

[15] Defenders of these non-evidentialist theories who also endorse contextualism will say that the responses described in the remainder of this section apply only to attributions of "knowledge" in ordinary contexts. Defenders of these theories who do not endorse contextualism would say that the responses apply to all contexts.

them. We will then turn to a general assessment of the merits of these responses to the arguments. It is worth noting at the outset that evidentialists are likely to agree with non-evidentialists about which premises of these arguments are mistaken. The disagreements concern the reasons for their falsity.

Consider first (PE). All of us will concede that, in some relevant sense, each external world belief could be mistaken. Non-skeptics will, however, contest the other premise:

2. If a belief could be mistaken, then it is not knowledge.

The phrase "could be mistaken" is susceptible to a number of interpretations. According to one, a belief could be mistaken provided it is possible for the belief to be falsely held. Virtually all beliefs in contingent propositions could be mistaken in this sense. This possibility seems just irrelevant to whether the proposition is ever believed under conditions that yield knowledge. It is difficult to see why anyone would accept (2) when it is interpreted in this way.

According to a better interpretation of (2), its antecedent says that it is possible for the belief to be false while holding fixed epistemically relevant aspects of the circumstances. The central idea is that it is possible for the belief to be false while it is believed on its actual basis. Thus, for example, one's belief about what color one is seeing "could be mistaken," if it is possible for that belief to be false when it is held on the basis of the visual experiences, background information about color classifications, and whatever else one uses to determine the color. We will make use of this interpretation of (2) in the remainder of this discussion.

Causal theorists can say that even though a belief on the same basis could have been mistaken, it may actually have the right sort of causal connection to the facts. Tracking theorists can say that it is consistent with this sort of possibility of error that people actually track the truth of the proposition. Reliabilists will be quick to point out that beliefs formed on any given internal basis could be false while caused by reliable processes. Proper function theorists will respond that a belief that results from a properly functioning cognitive system could have been false even if held on the same basis. A safety theorist will observe that a safe method of adopting a true belief need preserve truth only among the sufficiently likely possibilities. This allows for the possibility of false belief using the same virtuous method. Thus, defenders of all these theories will dismiss (2) as obviously false. Knowledge as they characterize it is plainly compatible with the possibility of error.

The initial appeal of the skeptical argument (PE) remains to be understood, however. In particular, there is significant credibility in some thought to the effect that we do not know anything about which we could be mistaken. On its

current interpretation, the thought is that one does not know a proposition when it is possible to believe the same proposition falsely on the same basis. This is credible because if one believed the same proposition on the same basis, then things would seem just the same. This duplication of how things seem, accompanied by false belief, appears somehow to undermine the capacity of the basis to give knowledge.

The invariantist versions of the non-evidentialist theories of knowledge neither accommodate this appearance nor explain it away. None of those theories implies a condition for knowledge that is even weakened by the possibility of error on the same basis.[16] None of the theories provides a way to find a plausible error in the impression that this possibility excludes knowledge.[17]

A correct theory of knowledge that is incompatible with premise (2) of (PE) need not readily explain its plausibility. A correct theory need only give some informative account of what knowledge is. However, something about knowledge makes external world skepticism a perennial temptation. It is reasonable to expect that an illuminating account of the conditions that constitute knowledge will shed light on sources of that temptation.

This is a strength that evidentialist theories have. They make the bearing on knowledge of the skeptic's possibility easy to see. Any known proposition must be adequately supported by the person's evidence. It is at least initially credible that the possibility of being a basis for a false belief is a flaw in the evidence. It is clearly reasonable to suspect that justification constituted by flawed evidence is not adequate for knowledge.

It can be acknowledged by evidentialist theorists that this fallibility is an imperfection in the evidence. So the skeptic need not be wrong that the

[16] A broadly reliabilist explanation of the credibility of the premise might be attempted. It would have to use a notion of reliability very different from the one employed by reliability theorists. Perhaps the possibility of error shows that the process used is not necessarily completely reliable. This affects how reliable the process is only if accuracy in remote non-actual cases affects the measure of reliability. It is possible for a reliabilist to claim that the plausibility of (PE) derives from assuming wrongly that all possibilities for a process affect its reliability. It is contrary to the spirit of reliabilist theories to allow such factors to affect measures of reliability. Reliabilists emphasize contingent reliability as what is crucial. So anyone who would count such remote factors toward reliability is using something that does not resemble a standard reliabilist condition on knowledge. (We thank John Bennett for his comments on this topic.)

[17] As noted earlier, some contextualists contend that attributions of external world "knowledge" are almost always false in contexts in which skeptical arguments are discussed. Thus, they do not reject (2). For example, Lewis, in "Elusive Knowledge", held that knowledge is not compatible with the possibility of error. However, he claimed that what possibilities there are varies with context. Discussion of skepticism expands the set of possibilities. On his view, then, there is no error, hence no plausible error in the view that knowledge is incompatible with the possibility of error. Relative to ordinary contexts in his view, (1) of (PE) is mistaken.

justification has a kind of weakness. The task of opposing (PE) by using an evidential theory is to explain how the evidence manages nonetheless to be adequate for knowledge. In contrast, defenders of invariantist non-evidential theories get no help from their theories in constructing an account of why the possibility of error might even be thought to bear on whether one has knowledge.

Consider next (II). Here non-evidentialists usually object to its second premise:

2. There cannot be cases of knowledge that are introspectively indistinguishable from cases of false belief.

Their reasons for rejecting (2) are almost identical to their reasons for rejecting the disputed premise of (PE). There need not be any introspectable difference between cases in which the right causal connection is present and cases in which it is not. Similarly, cases in which someone tracks the truth, believes as a result of a reliable process, functions properly, or has safe beliefs, can be indistinguishable from cases in which the person does not do these things.

Again, the initial force of the argument is not accounted for. Why would it be thought for a moment that knowledge is excluded by the fact that introspection finds no difference, if knowledge turns on proper causal connection, or truth tracking, or reliably caused belief, or a properly functioning belief-producing mechanism, or safety? The findings of introspection under very different possible circumstances where the belief is false would have no discernible relevance. The skeptic might as well have proclaimed that someone knows that something is true only if the person is the sole cause of its truth. This idea has nothing going for it from the start. We would not expect a satisfactory theory of knowledge to accommodate this idea or to show the reasonable mistake in it. By contrast, the absence of an introspectable difference between the actual case and a possible case of false belief does seem at least threatening to knowledge. Non-evidentialist theories do not help us to understand why.

