

EPISTEMIC PRAGMATISM: AN ARGUMENT AGAINST MODERATION

Juan Comesaña

1 INTRODUCTION

A specter is haunting epistemology—the specter of pragmatism. Epistemic pragmatism in general is 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. Two varieties of epistemic pragmatism will occupy us in what follows. According to what I shall call *moderate* epistemic pragmatism, how much justification 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.

In sections 2 and 3, I introduce the two kinds of pragmatism to be discussed and compare them with more traditional epistemological theories. As we’ll see in section 4, moderate epistemic pragmatism can be motivated by plausible arguments. However, as we’ll see in section 7, extreme epistemic pragmatism has some very implausible consequences. It would be good, therefore, if one could be a *merely* moderate epistemic pragmatist—that is, it would be good if one could hold on to moderate epistemic pragmatism while denying extreme epistemic pragmatism. In sections 5 and 6 I argue, however, that that intermediate position is unstable: given a plausible epistemic principle, moderate epistemic pragmatism entails extreme epistemic pragmatism. Therefore, either even moderate epistemic pragmatism is false, or else extreme epistemic pragmatism is true.

2 MOVING THRESHOLDS AND SHIFTY CONTEXTS

Justification comes in degrees, knowledge doesn't.¹ You can be more or less justified in believing 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 justification 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 in Foxboro you have to be justified in believing it to a high degree.² How high? We can ask the same question using a slightly different terminology. Say that when a subject's justification for believing that p is sufficient for knowledge,³ then that subject's justification for believing that p has crossed the "knowledge threshold." The question now is: how high is the knowledge threshold?

We can partition the space of 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 degree of justification that is needed for knowing any proposition in any context. We can factor whether we have knowledge-level justification for believing that p into two components: how much justification *we have* for believing that p , and how much justification *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 justification depends exclusively on how much justification 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

¹Two complications. First, some people (Lewis (1996), the entry for "Knowledge" in Quine (1987), Hetherington (2001), Alan Sidelle in personal communication) think that knowledge does come in degrees. But even they must admit that there is a level of justification below which you don't count as knowing, and the question that sets up our problem is about the location of this lower bound. Second, there is also a notion of justification that doesn't come in degrees—if you are justified to a sufficiently high degree, then you are justified, period. Knowing entails being justified in this binary sense (and, indeed, that is why knowledge itself is binary).

²Some people (for instance, Lewis (1996) again) think that justification is not (always) necessary for knowledge. I would argue (but not here) that this position stems from an overly intellectualistic conception of justification. 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 be replaced by the notion of justification throughout.

³That is, if the subject doesn't know that p it is not because he lacks justification for believing that p . Some authors say that, in that situation, the subject has "knowledge-level" justification for believing that p .

Foxboro or is instead the “express.” Just a moment ago both Matt and Jeremy heard the announcement that the train does indeed stop in Foxboro. Matt says to himself: “Oh, that’s right: they just announced that the train stops in Foxboro, so I know that it does”, 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 there that will determine the course of the rest of his life. Jeremy says to himself: “They just announced that the train stops in Foxboro, but the meeting is so important that I better double-check with an agent. Until I do, I do not know that the train stops in Foxboro.”

Fixed-threshold theories entail that one of Matt or Jeremy is wrong. 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 be justified in telling 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 are justified to the same degree in believing that the train stops in Foxboro, Jeremy’s knowledge threshold is such that this amount of justification is not enough for knowledge, whereas Matt’s knowledge threshold is such that that very same amount of justification is enough for knowledge.

One influential moving-threshold view is contextualism.⁴ 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 in the scale, and there are more strict knowledge relations that locate the knowledge-threshold higher in the scale. A

⁴Three 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.

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 first a "high-standards" context and the second a "low-standards" one).⁵ Let's also suppose that, in each of these contexts, someone says "Matt knows that the train stops in Foxboro" and someone else says "Jeremy knows that the train stops in Foxboro." Tables 1 and 2 represent what a contextualist and a traditional invariantist (non-skeptic) theories would say about the truth-values of those utterances.

