

EPISTEMIC PRAGMATISM: AN ARGUMENT AGAINST MODERATION

Juan Comesaña

Abstract: By “epistemic pragmatism” in general I will understand the claim that whether propositions instantiate certain key epistemic properties (such as being known or being justifiably believed) depends not just on factors traditionally recognized as epistemic, but also on pragmatic factors, such as how costly it would be to the subject if the proposition were false. In what follows I consider two varieties of epistemic pragmatism. According to what I shall call *moderate* epistemic pragmatism, how much evidence we need in favor of a proposition in order to know that the proposition is true depends on our preferences. According to what I shall call *extreme* epistemic pragmatism, on the other hand, our preferences influence our epistemic position at a more basic level, because they help determine how much justification we actually have in favor of the proposition in question. Simplifying brutally, moderate epistemic pragmatism has it that the more worried we are about a proposition’s being false, the more justification *we need* in order to know it, whereas extreme epistemic pragmatism has it that the more worried we are about a proposition’s being false, the less justification *we have* for it. Recently, Fantl and McGrath have presented an interesting argument for moderate epistemic pragmatism, an argument which relies on the principle that (roughly) knowledge is sufficient for action (KA). In this paper I argue that KA, together with a plausible principle about second-order evidence, entails extreme epistemic pragmatism.

1 Introduction

By “epistemic pragmatism” in general I will understand the claim that whether propositions instantiate certain key epistemic properties (such as being known or being justifiably believed) depends not just on factors traditionally recognized as epistemic, but also on pragmatic factors, such as

how costly it would be to the subject if the proposition were false. In what follows I consider two varieties of epistemic pragmatism. According to a rough formulation of what I shall call *moderate* epistemic pragmatism, how much evidence we need in favor of a proposition in order to know that the proposition is true depends on our preferences.¹ According to an equally rough formulation of what I shall call *extreme* epistemic pragmatism, on the other hand, our preferences influence our epistemic position at a more basic level, because they help determine how much evidence we actually have in favor of the proposition in question. Simplifying brutally, moderate epistemic pragmatism has it that the more worried we are about a proposition's being false, the more evidence *we need* in order to know it, whereas extreme epistemic pragmatism has it that the more worried we are about a proposition's being false, the less evidence *we have* for it.

In sections 2 and 3, I introduce the two kinds of pragmatism to be discussed and compare them with more traditional epistemological theories. Jeremy Fantl and Matt McGrath have recently argued for moderate epistemic pragmatism.² As we'll see, Fantl and McGrath's argument for moderate epistemic pragmatism depends on two crucial premises: the falsity of extreme epistemic pragmatism and a principle to the effect that (roughly) knowledge is sufficient for action (Fantl and McGrath call this principle "KA"). In sections 4 through 6 I argue, however, that (given a plausible principle of second-order evidence) KA entails extreme epistemic pragmatism. Therefore, the argument for moderate epistemic pragmatism fails, for two of its crucial premises are incompatible with each other. There is no stable intermediate position between the traditional view that practical matters are irrelevant to epistemology and the extreme view that practical matters determine the strength of our evidence.

2 Moving Thresholds and Shifty Contexts

One can have more or less evidence for a proposition, but either one knows it or one doesn't know it.³ You can have more or less evidence for the proposition that the train that you are about to board stops in Foxboro, but either you know that it stops in Foxboro or you don't. Although knowledge and evidence are thus notions of different kinds (one comes in degrees, the other doesn't), they are connected: in order to know that the train stops

¹Here and in what follows I assume that only evidence can justify, so that if knowledge entails justification then knowledge entails the having of evidence.

²See Fantl and McGrath (2002), Fantl and McGrath (2007), and Fantl and McGrath (2009).

³Some people (Lewis [1996], the entry for "Knowledge" in Quine [1987], Hetherington [2001]) think that knowledge does come in degrees. But even they must admit that there is a minimum amount of evidence that you must have in order to count as knowing, and the question that sets up our problem is about the location of this lower bound. I also ignore the vagueness of "knows" here.

in Foxboro you have to have enough evidence for that proposition.⁴ How much evidence? We can ask the same question using a slightly different terminology. Let us say that when a subject's evidence for p is sufficient for the truth of an attribution of knowledge (that is, if in such a situation the attribution is false that is not because the subject lacks evidence), then that subject's evidence for believing that p has crossed the "knowledge threshold." The question now is: how high is the knowledge threshold?

We can distinguish the possible answers according to whether they hold that the knowledge threshold is fixed or mobile—giving us "fixed-threshold" and "moving-threshold" theories. Traditional epistemological theories are fixed-threshold theories: according to them, there is a fixed amount of evidence that is needed for the truth of a knowledge attribution in any context. We can factor whether we have crossed the knowledge threshold for believing that p into two components: how much evidence *we have* for believing that p , and how much evidence *we need* for believing that p . According to fixed-threshold theories, the second factor is constant across all cases, and thus whether we have knowledge-level evidence depends exclusively on how much evidence we have.

Fixed-threshold theories have recently come under attack. The objection is that they simply cannot do justice to our practices of knowledge attribution. Suppose, for instance, that Jeremy and Matt are both about to board the train from Boston to Providence. They both wonder whether the train stops in Foxboro or is instead the "express." Just a moment ago both Matt and Jeremy heard the announcement that the train is indeed the express. Matt says to himself: "Oh, that's right: they just announced that the train is the express, so I know that it is", and stops worrying about the matter—there is not much riding, for Matt, on whether the train stops in Foxboro or not. By contrast, *a lot* is riding for Jeremy on whether the train stops in Foxboro or not. If it does, then he won't be on time for a meeting that will determine the course of the rest of his life. Jeremy says to himself: "They just announced that the train is the express, but the meeting is so important that I better double-check with an agent. Until I do, I do not know that the train is the express."⁵

⁴Some people (for instance, Lewis [1996] again) think that justification is not (always) necessary for knowledge. Given my previous stipulation that evidence is what justifies, these people would disagree with me that knowing that p entails having evidence for p . I would argue (but not here) that this position stems from an overly intellectualistic conception of justification (and, correspondingly, of having evidence). But never mind: even those who think that justification isn't necessary for knowledge have to admit that a subject cannot know that p unless his epistemic position with respect to p is strong enough. This notion of the *strength of one's epistemic position* can replace the notion of having evidence throughout.

⁵The case is taken from Fantl and McGrath (2002), who in turn take it from early discussion of knowledge ascriptions by Cohen (1988). It is important to keep in mind that part of the story is that Jeremy still believes, to the same degree that does Matt, that the train is the express.

Fixed-threshold theories entail that either Matt is wrong in attributing knowledge to himself or Jeremy is wrong in attributing lack of knowledge to himself. Non-skeptic fixed-threshold theories (that is, fixed-threshold theories according to which the knowledge threshold is moderately easy to achieve) have it that Matt is right in attributing knowledge to himself and Jeremy is wrong in attributing lack of knowledge to himself, whereas skeptic fixed-threshold theories (that is, fixed-threshold theories according to which the knowledge threshold is almost impossible to achieve) have the opposite consequences. But, the objection goes, what Matt and Jeremy tell themselves *sounds right in both cases*: it is what we would tell ourselves if we were in their (respective) shoes. The cases can be easily multiplied. Therefore, fixed-threshold theories entail widespread error in our common knowledge-attributing practices. To this extent, it is argued, they are defective.