If what introspection reveals is one's fundamental evidence, or at least a crucial portion of it for external world propositions, then it is no wonder that it matters to knowledge which possibilities introspection allows. If evidence determines justification, then it is no wonder that allowing the falsehood of the belief renders suspect the strength of the justification. Even if non-entailing evidence is enough for knowledge, the skeptic seems to be on to something.

Consider next (AH). All non-evidentialist theories will happily reject premise (2):[18]

2. If experiences do not provide better reason to believe one external world hypothesis than to believe another, then people are not justified in believing the one.

There can be causal connections, truth tracking, reliable belief formation, proper function, and safety, in the absence of better reasons for a hypothesis. Details may matter. We can imagine defenders of the skeptical argument contending that proper function requires evidential support of the sort described in the argument. However, there is plausibility to the claim that people are designed to believe certain propositions even if experience does not provide reasons to believe those propositions. In any case, it is clear that there is no necessary connection between this sort of support by reasons and proper function. So if the proper function theory is correct, then this premise could easily have been false.

Once again, these theories leave us in the dark about the rational attraction of the argument. The comparative strength of one's reasons for accepting ordinary propositions about the world is given no relevance. Yet it seems clearly requisite to knowing about the external world that one's beliefs about it are supported by adequately good reasons. At a minimum, if some extraordinary skeptical possibility is equally reasonable in the presence of ordinary experiences, then people seem to have at best weak reason to accept ordinary thinking on the topic. Having at best weak reason seems to exclude knowing.

Christopher Hill has argued that reliabilists can make good sense of the appeal of an argument like (AH).[19] Our premise (2) spells out the idea that the justification of our external world beliefs requires that one's experiences give better reasons to accept those beliefs than to accept the skeptic's alternatives. Hill's formulation includes a similar premise:

2*. In determining whether a person is justified in believing ordinary external world propositions rather than the skeptical alternatives, it is appropriate to set all non-sensory evidence aside, and to focus exclusively on facts involving the person's sense experiences and their purely sensory characteristics.[20]

[18] Or they will happily deny that knowledge requires justification. Either way, the same explanatory weaknesses arise.

[19] In "Process Reliabilism and Cartesian Skepticism", *Philosophy and Phenomenological Research* 56 (1996), 567–81; reprinted in Keith DeRose and Ted A. Warfield (eds.), *Skepticism: A Contemporary Reader* (Oxford: Oxford University Press, 1999), 115–28.

[20] We have reworded Hill's premise slightly. As Hill understands (2*), considerations of explanatory power or other "superempirical virtues" are set aside. Only the purely sensory characteristics of experiences have any bearing on justification. Hill acknowledges that some philosophers may think

Hill's central contention is that the appeal of (2*) results from conflating the plausible claim

A. If one wishes to determine whether a person is justified in preferring the real world hypothesis to the skeptical alternatives, it is permissible to set all of the person's extramental evidence aside, and focus exclusively on the person's sensory evidence

with the less plausible claim

B. If one wishes to determine whether a person is justified in preferring the real world hypothesis to the skeptical alternatives, it is appropriate to set all of the person's non-sensory evidence aside, and focus exclusively on facts involving the person's sensory experiences and their purely sensory characteristics.

Both (A) and (B) restrict evidence to sensory experiences. (B) further restricts evidential considerations to facts involving sensory characteristics. (A), by contrast, allows consideration of all facts of any sort about sensory evidence, including the reliability of its connection to the truth of external world propositions. One reason why the skeptical argument is supposed to seem compelling, then, is that philosophers slip easily from (A) to (B).

Hill says that "to do full justice to the seductiveness"[21] of the skeptical argument, a deeper explanation is needed. We agree. Conflating (A) and (B) would explain at most the appeal of just those skeptical arguments that turn on which features of experience provide external world evidence. We do not see that any such conflation is needed to find (2*) plausible. But in any case, not all similar skeptical arguments employ anything like (2*). Indeed, our (AH) is not so formulated, and it is difficult to see how anything like (A) and (B) are helpful when applied to (AH). In order to make the corresponding point about (AH), one would have to hold that the quality of the reasons that experience provides for external world beliefs is subject to the counterpart confusion. This would be a confusion about which features of experience are relevant to our reasons. We see no credibility in that suggestion. It is plausible that brains in vats can have the same reasons for their beliefs as those that any ordinary person has, and it is not plausible that their reasons must differ because their beliefs are mostly untrue. There are no good grounds to think that the appeal of (AH) depends on conflating propositions along the lines of (A) and (B).

that this restriction "limits the plausibility of the argument", but he chooses to discuss an argument with this premise nevertheless (ibid. 118). His view about what goes wrong with the premise in his argument can be applied to our own formulation of the argument, so we will not discuss further the restriction he imposes.

[21] Ibid. 122–3.

Hill offers a second reliabilist-friendly explanation of the appeal of skepticism. He says that people have two conceptions of evidential support. One of them is an internalist notion, and the other is an externalist notion.[22] According to the externalist concept, an experience provides evidential support for a belief provided there is a reliable cognitive process that begins with the experience and yields the belief. According to the internalist concept, evidential support is understood in terms of influence on subjective probabilities. Since purely sensory features of sense experiences can influence the strength of beliefs about the external world, on this conception they can be evidence for external world propositions. Since brains in vats or victims of deceptive demons adjust their subjective probabilities in light of experiences in just the way in which normal people do, their experiences have the same (internalist) evidential value as the experiences of normal people. But the reliable connections differ, so their experiences do not have the same (externalist) evidential value as the experiences of normal people. In Hill's view, since it is the externalist evidential value that is actually crucial for knowledge, the skeptic's argument fails. Hill says that there is a tendency to confuse these two conceptions of evidential support, and this explains the appeal of the skeptical argument to non-skeptics.

Granting for the sake of argument that there are these two concepts of evidential support,[23] we find here no explanation of the appeal of the skeptical argument. Presumably, the appeal is supposed to derive from our considering the argument while thinking of the internalist notion of evidential support, subjective probability. (The argument obviously has no appeal while thinking of the externalist notion.) However, using the internalist notion, the argument is distinctly unappealing. Both ordinary people and their envatted counterparts find themselves strongly inclined to believe the ordinary external world propositions on the basis of their sensory experiences. None of them has any inclination to believe skeptical hypotheses on the basis of their experiences. Hence, if people were using the internalist notion in thinking about skeptical arguments, they would conclude that the experiences provide vastly better evidential support for ordinary external world beliefs than for skeptical hypotheses. So the skeptical

[22] Hill says that this second explanation somehow accounts for the tendency to conflate (A) and (B). We doubt that the existence of two notions of the evidential support relation (the second explanation) can explain a tendency to conflate two accounts of what the evidence is (the first explanation).

