

Advice for Fallibilists: Put Knowledge to Work

Matthew McGrath
University of Missouri
mcgrathma@missouri.edu

§1. Conceptions of Fallibilism

We make mistakes, and not only when we purposely ignore the evidence or coax ourselves into believing what we would like to believe. Our primary sources of information about the world are fallible in the sense that they sometimes lead us to false beliefs. This is true of memory, testimony, inductive reasoning, perception and even introspection. So much is uncontroversial, though perhaps not of obvious epistemological interest. But we can bring these general fallibility considerations closer to our epistemic home.

We sometimes make mistakes even when our evidence appears to be the same as it is now for many of the beliefs we now would call knowledge. I think I know that my wife taught class today. But there have been people with the same sort of evidence I have but whose spouse didn't teach that day, for some reason. They had evidence just as good as mine but they were wrong. Or consider that any teacher knows that her students will believe and take themselves to have knowledge based on her testimony about matters of fact. When I teach my students Chisholm, I usually tell them that I took Chisholm's last seminar at Brown, and that the seminar was on the topic of intentionality. However, if instead told them that the topic of the seminar was epistemology, they'd believe it every bit as much. (I'm not sure anyone at the conference could catch me on this one...) It appears, then, that for much of what we take ourselves to know, our evidence is no stronger than the evidence on the basis of which people form false beliefs.

We can go further, as any reader of the *Meditations* knows. At least there is some initial plausibility to the thought that a dreamer, in her bed, has no weaker evidence than I have now for thinking he's typing at a computer, which of course he's not. So, perhaps general considerations of fallibility touch on even paradigm cases of putative perceptual knowledge. To go to the full extreme, it seems there could be someone with no better evidence than I for many "obvious" claims, such as *I have a body*, *I have friends*, etc. but who is the victim of a massive deception from an evil demon or superscientist. Such a person would think he knows these things but his beliefs are false.

But in what respect is the evidence in the bad case no weaker than the evidence in the putatively good case? One very common answer given is that in both cases the evidence fails logically to exclude all possibilities of error, and indeed that it leaves open the same range of error-possibilities. And this answer is usually backed up by the claim that, if the bad case is chosen correctly, the subject in the bad case has the *very same evidence* as the subject in the good case. If it's the same evidence for the same proposition, then, since the proposition is false in the bad case, the belief in the good case is based on evidence that fails to entail (in the broadly logical sense) the proposition believed.

Postulation of a logical gap explains facts of defeasibility well: I think I know that Plato taught Aristotle, but I recognize that if the *New York Times* published an article, coauthored by a long list of prominent specialists on ancient history, stating that this is false (and providing documentation, etc.), I would rationally give up the belief. And it seems this could happen without eliminating any of my evidence for the belief. The structure of the situation, then, is this: my evidence is *e*, and *e* confirms *p* for me, but were I to have this defeating misinformation, *d*, from the *New York Times*, my evidence would be *e* & *d*, which does not confirm *p* for me. But if *e* entailed *p*, and if entailing evidence confirms what it entails, then this sort of defeat would be impossible.

If all this is right, considerations of actual and possible mistakes indicate that our evidence, for at least some and perhaps all of the things we think we know, fails to guarantee the truth of what we believe. We are left with a logical gap. But is such a gap epistemologically worrisome? Does it, for instance, call our *knowledge* into question? The standard view among epistemologists is that it does not.

Let us establish some terminology before proceeding. Let us say that *fallibilism* is the doctrine that one can fallibly know that something is the case. Now, epistemologists might disagree about how widespread fallible knowing is, some thinking that everything we know we know fallibly and others thinking only that some of what we know we know fallibly. Let us put aside questions about the *scope* of fallible knowing. What we want to get a handle on is what it is to fallibly know something to be the case. Each conception of fallibly knowing will yield a form of fallibilism about knowledge.¹

The slogan that one can know on the basis of non-entailing evidence gives us our first conception of fallibly knowing and so our first version of fallibilism:

To fallibly know that *p* is to know that *p* on the basis of evidence which fails to entail that *p*.

Logical fallibilism about knowledge (LF): knowing that *p* is compatible with having non-entailing evidence for *p*.

If you took a survey of self-described fallibilist epistemologists, asking them to state the doctrine that best encapsulates fallibilism about knowledge, I expect that LF would come out on top. At least it has been widely endorsed.²

But there is reason to think LF doesn't capture all that we want in an account of fallibly knowing. Here are two familiar problems. First, any proposition whatever entails a necessary truth. So there is no logical gap between our evidence and the truth whenever the proposition in question is a necessary truth. Does this mean that all our knowledge of necessary truths is infallible? Of course, Descartes had ways of denying this conclusion, since he thought that God could have made $2 + 2$ not 4. But it would be preferable not to follow him there. Granted, this problem wouldn't show that LF is false,

¹ The term 'fallibilism' originates from Charles Sanders Peirce, but the idea that knowledge can be in some sense fallible is not new with him. Lehrer finds it in Reid, and others claim it can even be found in the Platonic dialogues.

² Proponents are too many to list. They include Peirce, Moore, Chisholm, Feldman (1981), Conee and Feldman (2004), Cohen (1988), and Stanley (2005).

but it would show that the conception of fallibly knowing figuring in LF is not adequate to cover all cases of fallible knowledge.