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 there to be invariantist moving-threshold theories. According to this kind of theories, the semantic value of "knows" is invariant across contexts (the invariantist part), but the degree of justification a subject must have in order for him to know a certain proposition may vary from case to case. Table 3 represents that kind of position.

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

Table 1: Contextualism

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

Table 2: Traditional invariantism

⁵Different contextualist theories will differ in the identification of the mechanisms by which certain features are made contextualist salient.

	Matt knows	Jeremy knows
High standards	True	False
Low standars	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 are justified to the same degree in believing the same 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 for there to be moving-threshold invariantist theories. Indeed, one of the versions of epistemic pragmatism that occupies me below is just such a theory.⁶

3 MODERATE AND EXTREME EPISTEMIC PRAGMATISM

Like knowledge, practical rationality doesn't come in degrees. Either you are rational to act as if the train stops in Foxboro or you are not.⁷ Like knowledge, too, practical rationality is connected to justification: in general, in order for you to be rational to act as if the train stops in Foxboro you have to be justified to *some* degree in believing it.⁸ How high a degree of justification must you have? In this case, we have a well-developed answer: as high as required by your preferences. Matt's practical concerns are such that he has enough justification for believing that the train stops in Foxboro to make it rational for him to act as if the train stops in Foxboro. Jeremy's practical concerns, however,

⁶Is it also possible for there to be a fixed-threshold, invariantist 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.

⁷You 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.

⁸I 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 justification at all for believing that p is needed for you to be rational in acting as if p . See also note 9.

are such that his level of justification for believing that the train stops in Foxboro (which is, remember, the same as Matt's) is *not* enough to make it rational for *him* to act as if the train stops in Foxboro.

So there are two different binary notions, knowledge and practical rationality, both of which are connected to a degree notion, justification. We have a relatively clear understanding of how practical rationality is connected to justification, and no clear understanding of how knowledge is related to justification. Why not explain the unclear by the clear, at least partially? Why not say that, in order to know that p , the degree to which you are justified in believing p has to be at least high enough to make it rational for you to behave as if p ?⁹ That thought can be captured in the following principle, which we will take as the official definition of moderate epistemic pragmatism:

Moderate epistemic pragmatism: For any possible subject S , if S knows that p , then S is rational to act as if p .¹⁰

Under the assumption that knowledge entails justification, moderate epistemic pragmatism amounts to the claim that the knowledge-threshold depends on your preference structure.¹¹ If there isn't much riding, for instance, on whether there is a glass of water on the table, then moderate epistemic pragmatism 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. 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 and the liquid in the glass is the only available liquid. In that case, moderate epistemic pragmatism entails that, in order to reach

⁹I say "at least high enough" and not "just as high" because it is not plausible to suppose that if your justification for p is high enough 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 justification 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 justification for believing that p is necessary for you to be rational to act as if p . But some (not very low) level of justification is always necessary for you to know that p .

¹⁰This is the principle that Fantl and McGrath (2002) call "PCK" (for "pragmatic condition on knowledge"). Stanley (2005), Hawthorne (2004) and Hawthorne and Stanley (forthcoming) are, to different degrees, also friendly to pragmatic encroachment, although they don't explicitly endorse moderate or extreme epistemic pragmatism.

¹¹Strictly speaking, moderate epistemic pragmatism doesn't entail that the knowledge-threshold varies according to your practical situation—it is compatible with moderate epistemic pragmatism to say that the knowledge-threshold is fixed, but that there is an *independent* condition on knowledge according to which you can know a proposition only if it is rational for you to act as if it is true. It is not clear what the motivation would be for holding this view, though, and so I will continue assuming that the moderate epistemic pragmatist thinks that the knowledge-threshold varies with the practical situation of the subject.

the knowledge-threshold, your justification for thinking that there is water in the glass needs to be very high indeed. How high? As high 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 high as it is necessary to make it rational for you to drink the liquid in the glass. Moderate epistemic pragmatism is thus an invariantist, fixed-threshold theory, a theory of the kind represented by Table 3 above.