[23] We doubt that these are pre-theoretical conceptions of evidential support. In our view, the proposed internalist conception of evidence is an implausible psychological reduction of the relation of evidential support, and the externalist conception is a reliabilist's invention.

argument would obviously fail.[24] This distinction, then, does not help to explain the appeal of skeptical arguments.

What does explain the appeal of skeptical arguments such as (AH), we think, is a reason for thinking that experiential evidence is at best neutral concerning the choice between the competing explanations. The reason is that the explanations all entail the experiential data, and the usual standards for evaluating explanations like simplicity and comprehensiveness do not obviously favor ordinary beliefs.

Once more, evidentialist theories of knowledge at least make sense of the concern. It is initially credible that experience is the ultimate evidence for the nature of the external world. It is initially credible that experiential evidence provides better reason to believe ordinary external world propositions than the skeptical alternatives. If, instead, one's experiential evidence really is at least equally good reason to accept some extraordinary contrary view of the external world, then one's reasons for ordinary thought are correspondingly weak. Philosophical reflection on the alternatives reveals that it is at least not obvious that standard explanations of experience are better explanations. After all, various skeptical scenarios account for all of the experiential data in one way or another, usually with a much leaner ontology. This clearly threatens the capacity of one's experiences to justify ordinary external world beliefs and thus, according to evidential theories, one's ability to have knowledge of the external world.

Even if this threat can be overcome by an evidentialist account of justification, the evidentialist approach faces a challenge. The skeptic has identified something that needs explaining away. A full evidential theory must say how ordinary thinking is best supported by experience. As will be discussed below, we think that when external world skepticism is properly appreciated, accomplishing this explanatory task is the major challenge that such skepticism poses.

Matters become somewhat more complex in the case of (TK). Unlike the other arguments for skepticism considered here, (TK) has a premise that simply asserts a denial of some external world knowledge:

1. S cannot know that (SK) is false.

By contrast, the other arguments derive denials of knowledge from other factors, such as the possibility of error. Of course, things could be said in support of (1), but they are likely to be along the lines of the considerations put forward in the other arguments. (TK) seems to us, therefore, to be a derivative argument and not in need of independent assessment. This, however, runs counter to recent discussions of skepticism, in which (TK) has figured prominently.

[24] This point would apply equally well to any less reductive notion of probability that also plainly turns on an internal condition like degree of belief.

In any case, many non-evidentialist theorists will confidently reject (1). Perhaps causal theorists will say that there is an appropriate causal connection between the fact that you are not a brain in a vat and your belief that you are not a brain in a vat. Exactly how this works is not entirely clear, since the causal consequences of one's *not* being a brain in a vat are difficult to discern. But whatever difficulty the causal theory has with this example carries over to beliefs that have no special connection to skeptical issues. One may know that one is not a professional pole-vaulter. Perhaps the fact that one is not a professional pole-vaulter is a cause of this belief, though it is difficult to identify any role that a negative fact like that might have played in causing one's belief. Presumably, the causal theory can be developed in some way that deals with this example, and it will similarly yield the desired result with respect to premise (1) of (TK).

Proper function theory will most likely be developed to yield the result that one's cognitive system is functioning properly when one believes that one is not a brain in a vat. Again, it is not entirely clear why this is true. A typical basis for denying a vat hypothesis would be that it just seems absurd. Though this impression may be a result of properly functioning cognitive mechanisms, it is doubtful that it is a basis for knowing. The idea of proper function is sufficiently obscure to allow a variety of interpretations, at least as it applies to this case. For instance, the evidentially supported acceptance of this belief that we describe in section VI below can be regarded as a proper functioning of our cognitive system.

Details about the formulation of reliabilism will make a difference in its implications concerning (TK). Exactly what processes produce one's belief that one is not a brain in a vat is unclear.[25] We can assume that the belief is true, and that it results from some inferential process involved in the formation of general beliefs about one's place in the world. Perhaps it is a reliable process. Similarly, one's belief that one is not a brain in a vat is safe—one would not easily be wrong about that—and, presumably, it results from a virtuous process. Thus, defenders of all these non-evidentialist theories will reject (1).

The tracking theory was designed in part to provide a different response to the argument than those so far considered. Apparently, many tracksters are convinced that the first premise is true, and thus are motivated to find some other way to avoid the skeptical conclusion. This leads them to deny closure principles, represented in our formulation by premise (3):

3. If S knows that (O) is true, and that (O) implies that (SK) is false, then S can know that (SK) is false.

[25] For extensive discussion, see "The Generality Problem for Reliabilism", [Ch. 6 in this volume].

The idea is that people do track ordinary truths but not the denials of skeptical hypotheses. You track the proposition that you have hands, but not the proposition that you are not a handless brain in a vat. You track the former because if it were not true, you would not believe it. But, it is said, if the proposition that you are not a handless brain in a vat were not true, you would still believe it (because you would be the victim of a deceit).

Thus, the tracking theory is supposed to support (1) and deny (3) of (TK).[26] Whether it actually has these results is questionable. What implication the theory has concerning (1) depends in an odd way upon details of the believer's situation. This is partly because there is a generality problem for methods of belief formation.[27] If anything is naturally regarded as the typical method of forming the belief that one is not a brain in a vat ("BIV"), it is something with a fairly broad range of application. It is plausible that one forms the belief because its negation seems absurd. Another possibility is that the process involves seeing how well the belief fits one's well-justified background beliefs. Applied to this case, the method would be something roughly like this: one's noticing the fit (in this case, the lack of fit) between the skeptical possibility and anything approximating the general view of the world incorporated in one's background beliefs about one's situation that one has good reason to accept, while having no positive reason to think that the skeptical possibility is actual.

If the typical method of believing that one is not a BIV is anything like this, then it is at most contingent that the belief fails to track the truth when one uses this method. It might be that the fact about the conditions under which one would have been a BIV, were this to happen, is that one would have been an informed beneficiary of an emergency brain extraction made possible by future medical progress. This seems considerably more realistic than the deceptive envatments that are the stuff of skeptical scenarios. If so, then one would have possessed that background information while envatted. Had one used the same broad "method" of testing the fit of the BIV possibility with justified background beliefs, then in light of one's background knowledge of the medical procedure, one would *not* believe that one was not a BIV. So on these assumptions, the belief *does* track the truth. On the other hand, the belief does not track the truth if one would have been a deceived BIV, were one a BIV.