A second putative problem for LF is that there is at least some case to be made that much perceptual knowledge, memory knowledge, and testimonial knowledge, if it is based on evidence at all, is based in evidence that does entail the truth of what is believed. Thus, disjunctivists (e.g., Snowdon, McDowell) who claim that there is no highest common factor present in both veridical and non-veridical sensory experience will argue that in the good case one knows that there is, say, a table before one on the basis of evidence such as *I see that there is a table before me*, whereas in the bad case one's evidence is merely *I seem to see that there is a table before me*. Timothy Williamson (2000), though he does not endorse disjunctivism about sensory experience, agrees that there is no highest common evidential factor across good and bad cases. The victim of the evil demon has different evidence than you or I. This certainly clears the way for claiming that while you and I often base our belief on entailing evidence our "twins" in the demon case do not. If this is how it is for all of our knowledge, LF is false. And even if LF is retained because there is some knowledge based on non-entailing evidence, one might doubt whether the logical conception of fallibly knowing can do all the work we need done: do we really infallibly know, say, when we base our knowledge that we taught a Tuesday/Thursday schedule in the fall of 2002 on the evidence *that we remember that we did* (and so on entailing evidence)?³⁴

There are ways of replying to these objections. A promising strategy for handling the first is modify the conception of fallibly knowing conservatively (Reed 2002, Hetherington 1999): to fallibly know that p is to know that p on the basis of evidence which is such that by believing p on that basis one could have either failed to have a true belief or failed to have a non-accidentally true belief (that is, one could have been Gettiered). On this conception of fallibly knowing, it is clearly possible to have fallible knowledge of necessary truths: even though any evidence entails a necessary truth, you could get the truth accidentally believing on the basis of your evidence. Since the Gettier condition is associated with of the connection between evidence and truth, this form of fallibilism seems very much in the spirit of the original LF. I'll refer to it as the revised version of the logical account.

The objection inspired by disjunctivism and Williamson's epistemology, however, is not so clearly answered by appealing to the revised version of LF, at least not if Williamson is right about the claim that if p is evidence S has, then p is known to be true. Although it is hard to state an exceptionless closure principle, the following is a good start (Hawthorne 2004): if one knows p, p entails q, and one competently deduces q

³ Neta, too, claims that knowledge is based on entailing evidence. See Neta (forthcoming).

⁴ Some epistemologists might feel uneasy about calling perceptual experiences evidence. If so, they might prefer the following, still broadly logical, form of fallibilism:

To fallibly know that p is to know that p on the basis of a justification the having of which fails to entail that p.

Both of the objections to LF apply to this view, too. First, if p is necessarily true, then any justification one has is such that having it logically excludes not-p. Second, if when one knows, one knows on the basis of entailing evidence, then at least assuming the plausible claim that when e entails p having e entails p, then the this version of logical fallibilism is false, too.

from *p*, knowing *p* all the while, then one knows that *q*. So, suppose I know that *p* on the basis of the entailing evidence *I remember that p*. Is this knowledge fallible according to the revised logical account? If Williamson is right about propositional evidence implying knowledge, then how could I possibly come to believe *p* on basis of a deduction from this same entailing evidence, *q*, without knowing *p*?⁵

Now, the revision to LF might well withstand this sort of objection. But suppose it didn't. Suppose all knowledge is grounded in entailing evidence belief on the basis of which is guaranteed to get us non-accidentally to the truth. It surely wouldn't follow that all our knowledge is infallible! Infallible on this conception, yes, but not on every relevant conception. It seems that what we need is need is an *epistemic* conception of fallibly knowing, a conception which connects knowledge with how justified the proposition is for the subject.

Most epistemologists today agree with Conee and Feldman when they remark that the fact that there is some logical possibility of mistake with the same evidence (or of believing the truth by accident on the basis of the same evidence) is by itself no reason at all to think we are making a mistake and so no reason at all for doubt (2004, 301). But does that mean that we are maximally justified whenever we know, i.e., that knowledge requires maximal or perfect justification? Conee and Feldman admit that it does not. They concede that logical possibilities of error show that our justification isn't perfect (2004, 285-6). Still, they insist that this imperfection doesn't stand in the way of knowledge.

What all this suggests is that however we answer the question whether or not there is a logical gap between evidence and the truth, there remains a question about whether there is some sort of *epistemic* gap. Indeed, it seems to this lay reader that in the first Meditation Descartes was trying to argue precisely for an epistemic gap not a logical one. The relevant gap isn't between our evidence and the truth but between our actual justification and maximal justification.

We will consider two *epistemic* conceptions of fallibly knowing and so of fallibilism. First, a weak epistemic conception:

To fallibly know that *p* is to know that *p* despite *p*'s not being maximally justified for one.

Weak Epistemic fallibilism (Weak EF): knowing that *p* is compatible with *p*'s not being maximally justified for one.

If we call maximal justification epistemic certainty (using the latter as a technical term), then Weak EF allows for knowledge without epistemic certainty. Notice that Weak EF is perfectly consistent with both the acceptance and the denial of LF (and its revised version). This is because its conception of fallibly knowing is compatible *both* with

⁵ Perhaps all that would be needed is the claim that propositional evidence is factive. The disjunctivist or Williamsonian might claim that whenever we know we know on the basis of evidence that does not merely entail the truth of what we believe but also entails that we know it. Thus, *I remember that p*, *I perceive that p*, *I was just informed that p*, etc. all entail that one knows that *p*. And if one's evidence for a known proposition *p* entails that one knows it, then if propositional evidence is true, one *does* know *p* when one has such evidence, and moreover one couldn't have failed to know that *p* basing one's belief on that evidence.

infallibly knowing in these senses *and* with fallibly knowing in those senses. I would venture that not only do most proponents of LF and its revised version accept Weak EF, but so do the critics who lodge the objections we have considered. It is not that these logical conceptions of fallibly knowing are *misconceptions*, but only that they cannot do all the fallibilist work we want done.⁶

I expect that Weak EF is a fallibilism many would be happy with. But there are reasons for going still further and embracing Strong EF:

To fallibly know is to know despite the fact that there is a non-zero epistemic chance, for one, that not-p.

Strong Epistemic Fallibilism (Strong EF): knowing that p is compatible with not-p's having a non-zero epistemic chance for one.