Now, moderate epistemic pragmatism is of course compatible with the claim that how justified you actually are in believing the proposition 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 level of justification 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 level of justification 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 justified you are in believing 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 are thereby less justified in 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 except for the fact that whether p is true is more important to S than to S', then S' has more justification for believing p than S does.

Where does extreme epistemic pragmatism fit 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 knowledge entails justification, 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 justificational status. 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 are justified to the same degree in believing that the train stops in Foxboro—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 justificational 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 is just as justified as Jeremy, 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 is more justified than Jeremy.¹² 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 justification 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 justification 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 than moderate epistemic pragmatism, but rather on the fact that it claims that our practical situation can by itself affect *how justified we are*, and not merely how justified we have to be in order to have knowledge-level justification.

Fantl and McGrath (forthcoming) have argued, on the basis of moderate epistemic pragmatism, against the following principle:

Purism: For any two subjects S and S', necessarily, if S and S' have the same amount of justification for believing that *p*, then S is in a position to know that *p* iff S' is, too.¹³

¹²If, on the other hand, we retained the assumption that Matt and Jeremy have the same justification, 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).

¹³Fantl and McGrath formulate purism in terms of evidence instead of doing it directly in terms

Fantl and McGrath call their target principle “Purism” because it allegedly defines a dimension of *pure* epistemic evaluation, free from any encroachment of practical matters. In a standard terminology, the principle states that being in a position to know supervenes on evidence. What is it to be in a position to know a certain proposition? It is for the degree to which you are justified in believing that proposition to be high enough for knowledge—that is, it is for you to have knowledge-level justification.

But notice that purism deserves its name only under the assumption that extreme epistemic pragmatism is false. For suppose that it was true. Then, purism might well be true even if practical matters do influence whether you are in a position to know.

Only against the background of the falsity of extreme epistemic pragmatism, then, is an argument against purism an argument for the encroachment of practical matters into epistemology. Fantl and McGrath, then, clearly want to be *merely* moderate epistemic pragmatists—that is, they want to hold on to moderate epistemic pragmatism while denying extreme epistemic pragmatism. But I will argue that this combination of views is untenable. Let us first turn, though, to Fantl and McGrath’s argument for moderate epistemic pragmatism, and how they think that it can be brought to bear against purism.

4 THE ARGUMENT AGAINST PURISM

Fantl and McGrath don’t take moderate epistemic pragmatism 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 don’t stop at the supermarket to buy yams the day before Thanksgiving and my wife asks me why, I can sensibly reply “Because I know that there are yams at home.” This practice of ours of appealing to what we know in defense of what we and others do makes sense if we are (perhaps implicitly) aware that moderate epistemic pragmatism is true. Second, Fantl and McGrath also argue that a plausible closure inference entails moderate epistemic pragmatism with the aid of the following knowledge-to-action principle:

KA: 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.

of justification. But given that they are assuming an evidentialist conception of justification, their principle is equivalent to the one presented here.

KA really is extremely 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

In their argument against Purism, Fantl and McGrath appeal not only to moderate epistemic pragmatism, but also to a version of fallibilism:

Fallibilism: It is possible for a subject S to know that p even if there is a subject S' such that S' is justified in believing that p to a higher degree than S is.

I won't dispute fallibilism here. Fantl and McGrath are right to point out that if the choice is between purism and fallibilism, then that is a very good argument against purism.¹⁶ Let us now see how Fantl and McGrath use fallibilism, together with moderate epistemic pragmatism, to argue against purism.

