

One Wage of Unknowability

Dennis Whitcomb
Forthcoming in *Synthese*

0. Stage setting

Suppose for *reductio* that I know a proposition of the form p and I don't know p . Then by the factivity of knowledge and the distribution of knowledge over conjunction, I both know and do not know p ; which is impossible. Propositions of the form p and I don't know p are therefore unknowable.¹ Their particular kind of unknowability has been widely discussed and applied to such issues as the realism debate.² It hasn't been much applied to theories of the nature of knowledge. That is what I'm going to do here.

Most centrally, I'll draw out some implications of unknowability for a standard package of epistemological views: evidentialism and fallibilism. The basic idea behind evidentialism is that knowing requires having adequate evidence. This idea raises the question of what "adequate" means. Fallibilism answers that question. *Its* basic idea is that adequate evidence need not entail what we know, but only needs to *make probable* what we know. There are important subtleties concerning the best ways to state evidentialism and fallibilism. I will come back to those subtleties later. But for now, with the basic ideas roughly glossed, we can work with the following statement of the package combining them:

fallibilist evidentialism: there is some degree n of evidential support short of entailment such that necessarily, for all persons S and propositions p , if S 's evidence supports p to at least degree n , and S meets the non-evidential conditions on knowledge (such as properly basing her belief on her evidence while that belief is true and unGettiered), then S knows that p .

This paper does four things. First, it uses unknowability to argue that fallibilist evidentialism is false. Second, it considers some objections to this argument. Third, it inquires into how far this argument applies to additional, nonevidentialist forms of fallibilism. Fourth, it draws a lesson from the discussion.

1. The argument from unknowability

Let D be the proposition that the president is dead and I don't know the president is dead. Now consider an oracle. Suppose (a) that 90% of the oracle's assertions are true. Also suppose (b) that the oracle asserts D . Finally, suppose that I know both (a) and (b). Since I know these things, they are part of my evidence corpus. And they jointly constitute evidence for D . So, I have some evidence for D .

Now suppose that the oracle's accuracy is even greater; let it be anything short of 100%. Again suppose that I know how accurate the oracle is, and that I know that she asserts D. I again have some evidence for D. And since the oracle's level of accuracy can be anything short of 100%, the strength of my evidence for D can be anything short of entailment. Moreover, the non-evidential conditions on knowledge can also hold of me for D.

In sum then, I can have arbitrarily strong (short of entailing) evidence for D, while also meeting all the non-evidential conditions for knowing D. And yet I can't know D. Therefore, fallibilist evidentialism is false. If there is any degree n of evidential support such that one knows p whenever p is supported to at least degree n by one's evidence (and the non-evidential conditions on knowledge also hold), then n amounts to entailment.³⁴

2. Elaboration via objections and replies

The foregoing argument is very condensed, almost just the skeleton of an argument. Let me now put some meat on its bones. I'll fill in the details and bring out the motivations for the various claims the argument makes. The method for doing as much will be to raise and answer some objections.

2.1. The straw man objection

One objection is that fallibilist evidentialism as I've defined it is a straw man position that no one really holds. But in fact, the view has numerous adherents. These include Feldman and Conee, who write that

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.⁵

In answering to this first task, they propose that the required strength of evidence is weaker than that of mathematical proof. They also say that it is "fallible, in that false beliefs can satisfy it" – i.e. that one's evidence can support falsehoods with the strength needed for knowledge.⁶ Furthermore, Conee claims that this view is standard among nonskeptical evidentialists:

According to nonskeptical evidentialist accounts of the strength of evidence needed for knowledge, some sort of nonentailing evidence is sufficient.⁷

So fallibilist evidentialism is no straw man. Now as it happens, there are evidentialists who reject the view, including skeptics and even some non-skeptics.⁸ But many accept the view, including Conee and Feldman.

2.2. The objection from evidence possession

A second objection is that the argument assumes something that many evidentialists would reject, namely that everything one knows is part of one's evidence. But that assumption is inessential; let me now explain why.

The argument from unknowability focuses on testimonial evidence. Whatever it is to be testimonial evidence, the testimonial evidence one gets from the oracle can be as strong as any possible testimonial evidence, save testimonial evidence that entails what it is evidence for. So, *whatever* it is to be testimonial evidence, one can have testimonial evidence of arbitrarily high non-entailing strength, but nonetheless not know. A generalized version of the argument from unknowability therefore shows fallibilist evidentialism to be false regardless of whether everything one knows is part of one's evidence. That view about knowledge and evidence focuses the argument by making it less abstract, but the argument is sound even without it.