Moving-threshold theories can avoid this consequence, for they can claim that the knowledge threshold is different for Jeremy and Matt. Thus, even though Matt and Jeremy have the same evidence for the proposition that the train is the express, Jeremy's knowledge threshold is such that this evidence is not enough for a corresponding knowledge attribution to be true, whereas Matt's knowledge threshold is such that that very same evidence suffices for the truth of the corresponding knowledge attribution. One prominent moving-threshold view is contextualism.⁶ According to contextualism, a knowledge attribution of the form *S knows that p* can express different propositions in different conversational contexts, and this variance is due to the fact that the predicate "knows" can express different relations in different conversational contexts. These relations differ by where they locate the knowledge threshold. Thus, there is a very lenient knowledge relation that locates the knowledge threshold very low on the scale, and there are stricter knowledge relations that locate the knowledge threshold higher on the scale. A subject's belief may satisfy some but not all of these relations. Thus, a knowledge attribution may express a true proposition in one conversational context and a false one in another. Contextualism is, therefore, a "variantist" theory, a theory according to which the predicate "knows" can express different relations in different conversational contexts. Traditional fixed-threshold theories, on the other hand, are invariantist: they hold that there is only one relation that is the semantic value of "knows".

Let's consider two different conversational contexts, one where Jeremy's concerns are somehow made salient and one where they aren't (let's call the

⁶Three influential contextualists are Cohen (1988), DeRose (1995), and Lewis (1996). Of these, Cohen is the one whose terminology is closest to the one used in this paper: instead of talking of the level of justification needed for knowledge, DeRose and Lewis talk about possibilities that need to be eliminated.

first a “high-standards” context and the second a “low-standards” one).⁷ Let’s also suppose that, in each of these contexts, someone says “Matt knows that the train is the express” and someone else says “Jeremy knows that the train is the express.” Tables 1 and 2 represent what a contextualist theory and a traditional invariantist (non-skeptic) theory would say about the truth-values of those utterances.

	Matt knows	Jeremy knows
High standards	False	False
Low standards	True	True

TABLE 1. Contextualism

	Matt knows	Jeremy knows
High standards	True	True
Low standards	True	True

TABLE 2. Traditional Invariantism

Contextualism, then, is a moving-threshold theory *because* it is a variantist theory, and traditional theories have been both fixed-threshold theories and invariantist. This may suggest that a moving-threshold theory must be variantist. But this is not the case: there is logical space for invariantist moving-threshold theories. According to this kind of theories, the semantic value of “knows” is invariant across contexts (this is the invariantist part), but the amount of evidence a subject must have in order for a knowledge attribution to be true may vary from case to case. Table 3 represents that kind of position.

	Matt knows	Jeremy knows
High standards	True	False
Low standards	True	False

TABLE 3. Moving-Threshold Invariantism

A theory represented by Table 3 is invariantist because which proposition is expressed by a knowledge attribution doesn’t vary with the context of utterance, as shown by the fact that the rows in the table have the same truth-values. But a theory represented by Table 3 is a moving-threshold theory because subjects who have the same evidence for some proposition may differ in whether they know the proposition, as shown by the fact that the columns in the table have different truth-values. So, it is possible

⁷Different contextualist theories will differ in the identification of the mechanisms by which certain features are made contextually salient.

for there to be moving-threshold invariantist theories. Indeed, one of the versions of epistemic pragmatism that I discuss below is just such a theory.⁸

3 Moderate and Extreme Epistemic Pragmatism

Like knowledge, practical rationality is all-or-nothing. Either you are rational to act as if the train is the express or you are not.⁹ Like knowledge, too, practical rationality is connected to evidence: in general, in order for you to be rational to act as if the train is the express you have to have *some* evidence that it is.¹⁰ How much evidence? In this case, we have a well-developed answer: as much as required by your preferences. Matt's practical concerns are such that he has enough evidence to make it rational for him to act as if the train is the express. Jeremy's practical concerns, however, are such that his evidence (which is, remember, the same as Matt's) is *not* enough to make it rational for *him* to act as if the train is the express.

So there are two different binary notions, knowledge and practical rationality, both of which are connected to the same degree notion, evidence. We have a relatively clear understanding of how practical rationality is connected to evidence, and no clear understanding of how knowledge is connected to evidence. Why not explain the unclear by the clear, at least partially? Why not say that, in order to know that p , the amount of evidence for p that you have has to be at least high enough to make it rational for you to behave as if p ?¹¹ That thought can be captured in the following principle, which Fantl and McGrath (2002) call "KA":

⁸Is it also possible for there to be a fixed-threshold, variantist theory? Yes, it is. The kind of contextualism introduced in the text distinguishes the different knowledge relations that the knowledge predicate may refer to by where they locate the knowledge threshold, but there may well be a contextualist theory according to which all the different knowledge relations locate the threshold in the same place, but differ on some other dimension—for instance, with respect to what kinds of defeaters are relevant to a knowledge attribution.

⁹Following Fantl and McGrath (2002), we will say that you are rational to act as if p if and only if, for any pair of acts A and B available to you, you are rational to prefer A to B if and only if you are rational to prefer A and p to B and p . One complication: among the irrational acts that we might perform, it makes sense to think of some of them as "less irrational" than others—they have a suboptimal, but not minimal, expected utility. In *that* sense, practical rationality does come in degrees.

¹⁰I say "in general" because there may be cases where you do not care at all whether p . In those cases, for you to act as if p *just is* for you to act as if *not-p*, in which case no evidence at all for believing that p is needed for you to be rational in acting as if p . See also note 11.

¹¹I say "at least high enough" and not "just as high" because it is not plausible to suppose that if you have enough evidence for p to make it rational for you to act as if p then you are in a position to know that p . Sometimes you don't care one way or the other whether p , in which case no evidence at all is necessary for you to be rational to act as if p . Sometimes you care only very little, in which case very little evidence for believing that p is necessary for you to be rational to act as if p . But some (not very low) amount of evidence is always necessary for you to know that p .

KA: For any possible subject S , if S knows that p , then S is rational to act as if p .¹²

Fantl and McGrath don't take KA to be obviously true. They present at least two arguments for it. First, they point out that, when defending courses of action, we routinely appeal to what we know. Thus, for instance, if I plan to go swimming in the ocean at night with my glasses on and my wife asks me why, I can sensibly reply "Don't worry, I know that I won't lose them." This practice of ours of appealing to what we know in defense of what we and others do makes sense under the assumption that KA is true (and that we are, perhaps implicitly, aware of its truth). Second, Fantl and McGrath (2007) also argue that a plausible closure inference entails KA with the aid of the following knowledge-to-action principle (which they call "KB"):

KB: If S knows that A is the best thing for him to do in light of all of his goals, then S is rational to do A .¹³

KB really is appealing.¹⁴ To deny it one would have to hold, for example, that even though Jeremy knows that boarding the train straightaway is the best thing for him to do in light of all of his goals, he nevertheless isn't rational to board the train straightaway. It is hard to see how to make sense of that claim.¹⁵

Now, most epistemologists nowadays will accept that knowledge is *fallible*, at least in the sense described by the following principle:

Fallibilism: It is possible for a subject S to know that p even if there is a subject S' such that S' has more evidence for p than S has.¹⁶

KA and fallibilism together entail the following principle, which I will take as the official definition of moderate epistemic pragmatism:

Moderate epistemic pragmatism: There are possible subjects S and S' and there is a proposition p such that (i) S and S' have the same evidence for p , but (ii) S is in a position to know that p whereas S' is not.