[26] Not all versions of the tracking theory have the same implications. Contextualist versions of the theory hold that whether one has knowledge depends upon whether one tracks the believed proposition through the contextually determined worlds. If worlds in which one is a brain in a vat are not relevant in ordinary contexts, then perhaps this sort of tracking theory will also reject premise (1) and accept (3).

[27] Considerations such as those advanced in "The Generality Problem for Reliabilism" also apply to versions of the tracking theory that appeal to methods of belief formation.

Knowledge that one is not a BIV does not appear to turn on the relative likelihood of these two possibilities. It is the sheer possibility of BIV deception that threatens knowledge. So the tracking theory's account of the truth-value of (1) is implausible. And this makes its account of the force of the argument implausible. The trackster's intended explanation, again, is that people mistakenly accept the closure principle.[28] However, given the tracking theory, a person thinking about the argument who thinks she would be an informed BIV should not be moved by the argument. Such a person should reject premise (1). This is not realistic. Again, it is nothing more than the possibility of being a BIV that gives premise (1) such plausibility as it has.

There are also questions about how successful the other theories are in explaining why some people find (1) plausible. We have noted that a causal theory has some trouble in saying exactly why (1) is false. We have also noted that the trouble is not limited to skeptical possibilities like being a BIV. Whatever turns out to be the best causal theory of why one's belief that one is not now on Mars is knowledge, it is likely that the same explanation will apply to all negative location beliefs that are contrary to perceived fact, including the belief that one is not located in a vat. This will leave the intuitive difference unexplained: Why does it seem clear that one knows that one is not on Mars, but significantly less clear that one knows that one is not in a vat?

Reliability theorists would be hard pressed to find something plausibly thought to be unreliable in the case of the belief that one is not a BIV. Other theorists would be equally challenged to find something improper, or unsafe, in the formation of this belief. Thus, they have little to say about why (1) even seems true to anyone.

Non-evidentialist theories, we conclude, do a poor job of explaining the appeal that skeptical arguments clearly do have.

V. Fundamental Epistemological Questions

Our central contention so far has been that evidentialist theories do a much better job than non-evidentialist theories in making sense of skepticism. Knowledge requires good enough supporting evidence, and skeptical considerations at least cast doubt on whether the evidence we have is good enough. Non-evidentialist theories propose conditions on knowledge that are not even jeopardized by skeptical considerations. It is, of course, open to defenders of non-evidentialist theories to say that the reason people find skepticism

[28] Of course, we think it is no mistake to accept some form of the principle.

troubling is that they mistakenly think that knowledge requires strong evidence, and thus they mistakenly find significance in considerations that cast doubt on the quality of their evidence. There is, however, little reason to think that people would make this mistake if knowledge were what these theorists say it is. The mistake surely is not the result of philosophical indoctrination, since beginning philosophy students so often find skeptical considerations compelling without prior exposure to evidentialist writings.

Some philosophers find non-evidentialist theories of knowledge attractive, and not simply because of the easy denials of skepticism that they make available. In some cases, the starting point is to see human knowledge as an instance of some more general phenomenon like having a capacity to make various sorts of highly selective responses. That "knows" has some such broad application is suggested by sentences like "The thermostat knows that it is time to turn the heat on" or "The trees know when to shed their leaves." Presumably, there is no implication here that thermostats and trees have evidence. They just respond in some suitable way to their environment. Human cognitive systems facilitate human welfare largely by enabling people to improve their responses. This encourages the view that knowing is just having some such selective response capacity.

We find this view about knowledge ill-motivated. Sentences like those about the thermostat and the trees are not literally true. These uses of "know" serve approximately the same practical function as do plainly non-literal uses of psychological expressions other than "know", such as, "Nature *abhors* a vacuum," and "Water always *seeks* the lowest level."

In any case, this view of "knows" also misses the most challenging questions that skepticism raises. Suppose that the truth conditions for the word "knows" are conceded to some non-evidentialist theory or other. Important philosophical questions raised by skepticism remain to be addressed. How good are our reasons for believing the things ordinarily believed about the external world? Does the possibility of error, the existence of introspectively indistinguishable situations in which our beliefs are false, or the availability of alternative hypotheses consistent with ordinary experience, show that ordinary reasons are not good enough to justify the beliefs? If not, how do the reasons justify? If so, how do the reasons fall short, and how far short do they fall? Either way, what makes a reason any good at all?

These are the fundamental epistemological questions that are raised by the classic skeptical arguments. The traditional skeptical arguments are formulated as arguments about knowledge largely because it is assumed that knowledge requires good reasons. We see no good reason to deny that assumption. If it is jettisoned, the questions do not go away. Instead, they must be formulated more directly as questions about reasons or justification, not as questions about knowledge.

To confront the philosophical issues raised by external world skepticism, then, one must address questions about the quality of our reasons. Traditionally, this has been done by thinking about knowledge. This tradition warrants our continued adherence. Even if the tradition were abandoned, though, the questions that underlie skepticism would not have been thereby resolved. Philosophers would have to restate them as questions about the quality of the ordinary reasons for external world beliefs. Such philosophers would face the unenviable further task of explaining why an extensive tradition wrongly links these issues so closely to knowledge.

VI. Evidentialist Theories and the Skeptical Arguments

The threat of skeptical arguments is best understood as an attack on the sufficiency of the available justification for ordinary external world beliefs. This in turn is best understood as an attack on the sufficiency of the evidence for those beliefs. Taking the threat this way, it is easy to make sense of skepticism. In doing this, it should first be noted that an evidentialist theory of knowledge is not automatically anti-skeptical. For instance, it is open to an evidentialist to hold that one's basic experiential evidence for a proposition must entail it in order for the proposition to be well enough justified to be known.[29] Then arguments like (PE) and (II) can be used to show that no one has evidence that is this good.[30] Thus, one way that an evidentialist theory of knowledge can make sense of the external world skeptic's central concerns is to substantiate the skeptical conclusion.

This is not the best evidentialist position. A fallibilist evidentialist theory enables us to explain how people have external world knowledge and to account for the appeal of external world skepticism.

External world skeptical arguments give non-skeptical evidentialist theories of knowledge two main assignments. A first task is to give an informative account of the strength of evidence that is needed to have knowledge. Having just some slight reason to believe is insufficient, while having entailing basic evidence is not required. The problem is to identify and defend something in

[29] We assume that the entailment relation is construed broadly enough to enable experiential evidence to have entailments.