2.3. The objection from evidential strength

A third objection questions, not what evidence I *possess* in the oracle cases, but the *strength* of my possessed evidence in those cases. According to this objection, my evidence for the unknowable proposition is not arbitrarily strong short of entailing. I know that the oracle asserts the proposition, and I know her accuracy to be arbitrarily high short of 100%, and these items of knowledge are part of my evidence. Yet I do not have arbitrarily strong (short of entailing) evidence for the proposition in question. Or so claims this objection. Let me motivate this claim and then explain why, despite that motivation, the claim is mistaken.

We cannot validly infer, from the mere assumption that my evidence includes *the oracle is arbitrarily accurate short of 100% and she asserted p*, that my evidence supports *p* arbitrarily strongly short of entailment. For there might be reasons to doubt her in the particular case. For instance, suppose she asserts *the president is dead and she is always wrong about presidents*. Here there is a reason to doubt her, to wit: her assertion cannot be true given that she makes it. Only absent such countervailing reasons does *the oracle is arbitrarily accurate short of 100% and she asserted p* render *p* supported arbitrarily strongly short of entailment.

You might think that such countervailing reasons are present in cases at work in the argument from unknowability. If that thought is right, then I do not in those cases have arbitrarily strong nonentailing evidence. Hence the thought motivates our third objection.

However, it is not clear what reason there could be to doubt the oracle, when she says that *p* and *I don't know it*. That proposition, unlike *p* and *she is always wrong about p*, can be true given that she asserts it. What then could be the reason to doubt it? One answer is that each conjunct of *p* and *I don't know it* is evidence against the other. If *p* is true, that fact is evidence that I know *p*; it rules out one of the ways I might fail to know *p*. And if I fail to know *p*, that fact is evidence against *p*; it rules out one of the ways in which *p* could be true, namely its being true-and-known-by-me. Is this internal

disconfirmation, this disconfirmation of each conjunct by the other, a reason to doubt *p* and *I don't know it* when the oracle asserts that claim?

No, it is not. For suppose my friend tells me the president is dead but not assassinated. The first conjunct of this claim is evidence against the second, and *vice versa*. But that is no reason to doubt my friend. Similarly with the oracle, then. The disconfirmation of *p* by *I don't know p* and *vice versa* is no reason to doubt the oracle.⁹

Could there be some *other* reason to doubt her? You might think that such a reason is illustrated by the following line of thought:

In listening to the oracle I come to be *justified to believe* each conjunct of the conjunction she asserts.¹⁰ Moreover, each of those conjuncts might be true. Hence, if I were to believe them, my beliefs might well be justified and true. Those beliefs might even be *knowledge*. Hence I can have a justified true belief in, and indeed *knowledge of*, each conjunct the oracle asserts. Hence I can have a justified true belief in, and indeed *knowledge of*, their conjunction.

This line of thought is mistaken, because it entails that I can know an unknowable proposition. Where then is the mistake? One view is:

The mistake is in the claim that I am justified to believe the oracle in the first place. If I were, then since the rest of the line of thought is correct, it would follow that I can know an unknowable conjunction. Hence I am not justified to believe her in the first place. *Hence it is false that the oracle gives me arbitrarily strong nonentailing evidence.*

If this view were correct, it would constitute a reason to doubt the oracle in the cases at work in the argument from unknowability. However, the view is not correct. For there are other more plausible views about where the foregoing line of thought is mistaken. For one thing, it is (plausibly) false that whenever I am justified to believe *P* and justified to believe *Q*, I am justified to believe *P&Q* – as the preface paradox suggests.¹¹ Also and more saliently, it is (plausibly) false that whenever I can have a justified true belief in a conjunction, I can know that conjunction – as is suggested by Gettier cases involving conjunctions. The foregoing line of thought mistakenly assumes both of those falsehoods; hence it does not illustrate a reason to doubt the oracle.

We've been considering the objection that I do not get arbitrarily strong nonentailing evidence from the oracle. This objection is correct only if, in those cases, I always have a reason to doubt her. Several attempts to locate such reasons failed; I conclude that they don't exist.

2.4. The objection from nonevidential conditions

A fourth objection targets the claim that the nonevidential conditions on knowing can jointly obtain in the oracle cases. If in these cases my belief always fails to meet some nonevidential condition on knowing, then my lack of knowledge is consistent with fallibilist evidentialism. And so this fourth objection claims that the nonevidential conditions cannot jointly hold in the oracle cases.

But *prima facie*, it is hard to see how that claim could be right. Perhaps the president is dead and not known by me to be dead; then my belief would be true. Perhaps I base my belief on the oracle's testimony, combined with my knowledge of her reliability. Then my belief would be properly based. And perhaps there are no Gettier high jinks either: no misleading testimony just out of earshot, no oracle holograms surrounding the oracle, and so on and so forth. Why couldn't all this happen? It is hard to see how it couldn't.