¹²In later work Fantl and McGrath replace "S is rational to act as if p " with " p is warranted enough to justify you in ϕ -ing, for any ϕ ." (Fantl and McGrath [2009, 66]) This replacement doesn't affect the arguments in this paper.

¹³For a defense of related principles, see also Fantl and McGrath (2009, chapter 3).

¹⁴Except, as Fantl and McGrath note, in those cases where whether S does A causally affects what is the best thing for S to do.

¹⁵But see the last paragraph of section 7 for a suggestion.

¹⁶See Fantl and McGrath (2009, chapter 1) for an illuminating discussion of the problems surrounding how to define fallibilism. Hawthorne and Stanley (2008) argue that "knowledge delivers probability 1." It is not clear whether they mean to deny fallibilism. Williamson (2000) thinks that if S knows that p then S 's evidential probability for p is 1. However, Williamson also argues that knowledge is defeasible by misleading evidence, which means that the evidential probability of p can fall below 1 for a subject.

What is it to be in a position to know a certain proposition? It is for the amount of evidence for that proposition to be high enough for knowledge—that is, it is for you to have knowledge-level evidence.

Under the assumption that knowledge entails justification (and that evidence is what justifies), KA and fallibilism entail that the knowledge threshold depends on your preference structure. If there isn't much riding, for instance, on whether there is a glass of water on the table, then KA doesn't impose any very stringent requirement on what it takes for you to know that there is a glass of water on the table—the knowledge-threshold is set, in accordance with fallibilism, at a level where a certain non-maximal amount of evidence is needed. But suppose that the issue of whether there is water or gasoline in the glass is of vital importance to you, because you are extremely thirsty. In that case, KA entails that, in order to reach the knowledge threshold, you must have enough evidence for thinking that there is water in the glass. How much evidence? As much as it is necessary to make it rational for you to act as if there is water in the glass—in the imagined case, as much as it is necessary to make it rational for you to drink the liquid in the glass. Crucially, this may mean that you need *more* evidence to know than you would have needed if the stakes had been lower. Moderate epistemic pragmatism is thus an invariantist, variable-threshold theory, a theory of the kind represented by Table 3 above.

We can further illustrate the argument from KA and fallibilism to moderate epistemic pragmatism with our familiar case: Matt knows that the train is the express, Jeremy has the same evidence that Matt has for believing that the train is the express, but (because there is a lot riding on it for him) Jeremy is not rational to straightaway board the train. If we rejected fallibilism we could hold that Jeremy's situation is not possible, because he has the same evidence as Matt and Matt knows that the train is the express. Given the negation of fallibilism, this evidence must be maximal—it is impossible to have more evidence, in which case it is arguably impossible for it to be irrational for you to act as if the train is the express. But, given fallibilism, the argument is now the following:

- (1) Matt knows that the train is the express.
- (2) Jeremy is not rational to act as if the train is the express.
- (3) Jeremy and Matt have the same evidence for believing that the train is the express.
- (4) Jeremy is not in a position to know that the train is the express.
(From 2 and KA)
- (5) Moderate epistemic pragmatism is true. (From 1, 3, and 4).

The argument is valid. The step from 2 and KA to 4 requires some comment. All that immediately follows from 2 and KA is that Jeremy doesn't know that the train is the express. But, as the description of the case makes clear, Jeremy's lack of knowledge is not due to lack of belief, falsity of the proposition believed, or failure to satisfy any other of the traditional

conditions on knowledge. Therefore, Fantl and McGrath conclude, it must be due to the fact that Jeremy is not in a position to know. Although this move gives rise to interesting questions, I won't object to the move from 2 and KA to 4. I won't deny premises 1 and 2 either—that much can be seen as merely a stipulation about what case we are thinking of. Premise 3 will be the focus of the rest of this paper. I will make some preliminary comments on premise 3 after introducing extreme epistemic pragmatism.

Moderate epistemic pragmatism is of course compatible with the claim that how much evidence you actually have for the proposition, for instance, that there is water in the glass, is independent of your preferences. That is, someone who accepts moderate epistemic pragmatism is free to make the following claims about the case: your evidence for thinking that there is water in the glass is the same whether or not it is vitally important to you that there be water in the glass, but whether that preference-independent amount of evidence reaches the knowledge threshold *does* depend on how important it is to you that there be water in the glass.

Extreme epistemic pragmatism, on the other hand, holds that your preferences help determine how much evidence you have for a given proposition. For instance, the extreme epistemic pragmatist will say that if it is vitally important to you whether there is water in the glass, then you thereby have less evidence for thinking that there is water in the glass than someone who doesn't care much one way or the other. More generally:

Extreme epistemic pragmatism: For any possible subjects S and S' and proposition p , if S and S' are exactly alike in all non-epistemic matters except for the fact that S' is not rational in acting as if p (whereas S is so rational), then S has more evidence for p than S' does.

Given the very plausible claim that epistemic properties supervene on non-epistemic properties (there cannot be a difference in epistemic properties without a difference in non-epistemic properties), the antecedent of extreme epistemic pragmatism guarantees that if there is an epistemic difference between S and S' that difference will be due to the fact that there is a practical difference between S and S' , and its consequent says that there is indeed such a difference. We must be careful not to set up matters so that extreme epistemic pragmatism comes out as incoherent. If we stipulate that how much evidence a subject has is a purely epistemic matter, in the sense that it depends only on truth-related features of the subject's situation and not, for instance, on the subject's practical situation, then we make extreme epistemic pragmatism false by stipulation.¹⁷ But the difference between moderate and extreme epistemic pragmatism resides precisely on whether there is such a thing as a purely epistemic domain of evaluation. No stipulation will make this question go away.

¹⁷Fantl and McGrath (2007) and Fantl and McGrath (2009) make precisely this stipulation.