What do I mean by 'epistemic chance'?, you might ask. Depending on how am construing it, it is unclear what the implication relations are between Strong EF and LF or even Strong EF and Weak EF. Indeed, one might worry about the following dilemma. Either we take epistemic chance to measure of the degree to which one's evidence entails a proposition, i.e., as a measure of partial entailment of a proposition by one's evidence, or we take it to measure how justified the proposition is for you. If we read in the former way, then there is some non-zero epistemic chance that not-p just if one's evidence doesn't entail that p. So, on *that* reading of 'epistemic chance', Strong EF boils down to LF. (And hardly deserves to be called *Strong* EF, because it, being equivalent to LF, would be compatible with both Weak EF and its denial.) If we read 'epistemic chance' instead as a measure of justification, then Strong EF boils down to Weak EF, because there being some non-zero epistemic chance that not-p would amount to there being some gap between the justification p has and maximal justification. So, either way, Strong EF is equivalent to either LF or Weak EF. It adds nothing.

This dilemma assumes I am forced, from the start, into choosing which of the two specified conceptions of epistemic chance we are employing. But I am not forced to make this choice. There is an independently comprehensible notion of epistemic chance, distinct from both the concepts of partial entailment by the evidence and strength of justification. One might argue that the concepts are extensionally equivalent to one of the two above, but one cannot assume this from the start. This will be clearer after we consider several arguments for Strong EF. I now turn to those arguments.

As long as one keeps one's focus directly on paradigm instances of perceptual knowledge, one might be happy denying that knowledge is compatible with an epistemic chance of error. Is there really a chance for me that I'm not typing on something that at least looks like a computer? But when we turn to other paradigm instances of knowledge – I know that I was born in the United States, that my mother is over 60 years old, that Plato taught Aristotle – it is harder to maintain this resistance. I think there are two reasons why.

⁶ I am certainly not the first to specify an epistemic conception of fallible knowledge. One finds it in Feldman (2003, 122), at least suggested in Hetherington (2005), and elsewhere. I have found, though, that epistemologists often write as if fallibly knowing in this weak epistemic sense is obviously equivalent to fallibly knowing in the logical sense.

First, it can seem that such beliefs are lotteryizable.⁷ That is, the falsity of the proposition believed can be viewed as “winning a lottery,” where the other tickets are beliefs, whether one’s own or others’, backed with roughly the same evidence, and where one has good reason to think that some of these beliefs are false. If beliefs like I was born in the US are lotteryizable, it seems hard to resist the conclusion that there is some small chance that they are false, that they “win” the lottery. And if that’s right, then the price of claiming that knowledge requires no chance of error is a fairly strong form of skepticism. Perhaps some perceptual knowledge can be saved, and maybe some basic logical knowledge, but we’d have to concede that we don’t know much of we ordinarily think and claim we know.

Notice that the notion of epistemic chance employed here can’t be equivalent to that of partial entailment by the evidence. One’s belief might, as it happens, be in a necessary truth. Some such beliefs are surely lotteryizable if any are. But they are (wholly) entailed by any evidence. Matters aren’t as clear when we ask whether the notion of epistemic chance we’re working with is a measure of strength of justification. The issue is whether justification can improve even once it leaves no epistemic chance of error. Perhaps there are justified beliefs that aren’t lotteryizable and yet still fall short of maximal justification. We can make some progress on this question after considering a second argument for Strong EF.

Here is the second argument. Would you stake your life on the proposition that Plato taught Aristotle? Or if you don’t like gambles, do you think it would be reasonable for someone who was just as well justified as you are in thinking Plato taught Aristotle and who was not risk-averse, to stake her life on this proposition? It seems to me it would not. Perhaps it would be different for Plato himself when he was meeting with Aristotle for a lesson. Perhaps during the lesson Plato should be reasonable to accept such gambles. But our justification is not as strong as his.

What do these reflections have to do with epistemic chance? Here we draw on the connections between epistemic chance and gambles. Rational choice between gambles arguably depends, not on subjective degree of belief, but how genuinely epistemically likely the relevant propositions are for the subject, i.e., on epistemic chance. If all a person’s evidence points to the truth of *p*, but the person, irrationally, believes that *p* is very unlikely, then action on the basis of this belief inherits the irrationality.

Given the connection between epistemic chance and rational choice between gambles, it becomes more plausible to think that what I’ve called Strong EF really is stronger than Weak EF. It is obvious that Strong EF implies Weak EF ((justification which leaves a chance of error is not maximal justification). And there is a good case to be made that the converse implication fails. Perhaps there are or could be some beliefs which are so well justified that it would be rational to gamble on their truth no matter what the stakes. Still, this is compatible with saying that there is a room for improvement in one’s justification. Perhaps if one had a substantially better answer to a skeptic than Moore does, one’s justification even for “here is a hand” would be improved, but it is hard for me to believe that my epistemic chance for this proposition is anything other than 1.

⁷ This idea of lotteryizability is derived from Vogel 1990.

One final preliminary objection to Strong EF should be considered, before we move on to the main difficulties for the view. Philosophers versed in epistemic logic might be frustrated by my taking Strong EF seriously. Isn't epistemic chance, along with all epistemic modals – it's possible that p, it could be that p, it might be that p – simply *defined* in terms of knowledge in such a way that Strong EF must be false? On a standard definition, these are defined as true (relative to S) iff p is consistent with what S knows. If this is right, then clearly Strong EF is stronger than Weak EF, since the latter is pretty clearly true and the former is contradictory! Now, I do not deny that for certain purposes, e.g., in doing epistemic logic or doing formal semantics, it is fine to stipulate certain meanings or truth-conditions for epistemic modals. But no stipulative definition can resolve the question of whether the notion of epistemic chance, as it figures in our thought about gambles and lotteries, is definable in terms of knowledge. Compare the following consideration. On standard assumptions in epistemic logic, 'Kp' is true at a world w iff 'p' is true at every world compatible with the subject's knowledge. Plug in any logical truth for p and 'Kp' comes out true at any world. But this hardly shows that we all know all logical truths.