You might ask: but isn't one of the nonevidential conditions on S's knowing P simply *that P is knowable for S*? And isn't this condition unmet in the cases at work in the argument from unknowability?

Well, clearly, if S knows *p* then *p* is knowable for S. But the non-evidential conditions on knowing are conditions that can be adequately specified without reference to knowledge (they have to be, to play their role in the evidentialist attempt to define knowledge noncircularly). And the condition that P is knowable for S *cannot* be adequately specified without reference to knowledge. On pain of circularity then, the condition that *p* is knowable for S cannot be one of the non-evidential conditions on S's knowing that *p*.¹²

Now you might ask: but isn't there a straightforward way in which the nonevidential conditions fail to hold in the oracle cases? In particular, don't I always in those cases have a *defeater* for the proposition the oracle asserts - said defeater being the proposition that I can't know the former?¹³

I don't think this defeater-based idea will work. It takes *R is unknowable* to be a defeater for *R*. But why should that be so? One answer is that it is a brute first principle, an unexplained explainer. But if the defeater theorist gives *that* answer, she renders her theory of knowledge circular - because she defines knowledge in terms of defeat and defeat in terms of knowability. Hence she should give a different account of why *R is unknowable* defeats *R*. But what could the account be, such that it avoids circularity? She might reply: just this - "*R is unknowable* defeats *R* because its addition to one's evidence renders *R* unjustified by one's evidence. That's just *what it is* to be a defeater."¹⁴ But why should the addition of *R is unknowable* to one's evidence render *R* unjustified by one's evidence? In some cases at least, it is hard to see why it should - for instance our oracle cases. In these cases it seems that the status of *R* as justified by my evidence *remains the same* under the relevant addition to my evidence - for as I (in effect) argued above, said addition does not render *R* any less supported by my evidence. In these cases at least, the *support* relations between my evidence and *R* are invariant under the addition to my evidence of *R is unknowable*; and so plausibly, the *justification* relations between my evidence and *R* are similarly invariant.

In summary, there is nothing *prima facie* impossible about the nonevidential conditions on knowledge all holding in the oracle cases. This should lead us to doubt the power of the objection from nonevidential conditions, at least until we find convincing arguments that the objection holds. We examined two attempts to build such arguments: one from

the idea that knowledge requires knowability and one from defeater theory. The resulting arguments were both unconvincing.

2.5. The objection from the lottery

A fifth objection is that the oracle cases are nothing new, but just familiar lottery cases under a different guise. Familiarly, it is very implausible that we can know, simply on the basis of the probabilities, that a ticket in a large fair lottery is a loser - even if it *is* one. How are the oracle cases different from these familiar lottery cases?

Well, unlike the claim that a ticket is a loser, the oracle's claim is demonstrably unknowable given only the assumption that knowledge is factive and distributes over conjunction. The bullet-biting response to the lottery – the response saying that if the lottery is large enough we *do* know – is thus a thorough nonstarter with the oracle.¹⁵

2.6. The objection from infallible oracles

A sixth objection is that if the oracle cases I have described (in which the oracle's accuracy is arbitrarily high short of 100%) are possible, then so too are similar cases in which the oracle's accuracy is 100% - and that the latter possibilities are problematic.

Suppose that I knew that an oracle was infallible, and that I knew that she asserted *p* and *I don't know p*, and that from these items of knowledge I competently deduced that *p* and *I don't know p*, while retaining my knowledge that she is infallible and that she asserted *p* and *I don't know p*. Then I would competently deduce something from known premises, while retaining my knowledge of those premises, but not come to know the thing deduced. Closure would fail egregiously. This would be bad.

The sixth objection, then, is that if the near-infallible oracle cases I have described are possible, then so too is this infallible oracle case in which closure fails egregiously.

Clearly, we should consider infallible oracle cases alongside the near-infallible oracle cases in the argument from unknowability. If the *near-infallible* oracle cases are possible, then either

- (a) the *infallible* cases differ in some relevant way making *them* impossible, or
- (b) the infallible cases don't differ in any relevant way, and so closure fails egregiously.

Of course, one thing to say is that (a) and (b) are each so untoward that, since the possibility of near-infallible oracle cases entails their disjunction, near-infallible oracle cases are impossible. And this is just what a partisan of the sixth objection *would* say. But I don't think it is correct. In fact, I think, there is a relevant difference between infallible and near-infallible oracle cases, a difference that should lead us to accept the possibility of the latter but not the former.