The argument against purism can be illustrated with our familiar case: Matt knows that the train stops in Foxboro, Jeremy has the same justification that Matt has for believing that the train will stop in Foxboro, 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 is just as justified as Matt is and Matt knows that the train stops in Foxboro. Given the negation of fallibilism, this justification must be maximal—it is impossible to be better justified in believing that proposition, in which case it is arguably impossible for it to be irrational for you to act as if p . But, given fallibilism, the argument is now the following:

1. Matt knows that the train stops in Foxboro.
2. Jeremy is not rational to act as if the train stops in Foxboro.
3. Jeremy and Matt have the same justification for believing that the train stops in Foxboro.
4. Jeremy is not in a position to know that the train stops in Foxboro. (From 2 and moderate epistemic pragmatism)
5. Purism is false. (From 1, 3 and 4).

¹⁴Except, as Fantl and McGrath note, in those cases where whether S does A causally affects whether p is true.

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

¹⁶Hawthorne and Stanley (forthcoming) argue that "knowledge delivers probability 1." It is not clear whether they mean to deny fallibilism.

The argument is valid. The step from 2 and moderate epistemic pragmatism to 4 requires some comment. All that immediately follows from 2 and moderate epistemic pragmatism is that Jeremy doesn't know that the train stops in Foxboro. 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 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 deny the move from 2 and moderate epistemic pragmatism 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. But what is not merely a stipulation about the case is premise 3. That Jeremy and Matt have the same justification is a substantial claim about the case—it is a claim that depends on the denial of extreme epistemic pragmatism. But I will show that moderate epistemic pragmatism entails extreme epistemic pragmatism given a plausible assumption. Therefore, if Fantl and McGrath are right about moderate epistemic pragmatism, then they are wrong about purism.

5 A PRINCIPLE OF SECOND-ORDER EVIDENCE

Here is a quick argument for the claim that moderate epistemic pragmatism entails extreme epistemic pragmatism: Propositions can only be justified for a subject by propositions *known* by that subject. Therefore, if what is known is affected by the practical situation of the subject, what he is justified in believing is similarly affected by his practical situation—in other words, moderate epistemic pragmatism entails extreme epistemic pragmatism. There are interesting questions about the validity of this argument, but I won't pursue them here. I have argued elsewhere¹⁷ that the key premise in this argument, namely, that only propositions that a subject knows can justify him in believing further propositions,¹⁸ is false. But there is a better argument for the same conclusion.

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,

¹⁷Comesaña and Kantin (forthcoming).

¹⁸A claim defended by Williamson (2000) under the name of E=K.

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 justification:

SOE: For any two 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 is justified in believing that his belief that p amounts to knowledge to a higher degree than S' is justified in believing that her belief that p amounts to knowledge,
 THEN
 S is more justified in believing that p than S' is.

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 justification for believing in Goldbach’s conjecture than Joe does. In virtue of being more justified than Joe in 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* can be justified in two steps. First, if S is justified in believing that her belief that p

amounts to knowledge to a higher degree than S' is in believing that his belief that p amounts to knowledge, then S has more evidence for the proposition that she knows that p than S' does for the proposition that he knows that p . Second, if S has more evidence for the proposition that her belief that p amounts to knowledge than S' does for the corresponding proposition that his belief that p amounts to knowledge, and if S and S' are evidential cousins, then S has more justification for believing p than S' does.

Is the Gettier problem relevant to the truth of SOE? It might be thought that it is, for the following reason. 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 , and yet only one of them knows that p .¹⁹ If so, can't then subjects have evidence that they don't know a proposition *because* they have evidence that they don't satisfy some of the extra-evidential conditions on knowledge? In that case, wouldn't SOE be false? The answers are, respectively, "Yes" and "No."