In light of the arguments from lotterizability and gambles, I think Strong EF might very well be what it takes to maintain, against the skeptic, that the scope of our knowledge is roughly as wide as we think it is. It looks awfully plausible that a vast swath of what we take ourselves to know does not have probability 1 for us, i.e., that our justification leaves some small but non-zero epistemic chance of error.⁸

§2. The Madness of Fallibilism, at least of Strong Epistemic Fallibilism

Few epistemologists explicitly endorse Strong EF, hereafter strong fallibilism. And there is a reason for this. It's *mad*. Here is Lewis:

If you are a contented fallibilist, I implore you to be honest, be naïve, hear it afresh. 'He knows, yet he has not eliminated all possibilities of error.' Even if you've numbed your ears, doesn't this over, explicit fallibilism *still* sound wrong? (Lewis 1999, 419-20)

Now, talk of "eliminating" possibilities of error might seem too close to knowing them to be false, in which case this clashing conjunction is a denial of a principle of epistemic closure. But closure is not the real issue. If strong fallibilism is true, the following clashing conjunctions should be true in many cases:

"I know that p but there is a chance that not-p"
 "I know that p but it is possible (it might/could be) that not-p"

⁸ Others who seem, if in some cases hesitantly, to endorse Strong EF include Klein (1980, 197), Rysiew and Dougherty (forthcoming), Feldman (2003, 122) and Lehrer (1990, 177-8). Opponents are easier to find: Bach (forthcoming), DeRose (1991), Hawthorne (2004), Huemer (forthcoming), Stanley (2005), Williamson (2000), along with just about everyone working on epistemic modals in the philosophy of language.

Don't these just sound wrong, at least when one is careful to read both conjuncts as simultaneously endorsed, rather than reading the second conjunct as a correction to the first.

Appropriate third-person versions of these statements seem wrong, too:

“She knows that p but there is a chance, for her, that not-p”

“She knows that p but it's possible, from her perspective, that not-p”

If these sorts of statements are often true, as the proponent of strong fallibilism must say, why do they sound so wrong?

Consider, also, dialogues that seem to support the claim that “S knows that p” entails “There is no epistemic chance, for S, that not-p”:

Attorney for the Defense: Is there a chance that the man sitting here in the courtroom today is not the man you saw that night?

Witness: I know he's the guy. So, no, there is no chance.

Here the witness' knowledge claim answers the defense's question in the negative, which is just what we would expect if the entailment held. A second dialogue (cf. Hawthorne 2004) tests the contrapositive entailment from “There is an epistemic chance, for S, that not-p” to “S doesn't know that p”⁹:

Witness: Ok, I admit, there's a chance that the man sitting there isn't the guy I saw that night.

Defense: Ah, so you don't know this is the man you saw, do you?

Witness: No.

These dialogues give us reason to think that the clashing conjunctions above clash because they express genuinely contradictory propositions. And if this is right, and right no matter what the content of the claims in question, then strong fallibilism is false.¹⁰

Before we consider how our strong fallibilist might respond, let us get a second set of clashing conjunctions out on the table, the “Moorean” clashes:

p but there is a chance that not-p

p but it is possible (it might/could be) that p

As if these aren't bad enough, they have the following obvious entailments:

p is a proposition which, though true, has a chance of being false.

p is a proposition which, though true, might be false.

⁹ On standard conceptions of entailment, if P entails Q, then $\sim Q$ entails $\sim P$.

¹⁰ Other tests for entailment seem to support the claim, as well. The conditionals “If I know that p, there's no chance that not-p” and “If there a chance that not p, then I don't know that p” seem plausible, which is what is predicted by the entailment thesis.

These certainly sound wrong. But if strong fallibilism is true, these statements should express propositions which we can and should believe to be true, indeed we can *know* them to be true. But, as Michael Huemer (forthcoming) has recently argued, the claim that we can know such things seems obviously false.

When we turn to dialogues to test that the claim that “p” entails “the epistemic chance, for S, that not-p is zero,” we get a very different result than we got before. While the first dialogue is fine,

Defense: Is there a chance, do you think, that the man you saw that night isn't the man sitting here?

Witness: He's the man. So, no chance.

the second one, which tests the contrapositive entailment, is unsuccessful:

Witness: There's a chance that this is not the man I saw that night.

Defense: So, this isn't the man you saw last night, then.

Witness: Wait, I didn't say that!

This is decisive reason to reject the claim of entailment. We should expect as much, given that the other Moorean clashes (*p* but I don't believe that *p*, *p* but I don't know that *p*, etc.) do not express contradictory propositions.

This should not console the strong fallibilist. For, just as one cannot properly assert, nor rationally believe, let alone know, the propositions one would express by uttering the standard Moorean paradoxical sentences, “*p* but I don't believe that *p*” or “*p* but I don't know that *p*,” the same goes for “*p* but there is a chance that not-*p*”. And yet, if strong fallibilism is true, it is hard to see *why* this would be. I can know *p*, despite the fact that there is a non-zero chance that not-*p*. And why couldn't I, strong fallibilist that I am, also *know* that there is a non-zero chance that not-*p*? I might reflect on facts about gambles and lotteryability, it seems, to come to know this. Then I should be able to come to know the conjunction that *p* and there's a chance that not-*p*, shouldn't I?^{11 12}

So the strong epistemic fallibilist has to do several things. She has to explain why the first group of clashing conjunctions seem wrong, even though they can and presumably often are true on her view, and she must do this while also making sense of why the test dialogues seem to support the claim that there is a genuine entailment from “S knows that *p*” to “there is no chance, for S, that not-*p*.” In the case of the second group, the Moorean clashes, she must explain why these seem not only unassertable but even rationally unbelievable and unknowable (by the speaker or thinker) even though if strong fallibilism is true they should be very easy to know..