In particular, near-infallible oracle cases are (in at least one important way) analogous to lottery cases that are obviously possible. Whenever one is in cases of either of these two

sorts, one has very strong *nonentailing* evidence for a true proposition, but doesn't know that proposition. *Infallible* oracle cases are not like this; they feature *entailing* evidence. Thus we have a reason to believe in the possibility of *near-infallible* oracle cases (namely their similarity to lottery cases), which is *not* a reason to believe in the possibility of *infallible* oracle cases. Hence there is a relevant difference between infallible and near-infallible oracle cases; hence the sixth objection is unconvincing.

None of our objections refute the argument from unknowability. I conclude that the argument is sound. Let me now inquire into how far it can be generalized.

3. Extensions

Does the argument from unknowability refute even non-evidentialist forms of fallibilism? It is tempting to think that it actually refutes the general view that there is a non-maximal degree of justification such that, whenever one has that degree of justification and one also meets the nonjustificatory conditions on knowledge, one knows. If this general view is wrong, then versions of it that assume non-evidentialist theories of justification are also wrong.

However, if we try to actually apply the argument to those other theories, we find that they are too unspecific for the task. For instance, consider process reliabilism, the theory that a belief is knowledge-level justified if and only if it was produced by a reliable enough process. Do our oracle cases commit process reliabilists to the view that the only level of reliability sufficient for knowledge-level justification is perfect, 100% reliability?

It is not clear, because it is not clear what belief-forming process I use when trusting the oracle. Is that process believing an oracle whose accuracy is arbitrarily high short of 100%? If so, then our oracle cases refute fallibilist process reliabilism in the same way they refute fallibilist evidentialism. But why shouldn't we say that my belief-forming process is something else, like believing a confusing proposition?¹⁶ If *that* is the process, the oracle cases may not refute fallibilist process reliabilism after all. For, plausibly, that process is too unreliable to confer knowledge-level justification. Whether the oracle cases refute fallibilist process reliabilism thus depends on which beliefs are produced by which processes.¹⁷

Coherentism, according to which a belief is knowledge-level justified if and only if it is a member of a coherent enough belief corpus, fares similarly.¹⁸ There may be plausible characterizations of coherence on which *p and I don't know p* always introduces some incoherence into one's belief corpus. On such characterizations, that belief could not count as having arbitrarily strong non-maximal justification; and so it would not constitute a counterexample to fallibilist coherentism. However, there may be other plausible characterizations of coherence on which that belief is liable to have any level of coherence whatsoever. That belief is, after all, internally consistent. It might even be combinable with other beliefs to form a *maximally* coherent whole. If it can be so

combined (while also being true and unGettiered and so on), then it gives us a counterexample not just to fallibilist coherentism, but to coherentism *tout court*.

So, depending on the details, fallibilist (and infallibilist!) coherentism either may or may not be vulnerable to the argument from unknowability. It depends on the nature of coherence, just as with reliabilism it depends on the nature of belief forming processes. These views are simply not developed enough to directly confront the argument from unknowability.

In a way, this tells in favor of evidentialism. At least evidentialism is developed enough to confront the argument from unknowability directly. The other views escape the argument in an unhelpful way; they escape it through underspecificity.

4. A lesson

Let me try to draw a lesson from all this. I'll start by returning to the point that there are important subtleties concerning the best ways to state both evidentialism and fallibilism.

The argument from unknowability establishes that no nonentailing degree of evidential support is *sufficient* (given that the nonevidential conditions hold) for knowledge. It shows that for any nonentailing degree of evidential strength whatsoever, there is a case in which one has evidence of that strength, and the nonevidential conditions on knowledge hold, but one does not know. But consistently with the existence of these cases, there may also exist cases in which one knows a proposition that one's evidence does not entail. That is to say: it does not follow, from the mere fact that no nonentailing degree of strength is *always* enough for knowledge, that no nonentailing degree of strength is *ever* enough for knowledge. The argument from unknowability therefore leaves open the possibility that one can know things that one's evidence does not entail.

Fallibilist evidentialists should therefore do the following. They should change their fallibilism from the *sufficiency thesis* that there is a non-entailing degree of evidential support that is sufficient for knowledge given that the non-evidential conditions all hold, to the *possibility thesis* that it is possible to know things one's evidence does not entail.

This restatement circumvents the argument from unknowability, but it comes at a price. Contrary to the spirit of fallibilist evidentialism, it weakens the role of fallible evidence in the theory of knowledge.¹⁹ Or so I'll now argue.