[30] In general, entailing evidence is insufficient for justification. Any evidence entails all necessary truths. So any reasonable evidential theory will impose other conditions on which evidence is good enough evidence for a given proposition, whether entailed by the evidence or not. Whatever these further conditions are, they can only strengthen the grounds for external world skepticism when added to an entailment requirement.

between. The other principal work for an evidentialist theory is to defend the claim that the evidence people have for their external world beliefs often meets this condition. An adequate defense must accept the constant possibility of external world error on the same basis for belief, and grant the introspective indiscernibility of error on the same basis. Those imperfections in the external world evidence should be shown not to make it too weak for knowledge. Also, some explanation must be given of why concerns about transmissibility and the existence of skeptical hypotheses that are compatible with any experience do not show the evidence to be inadequate.

Concerning the first project, we shall make a proposal that allows fallible reasons to justify sufficiently for knowledge. Then we shall explain the failure of the skeptical arguments and defend the view that the evidence people have for many ordinary propositions about the external world is often adequate.

A. Knowledge-Level Justification

We turn first to the project of stating what sort of justification is needed in order to satisfy the justification condition on knowledge. We conceive of the evidential condition on justification along the lines of the legal standard for conviction in criminal cases, proof beyond a reasonable doubt. The legal notion of proof is weaker than a mathematical proof, but stronger than a good reason to believe. We will call it strong reason to believe. A belief is epistemically justified sufficiently for knowledge, according to the "criminal" standard that we endorse, when one has strong reasons in support of it, no undefeated epistemic reason to doubt it, and no undefeated epistemic reason to believe that one's evidence for it is unreliable. We will use the word "defeater" to refer both to reasons to deny a proposition and to reasons to believe that one's evidence for a proposition is unreliable. The criminal standard is fallible, in that false beliefs can satisfy it. It is an evidentialist standard, in that only evidence plays a role in determining whether a belief is justified.

This assertion of the criminal standard for the justification needed for knowledge is brief and abstract. It asserts, without argument, that certainty is not required for knowledge. It joins the non-evidentialist responses to skepticism in not arguing for its fallibilism.[31] It leaves many important questions about strength of evidence unanswered. According to the criminal standard, knowledge requires strong reasons, but the account leaves open exactly how strong those reasons must be. There is no conspicuously correct fact of this matter.

[31] While we do not argue directly that fallibilism is true, we offer below an explanation of why the possibility of error does not defeat justification.

Fortunately, no external world skeptical argument turns on the details of the required strength of evidence. In fact, the difficulty in resolving this indefiniteness can help to explain the difficulty in answering some puzzling questions about knowledge that arise independently of worries about skepticism. For instance, it is quite credible that any ordinary person's visual experiences under ordinary conditions provide evidence enough to justify many common classificatory beliefs.[32] For example, when a typical contemporary adult has typical perception of a car, the person has justifying evidence for the belief that he or she sees a car, and for some beliefs about its color and shape. These are features that it seems people can "just see." But unless one happens to know more than usual about cars, it does not justify precise propositions about the car's weight or age. Why?

The difficulty in resolving this sort of question is an asset of the view. It helps to explain some reasonable disagreements in this area to assume that people are tacitly guided in knowledge attributions by something like our conception of the justification required. For instance, doubts about the identity of what one "just sees," the propositions that are thereby evidentially supported, and the strength of evidence provided, help to account for disputes about what one knows by looking, disputes that intensify when the stakes rise for accuracy.

There is also plenty of room for reasonable dispute among non-skeptics about the extent of testimonial knowledge. For instance, it is not at all clear when testimonial evidence in support of a conclusion is sufficient (in the absence of any reasonable doubt) to satisfy the justification condition on knowledge. Does the believer need evidence of the reliability of the source about the topic of the testimony? Again, the openness of our proposal is an asset, since the extent of testimonial knowledge is likewise not at all clear.

B. The Justification of External World Beliefs

We turn next to the second project. According to our criminal standard, knowledge-level justification requires three things: (1) strong supporting evidence, (2) the absence of undefeated reasons to doubt, and (3) the absence of undefeated reasons to think that one's evidence is unreliable. The goal now is to defend the claim that ordinary beliefs often meet this standard. We will do this largely by showing why the skeptical arguments do not prove otherwise.

[32] This is not to say that the perception is by itself sufficient evidence. Background evidence may be required.

It will be useful in the following discussion to treat separately two positions skeptics might defend. According to the first, ordinary experiential evidence is acknowledged to provide better reason to believe ordinary world propositions than to believe skeptical alternatives, but the skeptical possibilities are defeaters for these propositions. So the evidence is strong enough, but defeated, and condition (2) or condition (3) is not satisfied. According to the second skeptical position, our experiential evidence does not even provide better reason to believe ordinary propositions than the skeptical alternatives. On this view, it is condition (1) that is not satisfied.[33]

B1. Skeptical Alternatives as Defeaters

We begin with the first of these positions. One might think that, although ordinary experiences do provide reasons for ordinary beliefs, the skeptical possibilities show that people have epistemic reasons to doubt the propositions about the world that they ordinarily believe, or that people have epistemic reasons to doubt that their ordinary reasons are reliable.

Three of the skeptical arguments can plausibly be interpreted as relying on just this point. Consider first premise (2) of (PE). If people can be mistaken about the external world propositions that they believe, as the skeptical possibilities show, then apparently people have a reason to doubt these external world propositions. The beliefs are therefore not justified and consequently not known. Similarly, premise (2) of (II) seems to rely on this point. The existence of introspectively indistinguishable cases in which ordinary beliefs are false would seem to undermine knowledge because their existence constitutes a defeater.[34] It may be that such credibility as is had by premise (1) of (TK) also relies on this point.

The skeptical possibilities do not show that people have a reason to doubt ordinary propositions. For instance, a BIV hypothesis is not a reason to doubt ordinary beliefs, unless there is reason to believe that the BIV hypothesis is true. A reason to doubt that a proposition is true is a reason to think that something incompatible with the proposition is true. Reason to think, or proof, that a BIV hypothesis is *possible*, while things are introspectively the same, is not a reason to think that a BIV hypothesis is true. By contrast, good reason to think that

[33] A third position which skeptics might defend is that our evidence better supports ordinary propositions than skeptical alternatives, but it does not provide strong enough support to satisfy condition (1). Our response to the second view also applies to this alternative.

[34] It is possible to argue that introspective indistinguishability shows that we have no strong reason at all for our ordinary beliefs. Our response to this way of defending (II) is the same as our response to (AH) below.

a proposition is *probable*, to any extent, given one's basic evidence, is at least some slight reason to doubt that the proposition is true.[35] Similarly, the mere possibility that a BIV hypothesis is true is not a defeater by virtue of being a reason to think that ordinary evidence is unreliable.