¹¹ Notice, no similar problem arises with the explanation of why one can't rationally believe or know the truth of *p* but I don't know that *p*. One standard line is that in rationally believing *I don't know that p* one has a defeater which prevents one from rationally believing *p*.

¹² Could the fallibilist say that knowledge is compatible with having a chance of error but that the recognition of this chance strips one of that knowledge? This suggestion is highly implausible. Knowledge becomes very elusive indeed: all you need to do is think about fallibilism as it applies to your beliefs and you lose knowledge left and right.

Lewis is no doubt right that, between fallibilism and skepticism, fallibilism is the “less intrusive madness” (419). But one would like, of course, to endorse fallibilism without being mad. The question is how to do this.

§3. How to be a Fallibilist

I will discuss three proposals for handling the troubling data.

Proposal #1: Scope confusions.

Could some sort of scope confusion, together with the factivity of knowledge, help to explain the clashes? No. While there is at least some plausibility to the thought that we find “if S knows then S can’t be wrong” compelling because we read the ‘can’t’ as having wide scope (see Feldman 2003, 124), there is no plausibility to the claim that ‘there is a chance that’ or ‘it is possible that’ take wide scope in any of the statements above. Scope-confusion isn’t the explanation. It’s much harder to get the epistemic modal to assume wide scope over a conjunction than over the antecedent of a preceding conditional. It is a *misreading* of “p but it might be that not p” to read it as meaning that it might be that: p but not-p. And in any case this would wrongly diagnose Moorean clashes as genuinely contradictory, which they are not.

Proposal #2: Contextualism

Of course, Stewart Cohen told us to be a fallibilists (Cohen 1988): accept contextualism. The idea might be that as soon as we mention that there is a possibility of not-p we drive up the standards for knowledge so high that only absolute certainty (probability 1) is enough to count as knowing. This explanation is more promising than the postulation of scope-confusion, but not particularly plausible. It is somewhat plausible to claim that when specific error possibilities are mentioned or taken seriously, possibilities which are consistent with and at least putatively explanatory of one’s apparent evidence, the standards for knowing rise so that to count as knowing one must be able to “eliminate” those possibilities (cf. Cohen 1999). So, if I mention the possibility that your car has been stolen and driven away, this might lead you to retract a knowledge-claim about the location of your car, and perhaps the retraction is correct, in the knowledge-claim is now false. However, the bare mentioning of “it might be that not-p” doesn’t conjure up any remotely specific possibilities of error. The well-informed reader might recall that Lewis himself insists on reading the noun ‘possibility’ in a certain rather technical way which ensures it has a certain degree of specificity (Lewis 1999, 423).¹³ But in the ordinary English statement “I know that p but it’s possible that not-p,” this technical reading is not in force.

Proposal #3. Pragmatic Infelicity

¹³ “A possibility will be specific enough if it cannot be split into subcases in such a way that anything we have said about possibilities, or anything we are going to say before we are done, applies to some subcases and not to others.”

The most plausible explanation, in my view, is to appeal to pragmatic infelicity. But it must be of a fairly deep kind. Again, there seems to be something wrong with believing them, with committing oneself to their truth.

What's going on, then?

We can rule out from the start the claim that these statements aren't assertable because they are so obviously true. Far from it, they seem positively wrong not trivially right.

Patrick Rysiew and Trent Dougherty (forthcoming) suggest a better pragmatic explanation, at least for the first set of clashes. The reason "I know that p but it's possible that not-p" seems wrong is that uses of 'it's possible that not-p' standardly pragmatically impart that there is a significant possibility or chance that not-p and not merely the sort that accompanies all fallible knowledge. Why should this explain the oddity? We are told: "a significant chance of error may well prevent one from knowing." [forthcoming, 9]. I will ignore the "may well" hedge, because if it's read too weakly the resulting account can't explain what it aims to explain.

The account I am interested in, then, has two parts:

Part 1: Assertions of 'it is possible that p' and the like pragmatically impart but do not require for their truth that not-p is a significant possibility.

Part 2: If not-p is a significant possibility for a person then the person doesn't know it.

What's wrong with asserting, "I know that p but it's possible that not-p," is that in one breath one is saying that one knows and in the next one pragmatically imparts that one doesn't know.

The dialogues are explicable because normally when one asks "is there a chance that not-p?" one is getting one's audience to focus on the question of whether there is a significant chance that not-p. The hearer can then answer this question by saying "No, I know that p." Similarly, in uttering, "yes, there is a chance that not-p," the speaker is pragmatically imparting that there is a significant chance that not-p. No wonder, then, that the hearer can reply, "so you don't know that p."

Rysiew and Dougherty go on to say that if one factors out this normal pragmatic implication the resulting statements don't sound so problematic. Thus,

"Of course there is always some chance that one is wrong, *anything is possible*, but we know that p."

"The possibility that not-p is ridiculous and not worth considering. We know that p."

do not seem so clearly wrong. I wouldn't say they seem clearly right either. But here is where the hard work in epistemology comes in. The lotteryizability and gambles argument could be used to defend the claim that such statements, at least in many cases, are true.

So far, so good. But what about the Moorean clashes? Here the Rysiew/Dougherty account needs supplementation. A natural thought is to claim that in

asserting that *p* one is “representing oneself” as knowing that *p*, or at least one is pragmatically imparting that any chance that not-*p* is so small as to be insignificant. If either of these proposals is correct, we can see why there is a clash. Although one is not stating a contradictory proposition, one is pragmatically imparting *both* that one knows and that one doesn’t, or perhaps alternatively *both* that there is no significant *chance* that not-*p* and that there is a significant chance that not-*p*.