As a first step, I'll state evidentialism precisely. Again, the basic idea is that *knowledge requires adequate evidence*. The notion of "adequacy" at work here is surprisingly tricky. So I'll start with the simpler idea that *knowledge requires evidence*. Even here, there are several different claims. One is that there is a positive amount of evidential support, such that anyone has that much evidential support for anything she knows. Another is that for anyone, there is a positive amount of evidential support, such that she

has that much evidential support for anything she knows. This second claim, unlike the first one, allows the relevant amount of support to differ across cases. Another claim still is that for anyone and any proposition, there is a degree of evidential support such that if she knows that proposition, her evidence supports it to that degree. This last claim, unlike the first two, allows the relevant amount of support to differ across propositions in the same case.

It helps to state these views in a way that highlights their logical form. This requires some terminology. Let “cases” be centered worlds (i.e. persons in worlds at times), “S” be the person in a given case, c range over cases, p over propositions, and n over nonzero degrees to which S’s total evidence supports p .²⁰ Let “Ksp” abbreviate “S knows P”. The views can then be stated as follows, in order of decreasing logical strength:

- (1) $\exists n \forall c \forall p$: If in c Ksp, then S’s evidence supports p to at least n
- (2) $\forall c \exists n \forall p$: If in c Ksp, then S’s evidence supports p to at least n
- (3) $\forall c \forall p \exists n$: If in c Ksp, then S’s evidence supports p to at least n

Again, these are just versions of the idea that knowledge requires evidence; we haven’t even tried, yet, to add the idea that this evidence must be “adequate”. Since these claims are so minimal, they are consistent with *infallibilist* evidentialism, which can be precisely stated as

- (4) $\exists n \forall c \forall p$: If in c Ksp, then S’s evidence supports p to n , and $n =$ entailment

Whatever else they end up doing, *fallibilist* evidentialists reject (4) while retaining some combination of (1)-(3). In other words, they hold a conjunction of claims, the first conjunct of which is some combination of (1)-(3), and the second conjunct of which is

- (5) $\exists c \exists p \exists n$: In c Ksp, and S’s total evidence supports p to at most n , and $n <$ entailment

It is precisely one (any one) of these three possible conjunctions that I am suggesting evidentialists retreat to, when I suggest that they retreat from the “sufficiency thesis” to the “possibility thesis”.²¹ I claimed that said retreat is contrary to the spirit of fallibilist evidentialism because it weakens the role of fallible evidence in the theory of knowledge. Now that we have a precise statement of the thing retreated-to, we’ve finished the first step in arguing for this claim. The second step is to get a precise statement of the thing retreated-from, that is to say a precise statement of the “sufficiency thesis”.

Again, close inspection reveals several versions of this thesis. One is that there is a nonentailing strength of evidence, which is always enough to know anything. Another is that there is always a nonentailing strength of evidence, which is enough to know anything. Unlike the first claim, this second claim allows the relevant strength to differ across cases. An even weaker claim allows that strength to differ across propositions in the *same* case. Letting “NEC” mean “the nonevidential conditions on knowledge” and “e” mean “entailment”, these three claims are:

- (6) $\exists n < e \forall c \forall p$: If in c , NEC hold and S’s evidence supports p to at least n , then Ksp

- (7) $\forall c \exists n < e \forall p$: If in c , NEC hold and S 's evidence supports p to at least n , then Ksp
 (8) $\forall c \forall p \exists n < e$: If in c , NEC hold and S 's evidence supports p to at least n , then Ksp

The argument from unknowability shows that in some cases, there are propositions for which *no* amount of nonentailing evidence yields knowledge when combined with the nonevidential conditions. This refutes (8); and so since (6) and (7) entail (8), it refutes them too.²²

It also refutes another view having some *prima facie* claim to being the “sufficiency thesis”. Roughly, this view is that some level of evidential support is “adequate” for knowledge but sometimes supports false beliefs.²³ Cleaning that idea up, we get

- (9) $\forall c \forall p \exists n$:
 (If in c , NEC hold and S 's evidence supports p to at least n , then Ksp)
 and
 ($\exists c \exists p$: in c , S 's evidence supports p to exactly n , and p is false)

Which is refuted by the argument from unknowability because, given plausible background assumptions, it entails (8).²⁴

It is hard to see what the “sufficiency thesis” might amount to, if not one of (6)-(9). But each of those claims is refuted by the argument from unknowability. Hence the sufficiency thesis is false no matter what it amounts to, false on every way we can make it precise. Hence it is *false*.²⁵ Hence fallibilist evidentialists should retreat away from it: they should reject (6)-(9).