Where does extreme epistemic pragmatism fit in with respect to the pair of distinctions introduced above—the distinction between fixed-threshold and moving-threshold theories and the distinction between variantist and invariantist theories? Even under the assumption that knowing that p entails having evidence for p , extreme epistemic pragmatism, as I have defined it, is compatible with any of the four possible kinds of theories resulting from the combination of those two distinctions. A theory that subscribes to extreme epistemic pragmatism can be variantist or invariantist, and it can posit a moving or a fixed-threshold. Unlike contextualism, extreme epistemic pragmatism is not primarily a theory about knowledge attributions, and unlike moderate epistemic pragmatism, it is not a theory about where the knowledge threshold lies. Extreme epistemic pragmatism is a theory about how practical matters affect evidence. Partly because of that reason, none of the three tables introduced above can unproblematically describe extreme epistemic pragmatism. According to extreme epistemic pragmatism, the cases that those tables aim to represent are not possible. Remember that part of the set-up of the case involving Matt and Jeremy is that, despite the differences in their practical situation, they have the same evidence—and that is precisely what the extreme epistemic pragmatist denies. If we concentrate our attention on *invariantist* theories that adhere to extreme epistemic pragmatism *and* we don't make the assumption of evidential equality, supposing only that Matt's and Jeremy's practical positions are different, then extreme epistemic pragmatism is compatible with the truth-values given in Table 3. But notice that, in this case, if moderate and extreme epistemic pragmatism give the same answers it may be that they do so for very different reasons: moderate epistemic pragmatism will be represented by Table 3 because it holds that even though Matt and Jeremy have the same evidence, Jeremy needs more evidence than Matt in order to know, whereas a theory that adheres to extreme epistemic pragmatism may be represented by that table for the same reason, but it also may be represented by that table merely because it holds that Matt has more evidence than Jeremy.¹⁸ The extreme epistemic pragmatist need not say, however, that Matt has knowledge and Jeremy doesn't. The extreme epistemic pragmatist is only committed to saying that the practical situation affects, by itself, how much evidence subjects have, but is free to remain neutral on the question whether there are cases such that the subject doesn't know only because his practical situation made it the case that his evidence didn't reach knowledge-level. Thus, extreme epistemic pragmatism doesn't entail moderate epistemic pragmatism. The relative extremity of extreme epistemic pragmatism doesn't reside on the fact that it is logically stronger

¹⁸If, on the other hand, we retained the assumption that Matt and Jeremy have the same evidence, leaving it open whether their respective practical positions change or not, then a theory that adheres to extreme epistemic pragmatism may instead agree with the truth-values depicted in Table 2 (depending on whether the extreme epistemic pragmatist theory is in addition a moving-threshold or a fixed-threshold theory).

than moderate epistemic pragmatism, but rather on the fact that it claims that our practical situation can by itself affect *how much evidence we have*, and not merely how much evidence we need to have in order to know.

Let us re-examine Fantl and McGrath's argument for moderate epistemic pragmatism in light of the distinction between it and extreme epistemic pragmatism. Fantl and McGrath call the denial of moderate epistemic pragmatism (roughly, the claim that if two subjects S and S' have the same evidence then S is in a position to know that p if and only if S' is), "Purism," because it allegedly defines a dimension of *pure* epistemic evaluation, free from any taint of practical matters. But notice that Purism deserves its name only under the assumption that extreme epistemic pragmatism is false. For suppose that extreme epistemic pragmatism was true. Then, Purism might well be true even if practical matters do influence whether you are in a position to know. In other words, if extreme epistemic pragmatism is true, then practical matters influence directly how much evidence we have, and so the purity ensured by the truth of Purism is only apparent. Moreover, premise 3 in Fantl and McGrath's argument for moderate epistemic pragmatism presupposes the falsity of extreme epistemic pragmatism.¹⁹ But I will argue that KA entails extreme epistemic pragmatism given a plausible assumption. Therefore, if Fantl and McGrath are right about KA, then they are wrong about extreme epistemic pragmatism.

4 A Principle of Second-Order Evidence

Evidence of evidence is evidence. If you have evidence that scientists have evidence for the proposition that there is life outside the solar system, then you thereby have evidence for the proposition that there is life outside the solar system. If someone is just like you evidentially except for the fact that she lacks that second-order evidence, then you have more evidence than she does for the proposition that there is life outside the solar system. *A fortiori*, if you have evidence that scientists *know* that there is life outside the solar system, then you thereby have evidence for the proposition that there is life outside the solar system. If someone is just like you evidentially except for the fact that she lacks that second-order evidence, then you have more evidence than she does for the proposition that there is life outside the solar system. Second-order evidence need not be about somebody else: the same principles hold if you have evidence that you have evidence (or that you know) that there is life outside the solar system. Again, if someone is just like you evidentially except for the fact that she lacks this second-order evidence about herself, then you have more evidence than she does for the proposition that there is life outside the solar system.

These considerations motivate the following principle of second-order evidence:

¹⁹Premise 3 is assumption (A2) in [Fantl and McGrath \(2009, 86\)](#).

SOE: For any possible subjects S and S' , necessarily,
IF

- (i) S has all the evidence for p that S' has;
- (ii) S' has all the evidence against p that S has; and
- (iii) S has more evidence for his belief that p amounts to knowledge than S' does for her belief that p amounts to knowledge,

THEN

S has more evidence for p than S' does.

Let us call any pair of subjects who satisfy clauses (i) and (ii) of the antecedent of SOE, “evidential cousins.” It is important that SOE contains this restriction to evidential cousins, for it would otherwise be very implausible. Suppose that Joe has just learned about Goldbach’s conjecture, and mistakenly but justifiedly thinks that there is a proof of it. Moe, on the other hand, is a mathematician who has spent his career working on Goldbach’s conjecture. Moe has evidence that he doesn’t know that Goldbach’s conjecture is true that Joe lacks, but we still want to say that Moe has more evidence for Goldbach’s conjecture than Joe does. In virtue of having more evidence than Joe for thinking that he doesn’t know that p , Moe *does* have evidence against p that Joe lacks, but that additional evidence is outweighed by his additional independent evidence *for* p . In this case, Joe and Moe are not epistemic cousins because Moe has evidence for p that Joe lacks. A similar case could be constructed that exploits a violation of clause (ii).

But if the pair of subjects in question are epistemic cousins, then SOE is indeed very plausible. I will not attempt here a full defense of SOE,²⁰ but I will point out the following: even if, as stated, SOE is false, that is not enough to let the friend of (merely) moderate epistemic pragmatism off the hook. There may well be counterexamples to SOE, counterexamples which may necessitate a fair amount of Chisholming to get a correct principle on the table. It may even be the case that it is impossible to formulate a true SOE-like principle in any neat way (although I don’t believe that it is). But the use of SOE in the next section is in the context of a very specific kind of case: a case where the only difference between the two subjects in question is in their practical situation. Thus, what the friend of KA needs to show in order to block that argument is not just that SOE is false, and not even that no true SOE-like principle is formulable, but rather that the basic idea that second order evidence that there is evidence for p is evidence for p fails in the kind of cases to which the argument applies it. I come back to this issue below, in section 6.

²⁰For a defense and application of a similar principle, see [Fantl \(2003\)](#).