Can the propositions expressed by Moorean clashes be reasonably believed and even known? Yes. They seem not to be because we turn our minds to the wrong proposition: *p but there is a significant chance, for me, that not-p*. That sort of proposition *can’t* be reasonably believed (for the same reason Moorean propositions can’t), though it can be true. But the proposition expressed by Moorean clashes can be reasonably believed and known because the speaker can reasonably believe that something is the case and also that there is an *insignificant* chance of that it isn’t the case, and an insignificant chance is still a chance. Keep separate the proposition expressed and the proposition that would normally be pragmatically imparted, and the problem seems to go away. (Again, though, I wouldn’t say that knowing or rationally believing these things is easy. One needs to know or rationally believe strong fallibilism!)

Finally, the pragmatic theory nicely explains the successful dialogue we considered before in which the defense asks whether there is a chance that the man the witness saw isn’t the defendant and the witness answers by saying “No, he’s the man.” Why is this successful if there is no entailment? A plausible story is that the question the defense aims to get the witness to answer is whether there is a significant chance of this. The witness can answer *that* question by saying, “He is the man,” because asserting this pragmatically imparts that any chance there is *isn’t* significant.

All this sounds promising. But there is just one remaining task: we need to give an account, or at least say *something*, about what makes an epistemic chance that not-*p* significant for a person.

One sort of account ties significance to the absolute probability of the proposition in question.¹⁴ So, if a proposition is exceedingly improbable, then that is enough to make it insignificant, but if a proposition is, say, fairly probable, it is significant. No doubt one would have to say that the threshold involved is vague.

This account will not do. In some cases, where much is riding on how one acts, depending on whether *p* is true or not, even a small chance of error must be taken seriously. In such situations, people will be prepared to say “Although it is very unlikely, it might be that *p*.” This use of ‘it might be that *p*’ does not pragmatically imply that there is a large chance of error, since it is well-known to all involved that there is only a very small chance. And notice that the successful dialogues we considered can and do take place even when it is clear to all parties involved that there is at best a very small chance that the relevant proposition – that the defendant isn’t the man the witness saw – is true. (The defendant is known to have been in the vicinity, etc.) And, of course, both sets of clashing conjunctions clash just as much when it is clear that the chance of error is very small. The clash can only be mitigated by making clear that the chance of error is not worth taking seriously.

This suggests a second account of significance. The chance that not-*p* is significant just in case it is not *idle*. Whether or not idle hands are the devil’s tools, those

¹⁴ Rysiew and Dougherty suggested this response in correspondence.

who are idle are not working. Perhaps the chance that not-p is significant just in case it is high enough to make it improper to *put p to work*.

What our long discussion of fallibilism suggests, then, is that one particularly plausible way to be a strong fallibilist is to connect knowledge to something broadly *pragmatic*, to something concerning the putting of what is known to work. In the final section of this paper, I want briefly to offer some ideas, developed more extensively in my work with Jeremy Fantl (2002, 2007, and forthcoming), about how to cash out these ideas.

§4. Putting Knowledge to Work.

Knowledge seems to be a basis for belief, for action, for terminating inquiry, for the reactive attitudes, for assertion, etc. What unifies this phenomena? Our thought is this: any time one can speak of a proposition *p* *justifying* someone in ϕ -ing, whether ϕ be an action, a belief, an emotion, an assertion, knowledge that *p* seems to be a basis for ϕ -ing. Knowledge, in other words, seems to be crucially related to *justifying*. This needs explanation. Let's start with justification for belief

Few would deny that propositions, when known, can justify believing other propositions. It does not follow from the mere fact that Brown is in Barcelona that Brown's being in Barcelona justifies you in believing that Brown is in Spain. Nor does this follow from the mere fact that you *believe* – or even believe truly – that he is in Barcelona; your believing that Brown is in Barcelona might itself be unjustified. But suppose you know that Brown is in Barcelona (and, of course, that Barcelona is in Spain). Then that he is in Barcelona justifies you in believing that he is in Spain. Even in cases in which it seems strained to say that what you know is part of your evidence, it seems quite natural to say that what you know can justify you in believing other things.¹⁵

Why not say something similar about action? Why not say that there are facts such that, when you know them, they can justify you in doing things? *That an act has the best consequences of the available options* certainly seems to be like this. But consider a more concrete example. Suppose, walking along the sidewalk with your son, you suddenly see a car careening onto the sidewalk, and on this basis come to know that your child will be hit by the car unless you grab him and pull him back in the next few seconds. It certainly seems that you are justified in grabbing your child and pulling him back! What justifies you? A fact about the relation between your son and the car, viz. *he will be hit by the car if you don't pull him back*. The fact that justifies you is the very fact you know. Notice, as before, that the fact alone would not justify you in grabbing your son, though it may entail that it is best to do so. Just as in the case of belief, were you to quite reasonably have no idea of the existence of this fact, you would be perfectly justified, it seems, in continuing your walk; you would be quite unjustified in grabbing your son. We don't think those who spontaneously grab their children and yank them around, without any inkling that this is a good thing to do, are models of rationality. Nor

¹⁵ Consider, for instance, knowledge about the future. Kucinich knows he won't be inaugurated as president in 2009. It's odd to say that this is part of his evidence (as critics of the K→E thesis point out. But surely *that he won't be inaugurated as president in 2009* justifies Kucinich in believing other things, e.g., *that he won't be legally entitled to pardon anyone in 2009*.

would the fact automatically justify you if you merely believed it – unjustifiedly – even if your belief were true. But if you know the fact, then it does justify you in grabbing your son. This is a perfect mirror of the situation with belief.