Bringing this all together, we get an important payoff. In particular, we get a unified and precise understanding of the various views and arguments on the table and the relationships among them. *All evidentialists* accept some combination of (1)-(3). *Fallibilist* evidentialists add (5). *Sufficiency* fallibilist evidentialists further add some combination of (6)-(9). The argument from unknowability shows that they should recant this last addition. But they can still retain both evidentialism and fallibilism, sticking to their favored combination of (1)-(3) with (5) sans (6)-(9). That is what I am suggesting they do, when I suggest that they retreat from the sufficiency thesis to the possibility thesis.²⁶

Now that I've made this suggestion precise, I'll explain why implementing it would weaken the role of fallible evidence in the theory of knowledge. The explanation analogizes epistemology to another normative domain, political philosophy. Suppose an egalitarian started with the view that adequate equality was necessary for justice, and then identified “adequate” equality with some nonmaximal amount of equality, and then added the point that some non-egalitarian conditions are required for justice too, for instance the lack of rights-violations. Suppose finally that she then claimed that the combination of these various conditions is *sufficient* for justice.

Now suppose that on reflection she decided to drop the final claim, replacing it with the claim that no nonmaximal amount of equality always yields justice when combined with

the nonegalitarian conditions. *Doing that would weaken the role of nonmaximal equality in her theory of justice.* And if she pointed out that there are still cases of justice with nonmaximal equality – well, the role she gives to nonmaximal equality would still seem weakened.

Similarly with the fallibilist evidentialist. Suppose such a theorist started with the idea that adequate evidence is necessary for knowledge, and then identified “adequate” evidence with evidence of some nonentailing strength, and then added the point that there are other necessary conditions on knowledge too, for instance truth. Suppose finally that she then claimed that the combination of these various conditions is *sufficient* for knowledge.

Now suppose that on reflection she decided to drop the final claim, retreating to the claim that no nonentailing degree of evidential support always yields knowledge when combined with the nonevidential conditions. Just as the similar retreat by the egalitarian would weaken the role of nonmaximal equality in the theory of justice, this retreat by the evidentialist would weaken the role of fallible evidence in the theory of knowledge. And similarly again, if our evidentialist pointed out that there are still cases of knowledge with nonentailing evidence – well, the role she gives to fallible evidence would still seem weakened.

We can see through analogies to other domains, then, that the retreat from the sufficiency thesis to the possibility thesis weakens the role of fallible evidence in the theory of knowledge. And the argument from unknowability brings fallibilist evidentialists to make this very retreat. Thus it happens that the argument from unknowability brings fallibilist evidentialists to begin severing the connection between knowledge and fallible evidence. This disconnect, this weakening of the role of fallible evidence in the theory of knowledge, is one wage of unknowability. (I hasten to add that similar wages may well be due from other forms of fallibilism, for example reliabilist and coherentist fallibilism, once those views are worked out in sufficient detail to confront the argument from unknowability directly.)²⁷

-
- Alston, William. 1988. “An Internalist Externalism.” *Synthese* 74, 265-283
- Comesaña, Juan. 2006. “A Well-Founded Solution to the Generality Problem.” *Phil Studies* 27-47.
- Conee, Earl. 2001. “Heeding Misleading Evidence”. *Philosophical Studies* 103: 99-120. Page references are to the reprint in Conee and Feldman 2004.
- Conee, Earl and Richard Feldman. 2004. *Evidentialism*. Oxford University Press.
- Feldman, Richard. 2003. *Epistemology*. Prentice-Hall.
- , 1988. “Having Evidence”. In David Austen (ed), *Philosophical Analysis* (Kluwer), pp. 83-104. Reprinted with an Afterword in Conee and Feldman 2004.
- Feldman, Richard and Earl Conee. 2004. “Making Sense of Skepticism”. In Conee and Feldman 2004.
- Goldman, Alvin. 1979. “What is Justified Belief?”. In G. Pappas (ed), *Justification and Knowledge* (D. Reidel). Page references are to the reprint in Goldman’s *Liaisons: Philosophy Meets the Cognitive and Social Sciences* (MIT Press 1992).
- , 2009. “Williamson on Knowledge and Evidence”. In D. Pritchard (ed), *Williamson on Knowledge*. Oxford University Press.

- Huemer, Michael. 2007. "Moore's Paradox and the Norm of Belief" in *Themes from G. E. Moore*, ed. Susana Nuccetelli and Gary Seay. Oxford University Press.
- Kelly, Thomas. 2008. "Evidence: Fundamental Concepts and the Phenomenal Conception". *Philosophy Compass*.
- Klein, Peter. 1981. *Certainty: A Refutation of Scepticism*. University of Minnesota Press.
- Lewis, David. 1979. "Attitudes *De Dicto* and *De Se*". *Philosophical Review* 88: 513-543.
- Neta, Ram. 2008. "What Evidence Do You Have?". *British Journal for the Philosophy of Science*. 59: 89-119.
- Salerno, Joe. 2008. "Knowability Noir: 1945-1963". In Joe Salerno (ed), *New Essays on the Knowability Paradox*. Oxford University Press.
- Sutton, Jonathan. 2005. "Stick to What You Know". *Nous* 39: 359-396
- . 2007. *Without Justification*. MIT Press.
- Unger, Peter. 1975. *Ignorance*. Oxford University Press.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.