Let us first recall one traditional Gettier case: Jones is justified in believing that Nogot owns a Ford (even though he 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*. Because what Jones's friend is telling Jones is that Jones's 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 he doesn't know it. 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

¹⁹This 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 purism, means that the level of justification for the belief in question is high enough for knowledge.

case is something that cannot be learned by that subject.²⁰

But Jones can have evidence that his belief doesn't satisfy some of the extra-evidential conditions the existence of which Gettier showed to be necessary for knowledge 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." 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 acquires some evidence against the proposition that someone in his office owns a Ford. It is not easy to say, without having an answer to the Gettier problem, what kind of evidence that is,²¹ but that he does acquire some evidence against his belief is revealed by the fact that Jones is less justified in believing that someone in his office owns a Ford after hearing his friend. Therefore, Jones and Smith are not epistemic cousins, because Jones has some evidence against his belief that Smith lacks. Therefore, Gettier cases don't represent a threat to SOE.

6 THE ARGUMENT FOR EXTREME EPISTEMIC PRAGMATISM

I will now argue that SOE and moderate epistemic pragmatism together entail extreme epistemic pragmatism. Therefore, Fantl and McGrath's argument against purism cannot be sound: if they are right about moderate epistemic pragmatism, then the third premise of that argument is false.

Recall the arguments for moderate epistemic pragmatism. One of them was that we (at least implicitly) rely on the truth of moderate epistemic pragmatism when justifying and criticizing actions. The second argument for moderate epistemic pragmatism relied on KA, and Fantl and McGrath's defense of KA appeals both to its *obvious truth* as well as to its role in the criticism and justification of action. If that is so, then realizing that we are not rational in acting as if p will be evidence that we don't know that p . Remember too 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. Therefore, whenever it is not rational for us to act as if

²⁰That is, "I am in a Gettier-situation with respect to p " is a *blindspot*—see Sorensen (1988).

²¹A defeasibility theorist, for instance, could say that Jones acquires an undercutting defeater for his belief.

p we will have evidence that we don't know that p that we would lack if it were rational for us to act as if p . Therefore, anyone who is not rational to act as if p is less justified in believing that she knows that p than someone who is rational in acting as if p (but who otherwise has the same evidence for and against p). In other words, anyone who is rational in acting as if p is more justified in thinking that she knows that p than someone who is not rational in acting as if p (but who is otherwise in the same evidential situation regarding p). Therefore, given SOE, anyone who is rational in acting as if p has more justification for p than someone who is not rational in acting as if p (but who is otherwise in the same evidential situation).

Let us consider once again the case of Jeremy and Matt. Remember that Matt is rational to act as if the train stops in Foxboro, whereas Jeremy isn't. If so, then Jeremy (but not Matt) has some evidence that he is not rational to act as if the train stops in Foxboro. Therefore, given Jeremy's at least implicit awareness of the truth of moderate epistemic pragmatism, Jeremy (but not Matt) has some evidence that he doesn't know that the train stops in Foxboro. Therefore, Matt is more justified in thinking that he knows that the train stops in Foxboro than Jeremy is. Therefore, by SOE, Matt has more justification for believing that the train stops in Foxboro than Jeremy does. Therefore, premise 3 in the argument against purism fails, as does any analogous premise about structurally similar cases.²²

7 CONSEQUENCES OF IMPURITY

If Fantl and McGrath are right, then moderate epistemic pragmatism is true. If I am right, then if moderate epistemic pragmatism 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 the case are such that Matt and Jeremy both believe justifiedly what the chief says, and on that basis also justifiedly believe that the train stops in Foxboro if and only if the coin landed heads. Because Matt doesn't much care whether the train

²²Notice that SOE is compatible with the position that we can be justified in believing p while also being justified in believing that we don't know that p . For instance, SOE is compatible with the position that "lottery propositions" (such as the proposition that a ticket in a fair lottery is a loser) can be justifiably believed but not known, a position recently defended by Hawthorne (2004).

stops in Foxboro or not, and given the evidence he has, he is justified more or less to the same degree in believing that it does stop as he is in believing that it doesn't (and this reflects the fact that he is justified more or less to the same degree in believing that the coin landed heads as he is in believing 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 justification for believing that the train stops in Foxboro than Matt does. So far, this is a direct consequence, however unpalatable, of extreme epistemic pragmatism. But remember that Jeremy justifiably believes that the train stops in Foxboro if and only if the coin landed heads. Therefore, extreme epistemic pragmatism has the consequence that Jeremy is justified in believing that the coin landed tails to a higher degree than he is justified in thinking 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 justification 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 moderate epistemic pragmatism and SOE (I assume that the other propositions needed for the entailment are beyond reproach). So, couldn't the friend of moderate epistemic pragmatism take the untoward consequences of impurity as a reason to reject SOE?