To think that knowing is enough in the theoretical case but not in the practical calls into question a basic feature of ordinary reasoning. When trying to determine what is true, we reason from propositions as premises. The same goes for trying to decide what to do. But we don't generally keep the two sorts of reasoning separate. We reason from premises knowing that we might draw all sorts of conclusions from them along the way, some practical and some theoretical.¹⁶ Suppose your sister calls you on the phone to tell you about her plans for her upcoming visit to see you. She tells you, and you thereby come to know, that she'll be arriving at the airport at 8:00am and will need a ride to your place. You might well take this proposition as a premise in your reasoning. At some point in your reasoning, you might draw a practical conclusion from it, e.g., "I'll be there at 8:00 sharp with my car," but you might also draw along the way any number of theoretical conclusions from the premise as well, e.g., that your sister will be tired when she arrives, that you'll not be able to drop the kids off at preschool, and so on. These theoretical conclusions, in turn, will be used as further premises from which other practical and theoretical conclusions are drawn. The bottom line is that we simply don't segregate premises by whether they are available for practical or for theoretical reasoning. But if knowledge that *p* was enough for having *p* as a theoretical reason but not enough for it as a practical reason, we should expect some degree of segregation, if not always, at least when something significant is at stake. But even when the stakes are high there is no segregation.

Fantl and I are not suggesting that there are no important differences between what justifies belief and what justifies action. We are happy to agree that what it takes for *p* to justify you in believing something might be very different from what it takes for *p* to justify you in doing something. Perhaps, justification for action, unlike justification for belief, has an essentially instrumentalist or teleological structure. We are not arguing for an all-purpose "unity" thesis about practical and theoretical justification. What we're arguing for is a unity thesis about the epistemic relation one must bear to a fact to suffice for that fact to justify: knowledge is enough.¹⁷

If you know that *p*, then your epistemic relation to *p* does not stand in the way of *p*'s justifying you in doing, believing, feeling, wanting, liking, hating, or intending anything at all. This will initially seem too strong: if you know that there are flowers on your desk, that there are flowers on your desk can justify you in jumping off your roof? It can justify you in believing that there are *no* flowers on your desk? Yes and yes. It can justify you in both of these things, because its failure to justify you in these things is not due to an epistemic lack with respect to the proposition that there are flowers on your desk. A helpful heuristic, here, is to think about the results of mere improvements of your epistemic position with respect to the proposition in question. Merely bumping up

¹⁶ Here it is not crucial precisely what a practical conclusion is. Whether it is an action, an intention, or plan, or even an "ought" judgment, still, we draw practical conclusions from the same premises from which we draw theoretical conclusions.

¹⁷ We're also not saying that knowledge is the minimal epistemic relation necessary for justifying either action or belief. Perhaps lesser epistemic relations than knowledge might, at times, suffice. But, if a subject knows that *p*, then that guarantees that the subject's epistemic relation to *p* is sufficient for justifying.

your epistemic position with respect to *there are flowers on your desk*, even to epistemic certainty, wouldn't by itself change the fact that *there are flowers on your desk* doesn't justify you in jumping off your roof nor in believing there are no flowers on your desk.¹⁸ And this is obviously because of the lack of any relevant connection between the proposition in question and what it's supposed to justify you in.

If you know something, then if it doesn't justify you in doing or believing something, this isn't because of an epistemic shortcoming in your relation to what is known. Knowing that *p*, as we will say, makes *p* *epistemically eligible* to justify you: if *p* doesn't justify you, then this isn't because your epistemic position with respect to *p* stands in the way.

The Knowledge-Justification principle we're recommending, more precisely, is:

(KJ) If you know that *p*, then *p* is epistemically eligible to justify you in ϕ -ing for any ϕ .

KJ provides a nice account of what makes an epistemic chance of not-*p* insignificant for a person. A chance for you that not-*p* is insignificant for you if it doesn't stand in the way of *p*'s being a justifier for you. And if *p* is a justifier for you, you can properly put it to work. Thus, we think that KJ should be welcome to the fallibilist. It gives her a principled way to accept strong fallibilism while avoiding the sort of madness that Lewis thought the fallibilist is stuck with.

§5 More Madness?

The only trouble is that if KJ and strong fallibilism are true, then it seems knowledge can vary holding fixed all truth-relevant factors in relation what is known. Call the profile of these truth-relevant factors *strength of epistemic position*. Call the thesis that knowledge that *p* can't vary holding fixed strength of epistemic position wrt *p* *purism* about knowledge.¹⁹ If KJ is true, then fallibilist purism is false.

Here is why. Suppose KJ and fallibilist purism are both true. Because fallibilism is true, there is some subject S1 who knows that *p* but for whom there is a non-zero epistemic chance that not-*p*. So, *p* is epistemically eligible to serve as a justifier for S1. Any non-zero chance that not-*p*, however, can make a difference to which gambles on *p* a person can rationally accept in certain possible choice situations. Let a subject, S2, be faced with such a gamble but remain just like S1 with respect to the chance for her that not-*p*, and indeed, we can assume that S2 is just like S1 wrt to strength of epistemic position wrt *p* generally. (We can assume further that S2 is not risk averse.) Now, by KJ, if S2 knows that *p*, then *p* would justify her in accepting the gamble on *p* (say for a small

¹⁸ The reason this is only a heuristic is that, like all proposed counterfactual analyses, it has exceptions!

¹⁹ We discuss purism in more detail in our 2007. The argument I give here for the incompatibility of KJ with fallibilist purism is very similar to the one given in that paper for the incompatibility of a similar principle about knowledge and action and fallibilist purism. The most important change in this paper is the explicit recognition that fallibilism is understood in the strong epistemic way rather than the weak epistemic way.

payoff if p and a very large penalty if not-p).²⁰ But it doesn't, we are assuming. So, by KJ, S2 doesn't know that p. Thus, S1 and S2 have the same strength of epistemic position wrt p, but only S1 knows that p. But this contradicts purism. Thus, KJ entails that fallibilist purism is false.