¹ An anonymous referee advised me to point out that this reductio argument does *not* show that we can't believe propositions of the form <P and I don't know P>. Nor, for that matter, does it show that we can't believe propositions of the form <P and I don't believe P>. The same referee further pointed out an interesting principle in the ballpark here, namely the principle that if I know that you know that P, then I know that P. If this principle is true, then whenever Q is unknowable for me, so too is the proposition that you know Q. For example, since I can't know *that the president is dead and I don't know the president is dead*, I also cannot know that you know that proposition.

² The literature on unknowability traces back to Alonzo Church, who (as an anonymous referee) suggested it to Frederic Fitch, with whom it is most associated. See Salerno (2008) and the other essays in the same volume.

³ You might worry that consistently with everything I've described, I can come to know D at some point in time after the oracle asserts it. After all, she just utters "P and *you don't* know it". She doesn't utter "P and *you'll never* know it". If you have this worry, then throughout the current paper, substitute "P and you'll never know it" for "P and you don't know it". I've left this nicety aside to simplify the text. Another nicety I am largely leaving aside is the fact that unknowability is agent-relative. What is unknowable for me is not identical to what is unknowable for you. For instance, *I can't* know that *p and I don't know it*, but *you can* know that proposition. Similarly, *you can't* know that *p and you don't know it*, but *I can* know that proposition. It would be theoretically perspicuous, then, to drop the term "unknowable" and replace it with the term "S-unknowable". But, for the most part, I will leave this nicety aside.

⁴ As an anonymous referee pointed out, there are versions of fallibilist evidentialism which take the upper scale of evidential support not as entailment, but as "conclusiveness". Plausibly, though, the argument from unknowability refutes those forms of fallibilist evidentialism as well. For just as the oracle can give us arbitrarily strong (short of entailing) evidence, so too can she give us arbitrarily strong (short of conclusive) evidence. Indeed, for *any* degree of evidential strength, she can give us arbitrarily strong (short of that degree) strength.

⁵ Conee and Feldman (2004: 295).

⁶ Conee and Feldman (2004: 296).

⁷ Conee (2001: 270).

⁸ Williamson (2000: 184-230), a non-skeptic, argues that knowledge requires adequate evidence and that adequate evidence is entailing evidence. Unger (1975: 197-226), a skeptic (or a skeptic *at the time* anyway), argues for the same claims.

⁹ An aside. It seems clear enough that I should believe what my friend says, namely that the president is dead but not assassinated. Or at least, that seems clear enough given suitable further specifications about the case (for example specifications that my friend is not mentally unstable or a serial prankster or anything like that). Does it follow that I should *also* believe what the *oracle* says, namely that the president is dead and I don't know it? It seems wrong that I should believe what the oracle says, despite seeming right that I should believe what my friend says. But how could this be so, given that I have *stronger* evidence from the oracle than from my friend? These issues about what we should believe are subtle and important. Without going much into it here, I'll just say two things about how it could be true that I should believe what my friend says, while I should not believe what the oracle says, even though I get stronger evidence from the oracle than from my friend. First: when we get evidence for an unknowable proposition, we should perhaps *raise our subjective probability* in that proposition, but we should not *believe* that proposition. Second: the reason we should not *believe* unknowable propositions is that *knowledge is the norm of belief*. My purpose in the current paper is not to defend these unorthodox claims. I mention them only to illustrate some of the applications our oracle cases have to epistemology more generally, beyond their application to fallibilist theories of the nature of knowledge. On the view that knowledge is the norm of belief see Williamson (2000: 11, 47, 208, 255-256), Sutton (2005, 2007), and Huemer (2007).

¹⁰ By “justified to believe *p*” I mean “propositionally justified to believe *p*”, a status one can have even if one does not believe *p*. Contrast this status with the state “doxastically justified belief that *p*”, a state one can be in only if one believes *p*.

¹¹ Of course, in denying that I am justified to believe P&Q whenever I am justified to believe P and justified to believe Q, I am denying multi-premise closure for propositional justification. More precisely, I am denying that propositional justification is closed under conjunction-introduction. So much the worse, then, for multi-premise-propositional-justification-closure.

¹² It should go without saying that this argument only applies to evidentialists who are in the business of giving noncircular definitions of knowledge. Evidentialists who aren't in that business can ignore it, and they can also ignore some of the similar arguments I'll make later in the paper.