5 The Argument Against Moderation

I will now argue that SOE and KA together entail extreme epistemic pragmatism. Therefore, Fantl and McGrath's argument for moderate epistemic pragmatism cannot be sound: if they are right about KA, then the third premise of that argument (that the two agents have the same evidence) is false. I will first present the argument and then comment on it:

- (1) KA: For any possible subject S , if S knows that p , then S is rational to act as if p .
- (2) SOE: For any possible subjects S and S' , necessarily, if (i) S has all the evidence for p that S' has; (ii) S' has all the evidence against p that S has; and (iii) S has more evidence for his belief that p amounts to knowledge than S' does for her belief that p amounts to knowledge, then S has more evidence for p than S' does.
- (3) If S is not rational to act as if p , then S has evidence that he is not rational to act as if p . [Premise]
- (4) If KA is true and S has evidence that he is not rational to act as if p , then S has evidence that he doesn't know that p . [Premise]
- (5) Matt and Jeremy are exactly alike in all non-epistemic matters except for the fact that Jeremy is not rational to act as if the train is the express (whereas Matt is so rational). [Assumption]
- (6) Jeremy has evidence that he is not rational to act as if the train is the express. [From 3 and 5]
- (7) Jeremy has evidence that he doesn't know that the train is the express. [From 4 and 6]
- (8) Matt has all the evidence for the proposition that the train is the express that Jeremy has; Jeremy has all the evidence against the proposition that the train is the express that Matt has; and Matt has more evidence for his belief that the proposition that the train is the express amounts to knowledge than Jeremy does for his belief that the proposition that the train is the express amounts to knowledge. [From 5 and 7]
- (9) Matt has more evidence for the proposition that the train is the express than Jeremy does. [From 2 and 8]
- (10) *Extreme epistemic pragmatism*: For any possible subjects S and S' and proposition p , if S and S' are exactly alike in all non-epistemic matters except for the fact that S' is not rational in acting as if p (whereas S is so rational), then S has more evidence for p than S' does. [From 5 through 9]

The only premises in the argument are KA, SOE, and 3 and 4. Why believe in 3, the claim that if S is not rational to act as if p , then S has evidence that he is not so rational? Remember that we are assuming that what is rational for you to do supervenes on your preferences and your beliefs. It doesn't take any strong thesis of privileged access to conclude

that whenever it is irrational for you to act as if p you will have at least some evidence that it is.

What about 4, the claim that (under the assumption that KA is true) someone who has evidence that he is not rational in acting as if p thereby has evidence that he doesn't know that p ? Here it is important to recall the arguments for KA. One of them was that we (at least implicitly) rely on the truth of KA when justifying and criticizing actions. The second argument for KA relied on KB, and Fantl and McGrath's defense of KB appeals both to its *obvious truth* as well as to its role in the criticism and justification of action. Thus, Fantl and McGrath's own arguments for KA strongly involve the claim that subjects are (perhaps implicitly) aware of the truth of KA. The contrapositive of KA says that if S is not rational in acting as if p , then S doesn't know that p . Premise 4, therefore, amounts to saying that (given that S is at least implicitly aware of the truth of that conditional) if S has evidence for the antecedent of that conditional, then S has evidence for its consequent.

This might seem like a dangerous claim, given the fact that "is evidence for" is not closed under known logical implication.²¹ But premise 4 relies on no general principle of evidential closure. That evidence fails to be closed is not a good argument for the claim that, in a particular case, evidence for p is not evidence for a q entailed by p . The cases where evidence fails to be closed all share a common structure, a structure which we have no reason to think is present in the cases to which the foregoing argument applies. The structure in common to all cases where evidence fails to be closed is most easily revealed by thinking in probabilistic terms (although one need not think of evidence probabilistically to appreciate the structure). Suppose that e is evidence for p for S only if e raises S 's epistemic probability for p . If p entails q , then the probability of q is at least as high as the probability of p . Therefore, e can raise S 's epistemic probability for p but not raise S 's epistemic probability for q (that is, e will fail to be closed under entailment) only if, before receiving e , S 's epistemic probability for q was higher than S 's epistemic probability for p . But there is absolutely no reason to think that, in the cases that interest us, this will be the case. That is, there is absolutely no reason to think that, before realizing that he is not rational to act as if p , S had more evidence for the proposition that he didn't know that p than for the proposition that he was not rational to act as if p . Therefore, given that there is no reason to think that the cases to which 4 applies share the structure of cases where evidence fails to be closed, the initial plausibility of 4 is not threatened by those considerations.

We could, if we so wished, replace premises 3 and 4 in the argument by a single premise: that if S fails to know that p in virtue of the fact that S 's evidence for p doesn't meet the knowledge-threshold, then S has evidence

²¹For instance, that there is a hairless animal in the bush is evidence for the proposition that there is a hairless dog in the bush, but is not evidence for the proposition that there is a dog in the bush.

that he doesn't know that p . The general principle that if S doesn't know that p then S knows that he doesn't know that p is of course false, but no such general principle is advocated in the proposed new premise. Rather, the idea is that in the particular case where the lack of knowledge is due to insufficient evidence, the subject has evidence that he doesn't know.²² Given that the subject S in question in the argument lacks knowledge precisely for this reason, we can get from this new premise to step 7 and continue as before.

Step 8 also deserves some comment. We can consider that step in the argument as the conjunction of two propositions: that Matt and Jeremy are evidential cousins, and that Matt has more evidence than Jeremy does for the proposition that his belief that the train is the express amounts to knowledge. Notice, first, that Fantl and McGrath themselves will agree with the claim that, in the case described, Matt and Jeremy are epistemic cousins—indeed, they think that Matt and Jeremy have the same evidence for the proposition that the train is the express, from where it follows that they are epistemic cousins with respect to that proposition. But the claim also follows from 5: given the supervenience of the epistemic on the non-epistemic, and given that the only non-epistemic difference between the subjects is that only Matt is rational to act as if the train is the express, it follows that they are epistemic cousins—for no party to this dispute (neither the moderate nor the extreme epistemic pragmatist, nor, obviously, the epistemic purist) thinks that the fact that S is not rational to act as if p means that he has *more* evidence for p than another subject who is rational in so acting. For analogous reasons, it follows also from 5 and 7 that the second proposition in 8 is true.

Remember that Fantl and McGrath's argument from KA to moderate epistemic pragmatism had as a premise the denial of extreme epistemic pragmatism (see section 3). The argument of this section shows that KA itself (together with SOE) *entails* extreme epistemic pragmatism. Therefore, Fantl and McGrath's argument for moderate epistemic pragmatism fails.

6 Evidential and Extra-Evidential Conditions on Knowledge

I have just argued that, given SOE, KA entails extreme epistemic pragmatism. I don't presume that, as it stands, SOE is free from counterexamples, but I do think that whatever Chisholm's SOE needs won't affect its use

²²Even the restricted principle may need refining. But keep in mind that the cases to which the premise will apply are ones where the subject in question (Jeremy) is not mistaken about anything. All we need, actually, is not any general principle, but the claim that, as things stand with the subject in this kind of case, he has evidence that he doesn't know that p . Premises 3 and 4, and the proposed revised premise, are reasons to think that this is true, but even if we doubt those reasons it still seems true that subjects in a situation like Jeremy will have evidence that they don't know.

in showing that KA entails extreme epistemic pragmatism.²³ But there is a worry about SOE which does not consist in claiming that it gets this or that isolated case wrong, but is rather the objection that it ignores a crucial distinction between evidential and extra-evidential conditions knowledge. SOE, the objection goes, should apply only to evidential conditions, and *being rational to act as if p* is an extra-evidential condition. In this section I consider three proposals as to how to draw the distinction between evidential and extra-evidential conditions on knowledge. I find specific problems with appealing to each of them to argue against my use of SOE. But I close the section by arguing more generally that, whatever the distinction between evidential and extra-evidential conditions on knowledge comes down to, SOE *should* ignore it. SOE applies to cases where subjects receive evidence that they don't know that *p*, not to cases where subjects receive evidence that they fail to satisfy some condition on knowledge. Even when the former kind of evidence is constituted by the latter, their evidential impact is still different. Therefore, the distinction between evidential and extra-evidential conditions on knowledge is irrelevant to SOE.