An illustrative example might be helpful. Suppose you fallibly know that the train in the station will stop in Foxboro. Then there is some non-zero chance for you that the train won't stop in Foxboro. But there is a possible subject, say me, for whom much hinges on whether the train will stop in Foxboro, who has the same strength of epistemic position wrt the proposition (including the same chance that the train won't stop in Foxboro) as you. Assume my stakes are such that I am not justified in boarding. Were there a significantly lesser chance for me that the train won't stop in Foxboro, I would be justified in boarding. My epistemic position wrt *the train will stop in Foxboro* is therefore standing in the way of its justifying me in boarding. Therefore, *the train will stop in Foxboro* is not epistemically eligible to justify me in boarding. By KJ, then, I don't know that the train will stop in Foxboro. But, you and I have the same strength of epistemic position wrt the proposition that the train will stop in Foxboro. This contradicts purism.

So, if KJ is true, then fallibilist purism is false. This shouldn't be too surprising. We were not far from denying fallibilist purism earlier when we claimed that:

- (1), knowledge that p is incompatible with there being a *significant* chance for one that not-p, and
- (2), if there is a non-zero chance that not-p, the significance of this chance that not-p is not simply a matter of not-p's probability, i.e., it can vary holding fixed not-p's probability.

If (1) and (2) are true, then fallibilist purism is false. We can run an argument just like the one we ran above from KJ. Assume fallibilist purism and (1) and (2) are true. Given fallibilism, there is an S1 who knows that p but has a non-zero epistemic chance for not-p. Given (1), S1's chance for not-p is insignificant. Given (2), there should be some S2 with the same probability for not-p that S1 has but for whom the chance that not-p is *significant*. Moreover, we can assume that S2 is alike with all the other relevant truth-related features contributing to strength of epistemic position wrt p. So S1 and S2 have the same strength of epistemic position wrt p, S1 knows that p, but S2's epistemic chance for not-p is significant. But, appealing again to (1), it follows that S2 doesn't know that p. So, S1 and S2 have the same strength of epistemic position wrt p but one knows and the other doesn't. This contradicts purism.

KJ is just a way of putting more flesh on the bones of the rough account of significance found in (1) and (2), and doesn't itself introduce any source of conflict with fallibilist purism.

I realize that it might be thought *mad* to deny purism. This is not the place to address all of the many objections that might be raised against the denial of purism. But the fallibilist who recoils at the thought of denying purism has not only to give up the

²⁰ We assume that the only thing which could stand in the way of p justifying S2 in accepting the gamble is the small chance that not-p. Since the small chance that not-p is a feature of S2's epistemic standing, it follows that p is not epistemically eligible to justify S2 in accepting the gamble.

plausible principle KJ, she must explain away the apparent madness of fallibilism without appealing to a conception of the significance of the chance of error which allows significance to vary holding fixed probability. I can only say that I see little hope down that road. My advice to fallibilists: if you don't want to budge on skepticism, budge on purism.²¹

Works Cited

Bach, Kent (manuscript). "Pragmatics in Epistemology," paper delivered to the University of Missouri Philosophy Department 2007.

Cohen, Stewart (1999). "Contextualism, Skepticism and the Structure of Reasons," *Philosophical Perspectives* Vol. 13, Epistemology: 57 – 89.

_____ (1988). "How to be a Fallibilist," *Philosophical Perspectives*, Vol. 2, Epistemology: 91-123.

Conee, Earl and Richard Feldman (2004a). *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.

Conee, Earl and Richard Feldman (2004b). "Making Sense of Skepticism," in Conee and Feldman 2004a.

DeRose, Keith (1991). "Epistemic Possibilities," *The Philosophical Review* 100: 581–605.

Fantl, Jeremy and Matthew McGrath (manuscript). *Knowledge in an Uncertain World*, under contract with Oxford University Press.

_____ (2007). "Pragmatic Encroachment in Epistemology," *Philosophy and Phenomenological Research* 75 (3): 558–589.

_____ (2002). "Evidence, Pragmatics, and Justification," in *Philosophical Review* 111 (1): 67-94.

²¹ Thanks go to Jeremy Fantl who contributed in obvious ways to the final two sections of the paper. His comments led to many improvements in the first three sections as well. He deserves to be listed as a second author and will whenever I get a chance to go back and replace first-person singulars with first-person plurals. I think it is fair to say that any errors remaining are not mine alone. A talk derived from the main ideas in this paper was delivered to the University of Wisconsin philosophy department. I thank members of that audience, and especially Juan Comesana, for helpful comments.

- Feldman, Richard (1981). "Fallibilism and Knowing that one Knows," *Philosophical Review* 90: 77-93.
- Feldman, Richard (2003). *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hetherington, Stephen (1999). "Knowing Failably," *The Journal of Philosophy* 96(11): 565-87
- Huemer, Michael (forthcoming). "Moore's Paradox and the Norm of Belief," in *Themes from G. E. Moore: New Essays in Epistemology and Ethics*, Susana Nuccetelli and Gary Seay (eds).
- Klein, Peter (1980). *Certainty*. Minneapolis: University of Minnesota Press.
- Lehrer, Keith (1990). *Theory of Knowledge*. Boulder, CO: Westview.
- Lewis, David (1999). "Elusive Knowledge," reprinted in his *Papers in Metaphysics and Epistemology*, Cambridge, UK: Cambridge University Press.
- Neta, Ram (forthcoming). "Defeating the Dogma of Defeasibility," in *Williamson on Knowledge*, Greenough, Patrick and Duncan Pritchard (eds), Oxford: Oxford University Press.
- Reed, Baron (2002). "How to Think about Fallibilism," *Philosophical Studies* 107: 143-57.
- Rysiew, Patrick and Trent Dougherty (forthcoming). "Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions," *Philosophy and Phenomenological Research*.
- Stanley, Jason (2005). "Fallibilism and Concessive Knowledge Attributions," *Analysis* 65.2: 126-31.
- Vogel, J. (1990). "Are There Counterexamples to the Closure Principle?" in *Doubting: Contemporary Perspectives on Skepticism*, M. Roth and G. Ross (eds.), Dordrecht: Kluwer Academic Publishers.
- Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.