Thus, the possibility of error on the same basis does not show that the basis fails to justify. Only evidence that supports the actual truth of some contrary proposition, or some undercutting proposition, defeats the justification. In skeptical scenarios, experiential evidence is misleading. But the fact that such deception is possible does not constitute a reason to think that the evidence actually misleads. The fact about what is possible would have to be supplemented with reason to think that the possibility actually obtains. This further requirement is easily overlooked. Some of the appeal of premises in the skeptical arguments results from the understandable mistake of thinking that the skeptical possibilities do provide reasons to doubt the propositions that are ordinarily believed. This argues against (2) of (PE) and (2) of (II).

The understandable mistake just described is one reason that premise (2) of (PE) may seem correct. In addition, there is a truth with which (2) might be confused. The premise can be casually formulated as "When you know, you have to be right." This sentence can express a truth. Although having knowledge-level justification for a belief does not guarantee the truth of that belief, knowledge does guarantee truth. The justification condition itself, however, allows falsehood, and it allows accidental truth of the sort that Gettier cases illustrate. Non-evidential theories of knowledge typically have no separate anti-Gettier condition. Proper causation, proper function, safety, or the like is supposed to exclude Gettier cases while otherwise filling the role filled by justification in an evidential theory. In accounting for how knowledge excludes the possibility of error, a separate condition to deal with the Gettier problem is an asset. Justification is not the only independent condition on knowledge. There is a "fourth condition" on knowledge. We do not have a specific proposal about what the fourth condition is. Our point here is that any solution will rule out as knowledge the "luckily true justified beliefs" that Gettier revealed. A condition that blocks this luck imposes a requirement of something approximating necessity for the truth of the belief. Justification alone does not accomplish this. Thus, when you know, you must satisfy the truth condition and you must satisfy the anti-Gettier condition. These two requirements for knowledge

[35] Suppose that a malicious or confused colleague asserts that you are a brain in a vat. Our view implies that either this undermines your knowledge or that this reason for doubt is defeated. We endorse the second disjunct for the most realistic versions of the example, although the details of the case could make the first option correct.

substantiate the intuitive truth that when you know, you have to be right. Thus, we can agree that there is some truth in this claim, without agreeing that justification is incompatible with the possibility of believing falsely on the same basis.[36]

We turn next to (TK). Premise (1) of (TK) denies that one can know the falsehood of skeptical scenarios. Premise (2) asserts that the knower sees the incompatibility of ordinary beliefs with the truth of the skeptical scenarios, and premise (3) is a closure principle connecting knowledge of ordinary beliefs with a capacity to know the falsehood of any proposition incompatible with a known fact.

Again, (TK) is an exceptional skeptical argument, in that its first premise is itself a skeptical proposition. Unlike the other skeptical arguments, (TK) is not clearly an attack on the adequacy of our justification. But the skeptical scenarios seem to function in (TK) in a way that is at least analogous to defeaters of the justification of ordinary beliefs. The alleged inability of people to know the falsehood of the scenarios is supposed to exclude knowledge of ordinary propositions for those who see the conflict between the two. Presumably, the exclusion arises from some vulnerability in ordinary evidence for the ordinary beliefs that makes skeptical scenarios defeaters of that evidence for those who see the conflict.[37]

There is no good reason to accept premise (1) of (TK), though. It is straightforwardly false. One's evidence for ordinary beliefs straightforwardly supports denials that the skeptical possibilities obtain. The evidence for ordinary beliefs argues against any proposition that one sees to conflict with the ordinary beliefs, while one has no new reason to doubt those beliefs or their support. This conspicuous implication by justified ordinary beliefs of the falsehood of the contrary propositions justifies denials of the contrary propositions. Consequently, one is justified in believing that the skeptical scenarios do not obtain. In this way, one can know that the skeptical scenarios are false.

The reasoning described in the previous paragraph seems to us clearly to refute premise (1) of (TK). To see this, it helps to note that this sort of reasoning is the way one usually knows typical facts about how things are not. For instance, this sort of reasoning is the usual way to know that one is not indoors, when according to one's basic experiential evidence one is outdoors. There are always indoor hypotheses that are metaphysically possible in conjunction

[36] The need for a fourth condition on knowledge makes room for a new skeptical concern only if there is reason to believe that the anti-Gettier condition cannot be met. There is no such reason. All promising approaches to a fourth condition on knowledge propose contingent conditions that we have good reason to believe actually obtain with roughly the frequency that knowledge appears to exist.

[37] Defenders of (TK) could defend its skeptical premise (1) by claiming that we have no strong grounds at all for our ordinary beliefs. Presumably, an argument such as (AH) would be used to support this contention. For our response, see section (B2).

with one's basic evidence. Still, one's experience and memory argue strongly enough that one is outdoors, and the conspicuous implication that one is not indoors is thereby amply supported. Absent some reason to believe the indoor proposition, it gives no reason to doubt the outdoor belief or the merits of one's evidence for it. The same goes for skeptical propositions, such as allegations of demonic deception. Absent some reason to believe that they are true, they give no reason to doubt ordinary beliefs or the merits of ordinary evidence for them.

Again, a possibility that is allowed by one's basic evidence, and incompatible with P, is not automatically a reason to doubt P. The possibility must have an epistemic status that bears on how things actually are, such as evidence that the possibility obtains, or at least evidence that gives the possibility some positive probability of being true. This fact about reasons to doubt is not just obviously correct. The lack of obviousness has significant explanatory utility. If people are guided by the conception that we defend of the justification needed for knowledge, according to which knowledge requires having no good reason to doubt what is justified, then they can easily be made to worry about any incompatible possibility. It conflicts with P, and there is nothing in the basic evidence anyone has that entails one rather than the other. In those important salient ways, a possibility incompatible with P is exactly like a good reason to doubt P. That can easily seem to be enough to render P unjustified, although it is not. Though this is not obvious, reason to doubt P is evidence that a possibility contrary to P obtains in fact. Evidence of sheer possibility is not that. But again, this is not an inescapable fact. Skeptics who fail to agree and locate reasons to doubt in other possibilities have made no manifest blunder.

We leave open numerous significant issues. There are questions about whether every known probability of error gives rise to a good reason to doubt. For instance, according to current physical theory, a brief time ago there was some minuscule objective chance that any given ordinary object, say, a teacup, would by now have virtually all of its matter dissolved by quantum fluctuations, save for the surface facing a perceiver.[38] Were that terrifically improbable event of dissolution actually to have happened, the perceiver would see, for a moment before its collapse made the cup's decomposition obvious, a mere cup facade, not a cup. Perceivers who know enough current physical theory to know about this objective chance thus seem rationally bound to assign some positive probability to the possibility that the things perceived at any given time are mere facades rather than ordinary objects.