¹³ Compare this objection to the remarks in footnote 5 of Sutton (2005). Thanks to <names deleted> for helpful discussion here.

¹⁴ This is the standard conception of “defeaters”; see e.g. Klein (1981: 137-141).

¹⁵ Maybe the bullet-biting response is a nonstarter with the lottery too. But in any case, that response is *more of a nonstarter* with the oracle than with the lottery.

¹⁶ Of course, the suggestion that this is the relevant process is in tension with Goldman's (1979: 116) conjecture that belief-forming processes are “content-neutral”. But we can leave this point aside.

¹⁷ An issue which is of course a standing problem for reliabilism – the “generality problem”. It is worth noting that standard attempts to solve this problem *do not* on their faces seem to render fallibilist reliabilism immune to the unknowability argument. For example, the argument seems to apply to fallibilist versions of Alston's (1988) “indicator reliabilism” and Comesaña's (2006) “well-founded reliabilism”.

¹⁸ Here I assume that coherentism is not a species of evidentialism, but rather its own different beast. Sometimes folks suggest that coherentism *is* a species of evidentialism; see Feldman (2003: 39-69).

¹⁹ Thanks to Dan Howard-Snyder for driving this point home to me.

²⁰ This is the standard way to understand “cases”; see Lewis (1979) and Williamson (2000: 52).

²¹ Why are there only three possible conjunctions here? Because I am assuming that every combination of claims is closed under consequence, and (1) is logically stronger than (2) which is logically stronger than (3). The only conjunctions of (5) with some combination of (1)-(3) are therefore: 5&1&2&3, 5&2&3, and 5&3. Later in the discussion as well, when I talk about other numbered claims up through (13), I will continue to assume that all combinations of claims are closed under consequence.

²² As astute readers will have noticed, (6) is the view I identified with fallibilist evidentialism in section 0.

²³ For an example of fallibilist evidentialists asserting such a claim, see our first quote from Feldman and Conee.

²⁴ In particular, (9) entails (8) given the assumption that evidence is factive – that if e is part of one’s evidence, e is true. For suppose we make that assumption. Then (9) tells us two things: first, that (for any case and proposition) there is a degree n to which (in some case) some truths support some falsehood; and second, that whenever anyone’s total evidence supports a proposition to *at least* that degree and the nonevidential conditions also hold, she knows that proposition. The relevant degree of support must always be less than entailment; otherwise it could never hold between some truths and a falsehood (here I assume the S5 principle that every world is accessible from every world – or, more exactly, the case-theoretic analogue of that world-theoretic principle). This shows that (9) entails

$$(10) \forall c \forall p \exists n < e:$$

(If in c , NEC hold and S ’s evidence supports p to at least n , then Ksp)

and

($\exists c \exists p$: in c , S ’s evidence supports p to exactly n , and p is false)

Which entails (8). Given the assumption that evidence is factive, then, we can show that (9) entails (8) because it entails (10) which entails (8). Obviously, I am opening up a big can of worms about the nature of evidence, when I claim that evidence is plausibly factive. For relevant work see Feldman (1988), Williamson (2000: 184-208), Goldman (2009), Neta (2008), and Kelly (2008).

²⁵ If these last sentences smack of supervaluationism about vagueness, ignore that. I’m not trying to be a supervaluationist. I’m just saying that whatever claim the sufficiency thesis amounts to, it is a false claim.

²⁶ There is one other thing fallibilist evidentialists might do instead of moving to the possibility thesis. In particular, they might restrict their theories of knowledge to *knowable propositions*. Their view would then amount to some combination of the following claims, listed in order of decreasing logical strength and analogous to (6)-(8):

(11) $\exists n < e \forall c \forall p$: If p knowable for S in c , then
if NEC hold and S ’s evidence supports p to at least n , then Ksp

(12) $\forall c \exists n < e \forall p$: If p knowable for S in c , then
if NEC hold and S ’s evidence supports p to at least n , then Ksp

(13) $\forall c \forall p \exists n < e$: If p knowable for S in c , then
if NEC hold and S ’s evidence supports p to at least n , then Ksp

The retreat to these knowability-restricted theories of knowledge is unsatisfying, and here’s why. The argument from unknowability gives us a series of cases showing fallibilist evidentialism to be false. The retreat simply stops theorizing about the problematic cases, and applies the old theory to everything else. This is a textbook example of *ad hockery*. Thanks to <name deleted> for helpful discussion here.

²⁷ Thanks to Frank Arntzenius, Chris Bryant, Andrew Cullison, Richard Feldman, Mikkel Gerken, Dan Howard-Snyder, Hud Hudson, Ernest Sosa, and Timothy Williamson for helpful comments on ancestors of this paper.