Let me start by trying to flesh out in intuitive terms what the objection is. According to SOE, when *S* has more evidence for the proposition that he knows that *p* than *S'* does for the corresponding proposition *and S* and *S'* are evidential cousins (which means that *S* has all the evidence for *p* that *S'* has and *S'* has all the evidence against *p* that *S* has), *S* thereby has more evidence for *p* than *S'* does. But, the objection goes, there are two ways for a subject to have more evidence than another subject for the proposition that he has knowledge. First, it may be that *S* has evidence that he has *more evidence* that *p* than *S'* does. But, second, it may be that *S* has more evidence than *S'* that he satisfies some *extra-evidential* condition on knowledge. If this second case obtains, the objection continues, then *S* and *S'* have the same evidence for *p*. This is so because one can come to lose some evidence for believing that *p* in one of two ways: one may acquire some evidence against *p*, or one may acquire some evidence that one's antecedent reasons for believing that *p* were not good enough (this is the distinction between "rebutting" and "undercutting" defeaters—see Pollock [1986]). But learning that one fails to satisfy some extra-evidential condition for knowing that *p* is not learning that there is some evidence against *p* or that one's reasons for believing that *p* are not good. Therefore, the mere fact that *S* has more evidence than *S'* for the proposition that he knows that *p* doesn't entail that *S* has more evidence than *S'* for *p*—not even when *S* and *S'* are evidential cousins. Moreover, the argument continues, that one is rational in acting as if *p* is clearly an extra-evidential condition on knowing that *p*. Therefore, the argument of the preceding section fails, because Matt has more evidence than Jeremy for the proposition that the

²³Something that we should note is that, in the cases where I apply SOE, the subjects have no mistaken belief about their epistemic position. Therefore, modifications of SOE that may be called for if it is to apply to subjects with such mistaken beliefs don't matter for my argument.

train is the express only because Matt's and Jeremy's preferences are such that Matt (but not Jeremy) is rational to act as if the train is the express, and so this extra evidence doesn't entail any extra evidence for believing that the train is the express.

The objection relies crucially on the distinction between evidential and extra-evidential conditions on knowledge. It is not obvious that the distinction can be made clear without prejudicing the case against SOE. I will consider three proposals about how to distinguish between evidential and extra-evidential conditions on knowledge and argue that none of them can be used against SOE.

First, the most straightforward way of understanding the distinction is as follows: something is an evidential condition on knowledge if and only if it is a condition on knowing that p having to do with how much evidence a subject needs to know that p . But on this understanding it is simply not true that being rational to act as if p is an extra-evidential condition, for a crucial determining factor of whether one is rational in acting as if p is how much evidence for p one has.

Second, we may want to try a supervenience thesis: something is an extra-evidential condition on knowledge just in case it *doesn't* supervene on evidence. In other words, if a condition is such that a subject may satisfy it even though another possible subject who has exactly the same evidence does not, then the condition is extra-evidential. Under this understanding, being rational to act as if p is an extra-evidential condition on knowledge. The question now is whether extra-evidential conditions, so understood, should be excluded from the scope of SOE. In other words, the question is whether having evidence regarding the satisfaction of an extra-evidential condition (under this conception of the distinction) can make a difference as to the strength of the subject's evidence with respect to his first-order beliefs.

According to this second characterization of the distinction between evidential and extra-evidential conditions, the anti-Gettier conditions count as extra-evidential. Gettier (1963) showed that justified true belief is not sufficient for knowledge. One way of putting Gettier's result is by saying that what makes the difference between knowledge and true belief can be extra-evidential: that is, two subjects can have the same evidence for p , while only one of them knows that p , even if they both truly believe that p .²⁴ Whatever other conditions are necessary for knowledge we can call "anti-Gettier" conditions. We can then ask whether having evidence about the satisfaction of an anti-Gettier condition can make a difference to the strength of one's evidence for first-order beliefs. That is, if two subjects differ only in that one of them has evidence that an anti-Gettier condition is not satisfied for him to know that p , do those subjects have the same

²⁴This doesn't mean, of course, that Gettier showed that Purism is false. One may think otherwise if one forgets that "being in a position to know," as used in the statement of Purism, means that the level of justification for the belief in question is high enough for knowledge.

evidence for p ? Fortunately for our use of SOE in the previous section, the answer is no.

Let us first recall one traditional Gettier case: Jones is justified in believing that Nogot owns a Ford (even though Nogot doesn't own a Ford), and infers from that belief and his knowledge that Nogot works in his office that someone in his office owns a Ford (which is true, because Havitt owns a Ford). Let us now suppose that Smith (who also works in the same office as Jones) is in the same evidential position as Jones: Smith also believes that Nogot owns a Ford, and infers from that belief and his knowledge that Nogot works in his office that someone in his office owns a Ford.

Now, what would happen if one (but only one) of our subjects were to become aware that he is in a Gettier situation? Suppose, for instance, that someone whom Jones trusts tells him "Your belief that someone in the office owns a Ford constitutes a Gettier case." Interestingly, even though what Jones's friend tells him was true *before* he told it to him, his telling him *makes it false*. For what his friend is telling Jones is that his belief that someone in his office owns a Ford is justified and true but doesn't amount to knowledge. But if Jones is to trust his friend that his belief is true, he cannot trust him that it doesn't amount to knowledge. Perhaps, then, Jones acquires information that his *original* basis for believing that someone in his office owns a Ford was faulty, but in the same breath he acquires a new basis for believing that someone in his office owns a Ford (a basis which is good enough for knowledge). That a belief of a subject constitutes a Gettier case is something that cannot be learned by that subject (although a subject can learn that a belief of his *used* to constitute a Gettier case, and somebody else can learn that my present belief constitutes a Gettier case).²⁵

But Jones can have evidence that his belief doesn't satisfy some anti-Gettier condition without thereby having evidence that his belief constitutes a Gettier case. For instance, Jones's friend could tell him: "I won't tell you whether it is true or not that someone in your office owns a Ford, but your belief that someone does is such that even if it were true you wouldn't know it" (or perhaps he could tell Jones "Your belief that someone in your office owns a Ford doesn't satisfy some anti-Gettier condition on knowledge", after explaining to him what an anti-Gettier condition is). Thus, it is possible for a subject to acquire evidence that his belief doesn't satisfy some of the extra-evidential conditions required for knowledge. But this doesn't constitute a counterexample to SOE. For, after hearing his trustworthy friend, *Jones is clearly less justified than he was before* in believing that someone in his office owns a Ford.