[38] The epistemic significance of this sort of physical possibility is discussed in John Hawthorne's monograph, *Knowledge and Lotteries* (Oxford: Oxford University Press, 2003).

Our view of the justification required for knowledge counts any probability for someone of the truth of any known contrary to a proposition as a reason for the person to doubt the proposition. So we seem committed to denying perceptual knowledge by the physically informed of the current existence of ordinary perceived objects.

We have nothing conclusive to say about this. We note that the reason to doubt is an inference from details of current physical theory, and there is inductive reason to doubt that the current theory is the full truth of the matter. It may also be epistemically significant that this sort of physical possibility has never actually occurred to any ordinary macroscopic object, as far as we know. Perhaps the total evidence of the physically informed on balance supports suspension of judgment about the existence of a positive probability of dissolution. This lack of any known case of the phenomenon stands in contrast to lotteries where it is known that some ticket, whose victory was epistemically highly improbable, nonetheless actually wins. It seems clear that lottery cases are not knowledge. Our account implies that there is inadequate justification, because of a reason to doubt. But it may be that not just any sort of basis for probability is a genuine reason to doubt. Perhaps a probability of an event in a reference class that has never been observed, according to an arguably transitional physical theory, does not give reason for doubt that undercuts justification, as do probabilities derived from actual events in the same reference class.

It also may be that those who know about this objective chance do have this minute probability of falsehood as a reason to doubt. So by ordinary perception they do not quite know of the presence of complete solid objects. If this is the case, and if they have no other reason to doubt, then these perceptual beliefs have such splendidly high probability for them, and are otherwise so well supported, that the practical difference of the beliefs from knowledge is nil. Their epistemic difference from knowledge, though notable, may be similarly small. Perhaps lottery beliefs based on high enough probabilities are likewise "almost known." Such issues plainly deserve further investigation.

This completes our discussion of the skeptical arguments other than (AH). Our ordinary beliefs are sufficiently justified for knowledge only if we have no undefeated reasons to doubt them. The three skeptical arguments just considered rely on the assumption that skeptical scenarios provide reasons to doubt. We have argued that they do not.

B2. Skeptical Alternatives and the Strength of the Evidence

Beliefs have knowledge-level justification only if we have strong evidential support for them. The second skeptical position described earlier in this section

contends that we lack strong evidence for external world beliefs. Premise (1) of (AH) relies on just this idea. It asserts that ordinary beliefs are not better supported by experience in the first place than are skeptical world views.

Skeptics give no good reason to believe this premise. And the premise does seem initially incredible. It surely seems that a typical full-blown experience of a warm summer day has to turn out somehow to give one better reason to believe that one is actually experiencing a warm summer day than to believe that one is a brain in a vat being fed warm summer day experiences. With premise (1) unsupported, (AH) does not argue successfully against our having external world knowledge.

It is a further worthwhile project to *show* that the premise is mistaken. To do this, some account is needed of what makes experiences provide basic evidence for ordinary beliefs. It seems clear that experiences do this. Explaining how they do so is a difficult project that remains to be accomplished. Two approaches are surveyed below. Each is found wanting, although not without merit. Showing that experiences really do provide basic evidence for ordinary beliefs is the main work that remains to be done to have a complete answer to external world skepticism.

It can seem so clear as to be a fundamental fact that typical experiential evidence is better evidence for ordinary external world beliefs than for any specific skeptical alternative, or for any disjunction of skeptical alternatives. One sort of view endorsing this position can be found in the work of Roderick Chisholm. In his *Theory of Knowledge* Chisholm proposes as fundamental epistemological principles such principles as the following:

> If S believes that he perceives something to have a certain property F, then the proposition that he does perceive something to be F, as well as the proposition that there is something that is F, is one that is *reasonable* for S.[39]

Chisholm does not derive this principle from some other more fundamental truths. He seems to assert, as a fundamental epistemological fact, that our perceptual beliefs do support our external world beliefs. More recently, James Pryor has defended a similar view:

> My view is that whenever you have an experience as of *p*, you thereby have immediate *prima facie* justification for believing *p*. Your experiences do not, in the same way, give you immediate *prima facie* justification for believing that you are dreaming, or being deceived by an evil demon, or that any of the skeptic's other hypotheses obtain.[40]

[39] Roderick Chisholm, *Theory of Knowledge* (Englewood Cliffs, NJ: Prentice-Hall, 1966), 45.

[40] James Pryor, "The Skeptic and the Dogmatist", *Nous*, 34 (2000), 517–49. The quotation is from p. 536.

Chisholm's principle makes the justification of external world beliefs dependent upon one's beliefs about perceptual experiences. Pryor's view makes the justification dependent upon the experiences themselves. As we interpret these views, both hold that it is a fundamental epistemological fact that our experiences, or our beliefs about experiences, provide immediate perceptual justification for our ordinary external world beliefs.

Many details of the Chisholm/Pryor view need to be worked out. A prominent issue concerns exactly which experiences and propositions any such view applies to. An experience of redness, say, may immediately justify (*prima facie*) the belief that there is a red object before one, but the experience of a ponderosa pine will not, for most observers, similarly justify the belief that there is a ponderosa pine present. Some explanation of the difference is needed.[41]

Furthermore, as a response to premise (1) of (AH) this approach is less than fully enlightening. Although it does indeed seem true that ordinary experiences justify propositions about ordinary physical objects better than they justify propositions about what an evil demon is doing, it also seems that there must be some reason why the experiences support the ordinary beliefs and not the rivals. It is not enough to say that the experiences "just do" support certain propositions. Chisholm and Pryor propose plausible general principles.[42] But simply asserting general principles that imply that experiences justify ordinary beliefs does not explain why they do this. We need either some explanation of why the experiences do justify the external world beliefs or some explanation of why the extension of the epistemic support relation turns out to be inexplicable in the relevant cases. Either way, the Chisholm/Pryor response does not complete an evidentialist reply to skepticism until it is itself completed.

A second evidentialist response to (AH) takes a different tack. It holds that the commonsense beliefs enter into better explanations of our experiences then do the skeptical hypotheses. It asserts that this makes the experiences better reason to believe the ordinary world views.[43] This is an initially promising approach. It is surely quite doubtful that the skeptical hypotheses provide explanations that are as good as the explanations provided by ordinary

[41] Ernest Sosa raises this issue in his contribution to *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*, co-authored with Laurence BonJour (Oxford: Blackwell, 2003). See pp. 122 f.