I don't think they could. True: both moderate epistemic pragmatism 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 moderate epistemic pragmatism, not SOE. It is moderate epistemic pragmatism 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 moderate epistemic pragmatism. On the other hand, moderate epistemic pragmatism has strange consequences (of the same kind as the conjunction of moderate epistemic pragmatism and SOE) all by itself.

Notice also that someone who, like Fantl and McGrath, wishes to retain moderate epistemic pragmatism but reject extreme epistemic pragmatism (and, thus, reject SOE) will be forced to say that a subject can be justified in believing that p to a higher degree than she is justified in believing 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 be more justified in believing that p than

she is in believing 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 moderate epistemic pragmatism and SOE are unbearable, there are strong reasons to think that it is moderate epistemic pragmatism 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 moderate epistemic pragmatism, 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 Matt knows that the best thing for him to do is to decline the bet, Matt is not sure that the best thing for him to do is to decline the bet, and so it is rational (given the potential payout) for him to take the bet.

8 CONCLUSION

The dialectical situation is the following. Fantl and McGrath argue against purism on the basis of moderate epistemic pragmatism. But the argument goes through only if extreme epistemic pragmatism is false. Therefore, the most that can be said against purism is that if the moderate, but not the extreme, variety of epistemic pragmatism is true, then it is false. But this argument is unstable, because, given moderate epistemic pragmatism, the very plausible SOE entails extreme epistemic pragmatism. Thus, purism is after all compatible with moderate epistemic pragmatism, given SOE. It is obvious, though, that this defense of the principle called "Purism" is not really a defense against pragmatic encroachment. On the contrary, by entailing extreme epistemic pragmatism, moderate epistemic pragmatism 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 moderate epistemic pragmatism, 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.

REFERENCES

- Cohen, S. (1988) "How To Be A Fallibilist," *Philosophical Perspectives* :2, 91–123.
- Comesaña, J., and H. Kantin (forthcoming) "Is Evidence Knowledge?" *Philosophy and Phenomenological Research*.
- DeRose, K. (1995) "Solving the Skeptical Problem," *The Philosophical Review* 104:1, 1–52.

- Fantl, J., and M. McGrath (2002) "Evidence, Pragmatics, and Justification," *Philosophical Review* 112:1, 47–67.
- Fantl, J., and M. McGrath (forthcoming) "Knowledge and the Purely Epistemic: A Defense of Pragmatic Encroachment," *Philosophy and Phenomenological Research*.
- Gettier, E. (1963) "Is Justified True Belief Knowledge?" *Analysis* 23, 121–123.
- Hawthorne, J. (2004) *Knowledge and Lotteries*, Oxford University Press.
- Hawthorne, J., and J. Stanley (forthcoming) "Knowledge and Action," *Journal of Philosophy*.
- Hetherington, S. (2001) *Good Knowledge, Bad Knowledge*, Oxford University Press.
- Lewis, D. (1996) "Elusive Knowledge," *Australasian Journal of Philosophy* :74, 549–567.
- Quine, W. V. (1987) *Quiddities*, Belknap/Harvard.
- Sorensen, R. A. (1988) *Blindspots*, Oxford University Press.
- Stanley, J. (2005) *Knowledge and Practical Interests*, Oxford University Press.
- Williamson, T. (2000) *Knowledge and Its Limits*, Oxford University Press.