It is, perhaps, surprising that this should be so, but learning that one's belief fails to satisfy an anti-Gettier condition on knowledge robs one of justification for that belief. Given that, together with Fantl and McGrath,

²⁵That is, "I am in a Gettier-situation with respect to p " is a *blindspot*—see [Sorensen \(1988\)](#). Notice also the parallel with Moore-paradoxical propositions such as "It is raining and I don't believe that it is raining."

we are assuming that how much justification one has supervenes on one's evidence, we are led to say that after learning that one's belief doesn't satisfy an anti-Gettier condition one is thereby less justified in that belief *because* one now has less evidence for it. This entails embracing the somewhat paradoxical-sounding (but not, of course, really paradoxical) claim that learning about some extra-evidential condition can by itself affect how much evidence one has. Therefore, Gettier cases do not represent a counterexample to SOE.

Given that learning that one's belief fails to satisfy some extra-evidential condition related to the Gettier problem *does* affect one's evidence for the first-order belief, the objection to the argument of the preceding section based on the idea that SOE fails in precisely this kind of case is flawed. Of course, it may well be that there are *some* extra-evidential conditions for which SOE fails, but the mere fact that the pragmatic condition on knowledge advocated by moderate epistemic pragmatist is extra-evidential (according to the characterization of the distinction that we are considering) is not by itself enough to show that SOE fails when applied to it.

Finally, consider a third proposal as to how to draw the distinction between evidential and extra-evidential conditions on knowledge. Let us say that an alleged condition on knowledge is *evidential with respect to p* if and only if it is such that knowledge of it can make a difference to how strong one's evidence for *p* is. The anti-Gettier conditions qualify as evidential under this definition, for, as we saw, knowledge of them can make a difference as to how strong one's evidence is. But, the objection goes, being rational to act as if *p* is not evidential in this sense, for knowledge of *it* doesn't make a difference to how strong one's evidence for *p* is. Given that we should exclude from the scope of SOE those conditions that fail to be evidential in this sense, we should exclude from the scope of SOE the condition imposed by KA.

There is at least an air of question-beggingness to this proposal. To be clear, what I suspect is question-begging is not the characterization of the distinction between evidential and extra-evidential conditions itself, but rather the claim that being rational to act as if *p* is extra-evidential according to this characterization. Isn't that precisely what is at stake between friends and enemies of (merely) moderate epistemic pragmatism?²⁶ I don't want to make much of this claim of question-beggingness, however, for I think that the proposal is flawed for deeper reasons. In fact, the flaw in the present proposal applies in general to all the objections to the effect that the scope of SOE should exclude extra-evidential conditions.

²⁶Indeed, one may wonder whether it is not begging the question against SOE to claim that there are *any* conditions on knowledge that would qualify as extra-evidential under the proposed characterization of the distinction. Some may propose as an obvious candidate the belief condition on knowledge. But I don't think it's obvious that knowledge that one fails to believe that *p* cannot make a difference to the strength of one's evidence for *p*.

SOE doesn't say that evidence that one fails to satisfy some condition or other on knowledge that p makes a difference to the strength of one's evidence for p . Rather, it says (approximately) that evidence that one *fails to know that* p makes a difference to the strength of one's evidence for p . All the proposals to restrict SOE's scope to evidential conditions would have SOE look into the structure of the subject's evidence that he fails to know that p . If the subject's evidence that he doesn't know that p is constituted by evidence that he fails to satisfy some extra-evidential condition then, the objection goes, we should exclude it from the scope of SOE. In other words, all the proposals considered so far would have SOE say (approximately) that evidence that one doesn't know that p makes a difference to the strength of one's evidence for p only when the evidence that one doesn't know that p is constituted by evidence that one fails to satisfy some evidential condition on knowledge.

The proposed reformulations of SOE are all unwarranted, however. Some epistemologists think that the project of coming up with illuminating conditions on knowledge is unfounded to begin with (see Williamson [2000]). But we don't need to go that far to agree that we just don't know which are the conditions on knowledge. *A fortiori*, we don't know which of the conditions on knowledge are evidential and which ones are non-evidential. But we are nevertheless many times sensitive to whether we know or not. There will be, therefore, cases where we have evidence that we fail to know some proposition without thereby having evidence about the nature (evidential or extra-evidential) or even the identity of the condition on knowledge that we fail to satisfy. Suppose that, nevertheless and unbeknownst to us, the evidence that we fail to know is constituted by what is in fact evidence that we fail to satisfy some extra-evidential condition on knowledge. The proposals considered in this section would exclude these cases from the scope of SOE. But that is the wrong result. Whether evidence that I don't know that p should make a difference to the strength of my evidence for p cannot depend on features of the evidence that I am unaware of. If all that I register is that there is evidence that I don't know that p , then that makes a difference to the strength of my evidence for p , even if the evidence that I don't know that p is in fact evidence that I fail to satisfy some extra-evidential condition on knowledge.

Of course, it may happen that I *do* know that the evidence that I don't know that p is constituted by evidence that I fail to satisfy extra-evidential conditions on knowledge.²⁷ Shouldn't SOE fail to apply in that case? I grant that if I know that the evidence that I don't know that p is constituted by

²⁷Of course, in order to know this it would have to be *true* that there are extra-evidential conditions on knowledge and I would have to *believe* that there are extra-evidential conditions on knowledge. But, as I said before, it is not clear that there are extra-evidential conditions on knowledge in the requisite sense, and even if there were only a vanishingly small minority of us would believe that there are (most people, for instance, would never even consider the question).

evidence that I fail to satisfy some extra-evidential condition on knowledge, then that shouldn't make a difference to the strength of my evidence for p . But notice that, in this case, I would have *two* pieces of evidence: that I fail to know that p , and that I fail to satisfy some extra-evidential condition on knowledge.²⁸ We are then free to say the following: the evidence that I don't know that p still makes a difference to the strength of my evidence for p , but this difference is neutralized by the evidence that I fail to satisfy an extra-evidential condition on knowledge.

Be that as it may, it doesn't really matter what we should say about SOE in cases where the subject is aware that the evidence that he doesn't know that p is constituted by evidence that he fails to satisfy some extra-evidential condition on knowledge. For it will not be the case that, in realizing that he has some evidence that he doesn't know that the train is the express, Jeremy also realizes that this is so because he fails to satisfy some extra-evidential condition on knowledge. It may be that Jeremy realizes that the evidence that he doesn't know is grounded in the fact that he is not rational to act as if the train is the express, without him realizing that this is an extra-evidential condition on knowledge.²⁹

To summarize: I have been considering the idea that we should distinguish between evidential and extra-evidential conditions on knowledge, and that we should restrict SOE to apply only to evidential conditions. That idea fails to take into account the fact that SOE applies not to evidence about conditions on knowledge, but to evidence about knowledge itself. Moreover, even if we did have to restrict SOE so as not to apply it to cases where the subject knows that he fails to satisfy an extra-evidential condition, this is not what happens to the subjects involved in my argument from the previous section.