[42] Chisholm's principle needs a "no defeaters" clause. He adds a condition along these lines in a revised principle in *Theory of Knowledge*, 2nd edn. (Englewood Cliffs, NJ: Prentice-Hall, 1977), 78.

[43] Evidentialist denials of premise (1) need not rest on either the claim that it is a brute fact that our experiences provide better reason to believe ordinary external world propositions than to believe skeptical alternatives, or the claim that ordinary beliefs are included in better explanations of experiences than are

beliefs.[44] The evil demon hypothesis, for instance, usually leaves entirely unexplained the existence, powers, and deceptive motives of the demon, and one's own existence. There is no obvious and economical explanation of these things. The alternative explanations bring in a kind of complexity—amazingly sophisticated computers, evil geniuses monitoring our thoughts and inducing lengthy undetected deceptions, implausibly orderly dreams. These explanations seem *ad hoc*, complex, and ridiculous. The pure chance skeptical hypothesis, according to which one's experiences just happen to occur with no causal explanation, implies less that needs explaining. But the apparent orderliness of the experiences and their coherence with memory makes a chance hypothesis seem an extraordinarily weak explanation. The ordinary external world explanation has none of these liabilities. It seems to be best by far, and good enough to justify.

It is well known that this response is not without difficulties.[45] There are important questions about exactly what counts as a better explanation and exactly why the skeptical hypotheses yield inferior explanations. It is also unclear whether it is sufficient for the ordinary beliefs to be justified that the explanation is in fact best (and good enough). It might be required, instead, that the believer realize that it has this status or that the believer's evidence support the proposition that it has this status (whether this proposition is believed or not). It is not implausible to think that our evidence does support this proposition, but we are not able to offer a compelling defense of this claim here.

These are not hopeless difficulties. Evidentialist responses along the lines of the two discussed here show promise of completing the remaining project, that of establishing that our evidence really does support our commonsense beliefs.[46]

skeptical hypotheses. If the ordinary beliefs are made "epistemically probable" by the experiences, while skeptical alternatives are not, then the experiences provide better reason for the ordinary propositions than the skeptical alternatives. Of course, the evidentialist utility of this appraisal depends on giving an evidentialist explanation of this epistemic sort of probability. That is no easy task. Since we think that the strength of "epistemic probability" must be explained by the strength of epistemic reasons, rather than vice versa, this sort of third approach does not seem promising to us.

[44] See Jonathan Vogel, "Cartesian Skepticism and Inference to the Best Explanation", *Journal of Philosophy*, 87 (1990), 658–66.

[45] See Richard Fumerton, *Metaepistemology and Skepticism* (Lanham, Md.: Rowman and Littlefield 1995), 207–14.

[46] Contextualism does not offer any help here. Any contextualist view that gives evidence a justifying role shares the task of saying how ordinary evidence justifies ordinary belief, at least in those contexts in which ordinary "knowledge" attributions are true. Contextualist views that deny evidence this role share the liabilities discussed above for non-evidentialist theories of knowledge in addressing the fundamental issues raised by skepticism.

If no response such as this were to be correct, then skepticism would be true after all. Justifying evidence is required for knowledge. If our evidence about the external world fails to give adequate support for external world propositions, then skepticism is correct.[47]

Evidentialism has the virtue of squarely facing skeptical challenges. It makes best sense of skepticism.

[47] There are important matters of relative detail. For instance, if it turns out that statistical evidence is insufficient for knowledge of inductively drawn conclusions, then a variety of commonly accepted external world propositions are not really known. Nothing in our criminal standard of justification for knowledge readily resolves this issue. But it would be a comparatively modest denial of knowledge, and not a version of external world skepticism. Skepticism with regard to the whole past, on the other hand, is clearly a significant form of external world skepticism. It is equally clear that by the criminal standard we have sufficient evidence to know many propositions about the past.

Works Cited

[1] Alston, William, "Internalism and Externalism in Epistemology", *Philosophical Topics*, 14/1 (1986). Reprinted in [5], 185–226.

[2] Alston, William, "Epistemic Circularity", *Philosophy and Phenomenological Research*, 47 (1986), 1–30. Reprinted in [5], 319–49.

[3] Alston, William, "The Deontological Conception of Epistemic Justification", *Philosophical Perspectives*, 2 (1988), 257–99. Reprinted in [5], 115–52.

[4] Alston, William, "An Internalist Externalism", *Synthese*, 74 (1988), 265–83. Reprinted in [5], 227–45.

[5] Alston, William, *Epistemic Justification: Essays in the Theory of Knowledge* (Ithaca, NY: Cornell University Press, 1989).

[6] Alston, William, "Foley's Theory of Epistemic Rationality", *Philosophy and Phenomenological Research*, 50 (1989), 135–47.

[7] Alston, William, *The Reliability of Sense Perception* (Ithaca, NY: Cornell University Press, 1993).

[8] Alston, William, "How to Think about Reliability", *Philosophical Topics*, 23 (Spring 1995), 1–29.

[9] Annis, David, "A Contextualist Theory of Epistemic Justification", *American Philosophical Quarterly*, 15 (1978), 213–19.

[10] Armstrong, David, *Belief, Truth and Knowledge* (Cambridge: Cambridge University Press, 1973).

[11] Audi, Robert, "Memorial Justification", *Philosophical Topics*, 23 (1995), 31–45.

[12] Audi, Robert, *Epistemology: A Contemporary Introduction to the Theory of Knowledge* (New York: Routledge, 1998).

[13] Audi, Robert, "Doxastic Voluntarism and the Ethics of Belief", in Matthias Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (Oxford: Oxford University Press, 2001), 93–111.

[14] Baergen, Ralph, *Contemporary Epistemology* (Fort Worth: Harcourt Brace, 1995).

[15] Bennett, Jonathan, "Why Is Belief Involuntary?", *Analysis*, 50 (1990), 87–107.

[16] Bergmann, Michael, "A Dilemma for Internalism", in Thomas Crisp, Matthew Davidson, and Dave Vanderlaan (eds.), *Knowledge and Reality: Essays in Honor of Alvin Plantinga* (Dordrecht: Kluwer, forthcoming).

[17] BonJour, Laurence, "Can Empirical Knowledge Have a Foundation?", *American Philosophical Quarterly*, 15/1 (1978), 1–14.

[18] BonJour, Laurence, "Externalist Theories of Empirical Knowledge", *Midwest Studies in Philosophy*, 5 (1980), 53–73.