7 Consequences of Impurity

If Fantl and McGrath are right, then KA is true. If I am right, then if KA is true, extreme epistemic pragmatism is true as well. So, if Fantl, McGrath and I are right, extreme epistemic pragmatism is true. But extreme epistemic pragmatism has some strange consequences.

Suppose that the chief of the train station announces that he will decide whether the train will stop in Foxboro or not by flipping a coin: heads, the train stops in Foxboro, tails, it doesn't. Let us stipulate that the details of

²⁸I don't mean to imply that the two pieces of evidence are independent of each other—we shouldn't, for instance, count them *both* against the proposition that I know that p .

²⁹In my argument from the previous section I appealed to the fact that, according to Fantl and McGrath themselves, we "rely on the truth" of KA in defense and criticism of actions. But this reliance on the truth of KA better not amount to a conscious and explicit belief in KA of the kind that will be necessary for ordinary folk to realize that it imposes an extra-evidential condition on knowledge. KA is a controversial philosophical thesis, knowledge of which is at best hard to come by.

the case are such that Matt and Jeremy both have good evidence to believe the chief, and therefore also good evidence for the proposition that the train stops in Foxboro if and only if the coin landed heads. Because Matt doesn't much care whether the train stops in Foxboro or not, and given the evidence he has, he has more or less the same evidence for the proposition that it does stop as for the proposition that it doesn't (and this reflects the fact that he has more or less the same evidence for the proposition that the coin landed heads as for the proposition that it landed tails). Now, because it is vitally important to Jeremy whether the train stops in Foxboro, extreme epistemic pragmatism has it that he has less evidence for the proposition that the train is the express than Matt does. So far, this is a direct consequence, however unpalatable, of extreme epistemic pragmatism. But remember that Jeremy has good evidence for the proposition that the train stops in Foxboro if and only if the coin landed heads. Therefore, extreme epistemic pragmatism has the consequence that Jeremy has more evidence for the proposition that the coin landed tails than for the proposition that it landed heads.

Many people will find this consequence unacceptable: merely being worried about whether a train stops in Foxboro or not cannot affect your evidence for thinking that a fair coin landed heads. I sympathize with this purist thought. If the consequence is sufficiently bad, then it might form the basis for a *reductio* of whatever entailed it—in this case, the conjunction of KA and SOE (I assume that the other propositions needed for the entailment are beyond reproach). So, couldn't the friend of KA take the untoward consequences of impurity as a reason to reject SOE?

I don't think they could. True: both KA and SOE are needed to complete the argument for extreme epistemic pragmatism, and it is extreme epistemic pragmatism that generates the untoward consequences. But it is clear that the principle primarily responsible for the untoward consequences is KA, not SOE. It is KA that injects impurity in the epistemic realm, SOE just spreads it around. SOE is a highly plausible principle about second-order evidence that has strange consequences only in the presence of KA. On the other hand, KA has strange consequences (of the same kind as the conjunction of KA and SOE) all by itself.

Notice also that someone who, like Fantl and McGrath, wishes to retain KA but reject extreme epistemic pragmatism (and, thus, reject SOE) will be forced to say that a subject can have more evidence for p than she does for the proposition that she knows that p . If the subject has some evidence that it would be irrational for her to act as if p , then she *thereby* has some evidence against the proposition that she knows that p , but she doesn't *thereby* have some evidence against p . If everything else is evidentially equal, then she will have more evidence for p than she does for the proposition that she knows that p . Maybe it could be argued that this is as it should be, but it is clearly a *prima facie* cost of this combination of views.

So, if the consequences of the conjunction of KA and SOE are unbearable, there are strong reasons to think that it is KA that has to be rejected. But

the rejection won't come easy. This is not the place to develop a way of coming to terms with the rejection of KA, but I believe that the best way of doing it would involve taking fallibilism seriously. If we do take fallibilism seriously, then we could perhaps come to accept that, even though Jeremy knows that the best thing for him to do is to board the train, Jeremy is not sure that the best thing for him to do is to board the train, and so it is rational (given the potential payout) for him to make further inquiries.

8 Conclusion

The dialectical situation is the following. Fantl and McGrath argue for moderate epistemic pragmatism on the basis of KA. But the argument goes through only if extreme epistemic pragmatism is false. Therefore, the most that can be said for moderate epistemic pragmatism is that if KA is true and extreme epistemic pragmatism is false, then it is true. But this argument is unstable, because, given KA, the very plausible SOE entails extreme epistemic pragmatism. Thus, by entailing extreme epistemic pragmatism, KA and SOE together entail that pragmatic encroachment goes all the way down. I take the upshot of the discussion to be that we must either reject KA, or else admit that extreme epistemic pragmatism is true. As in many other areas, so too in epistemology: purity is not easily lost only partially.

Juan Comesaña
Department of Philosophy
The University of Arizona
213 Social Sciences
1145 E. South Campus Drive
P. O. Box 210027
Tucson, AZ 85721-0027 USA
E-mail: comesana@arizona.edu

References:

- Cohen, Stewart. 1988. "How To Be A Fallibilist." *Philosophical Perspectives* 2: 91–123. <http://dx.doi.org/10.2307/2214070>.
- DeRose, Keith. 1995. "Solving the Skeptical Problem." *The Philosophical Review* 104 (1): 1–52. <http://dx.doi.org/10.2307/2186011>.
- Fantl, Jeremy. 2003. "Modest Informatism." *Canadian Journal of Philosophy* 33 (4): 537–562.
- Fantl, Jeremy and Matthew McGrath. 2002. "Evidence, Pragmatics, and Justification." *The Philosophical Review* 112 (1): 47–67. <http://dx.doi.org/10.1215/00318108-111-1-67>.
- Fantl, Jeremy and Matthew McGrath. 2007. "On Pragmatic Encroachment in Epistemology." *Philosophy and Phenomenological Research* 75 (3): 558–89. <http://dx.doi.org/10.1111/j.1933-1592.2007.00093.x>.

Acknowledgements Previous versions of this paper were presented at the University of Arizona, the Second Midwest Epistemology Conference, and the 2009 meeting of the Pacific APA. Many thanks to the audiences at those occasions and to Alan Sidelle and Elliott Sober for very helpful comments and criticisms. Special thanks to Jeremy Fantl, Matt McGrath and Carolina Sartorio, each one of whom read several versions of this paper and made it considerably better.

- Fantl, Jeremy and Matthew McGrath. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121–123. <http://dx.doi.org/10.1093/analys/23.6.121>.
- Hawthorne, John and Jason Stanley. 2008. "Knowledge and Action." *The Journal of Philosophy* 105 (10): 571–590.
- Hetherington, Stephen. 2001. *Good Knowledge, Bad Knowledge*. Oxford: Oxford University Press.
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (4): 549–567. <http://dx.doi.org/10.1080/00048409612347521>.
- Pollock, John. 1986. *Contemporary Theories of Knowledge*. Totawa: Rowman and Littlefield.
- Quine, W. V. O. 1987. *Quiddities*. Cambridge, MA: Harvard University Press.
- Sorensen, Roy. 1988. *Blindspots*. Oxford: Oxford University Press.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.