

Believing, Knowing, Acting

by

Daniel Blake Roeber

A Dissertation submitted to the

Graduate School-New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

written under the direction of

Ernest Sosa

and approved by

New Brunswick, New Jersey

May, 2013

ABSTRACT OF THE DISSERTATION

Believing, Knowing, Acting

By Daniel Blake Roeber

Dissertation Director:

Ernest Sosa

Interest-relative invariantism conjoins the interest-relativist thesis that knowledge depends in part on our interests with the invariantist thesis that ‘knows’ is not a context-sensitive word. Neither thesis entails the other, and interest-relativism is interesting in its own right. If interest-relativism is true, then knowledge depends in part on truth-irrelevant factors, since our interests will often be irrelevant to our grip on the truth. In Chapter 1, I steer the debate away from the invariantist thesis that ‘knows’ is not a context-sensitive word and toward the interest-relativist thesis that knowledge depends in part on our interests. Interest-relativism comes in two varieties: what I call ‘pragmatism’ and ‘intellectualism.’ Pragmatism is the view that *practical* interests can make a difference to knowledge, while intellectualism is the view that *intellectual* interests can make a difference to knowledge. Pragmatism and intellectualism might both be true, but neither view entails the other. While pragmatism has received considerable attention in the literature, intellectualism has scarcely been identified as a position in logical space, and it has no defenders. As a result, many philosophers think they can resist interest-relativism by simply resisting pragmatism. In Chapters 2 through 7, I show otherwise. First, in Chapters 2 through 6, I argue extensively against pragmatism. Then, in Chapter

7, I argue for intellectualism. Since purism is just the denial of pragmatism, and since intellectualism is a species of interest-relativism, Chapters 2 through 7 jointly defend *purist* interest-relativism. Knowledge depends in part on our interests, but not on our *practical* interests, I argue. Along the way, I sketch a theory of belief, apply this theory of belief to questions about the value of knowledge, and say how believing that p relates to one's credence that p .

Table of Contents

Abstract	ii
Chapter 1: Gricean Responses to High-Stakes Cases	1
The Rysiew/Brown WAM	6
A Worry about the Generality of Brown's WAM	10
A WAM with General Application	12
Conversational Implicatures and Cancelability	16
The Limits of WAMs	18
The Limits of Gricean Responses in General	25
Conclusion	27
Chapter 2: Pragmatism and the Nature of Belief	30
Ganson and Weatherson on Belief	31
Ross and Schroeder on Belief	41
A Close Neighbour of Ross and Schroeder's Theory of Belief	49
Fantl and McGrath on Belief	52
Rival Theories of Belief	56
Appendix 1	57
Appendix 2	63
Chapter 3: A Paradox for Justification, Knowledge and Practical Rationality	68
Preliminaries	69
A Purist Response to High Stakes	72
A Paradox for Justification, Knowledge and Practical Rationality	108
Responding to the Paradox	115
Pragmatism Unmotivated?	116
Chapter 4: Problems for Pragmatism	120
Introduction	120
A Counterexample to the Knowledge-Preference Principle	125
A Dilemma for Principle-Based Arguments	130
Case-Based Arguments	143

A Challenge for Purism	147
Conclusion	149
Chapter 5: Advice to Minimalists	151
The Rysiew/Brown WAM	154
A Worry about the Generality of Brown's WAM	158
A WAM with General Application	160
Conversational Implicatures and Cancelability	164
The Limits of WAMs	166
The Propriety of Keith's Utterance in High Stakes	173
Conclusion	176
Chapter 6: The Nature of Belief and the Value of Knowledge	177
The Threshold View	179
The Affirmative View of Belief	181
The Threshold View Defended	184
Is Knowledge Distinctively Valuable?	192
Back to the Affirmative View (or Something Near Enough)	194
The Platitudes <i>versus</i> the Distinctive Value of Knowledge	200
The Nature of Belief and the Valuable of Knowledge	203
The (Previously) Mysterious Interval $m < x < n$	209
Appendix	211
Chapter 7: Knowledge and Intellectual Interests	219
Interests, Practical <i>versus</i> Intellectual	222
Three Claims about Rationality	226
An Argument for Intellectualism	232
Objections and Replies	237
Concluding Remarks	246
References	248

Chapter 1: Gricean Responses to High-Stakes Cases

Suppose you go from a situation where it doesn't matter much whether p is true to a situation where it matters a lot whether p is true. Does this change in your practical interests *vis-à-vis* p make any difference to the strength of epistemic position required for you to know that p , or to the strength of epistemic position required for you to express a true proposition by uttering the sentence 'I know that p '?

The pragmatic encroachment debate pits contextualists against pragmatists, and minimalists against both of them. Pragmatism is the view that knowledge depends at least in part on our practical interests, in the sense that a mere difference in practical interests can entail a difference in knowledge.¹ If pragmatism is true, there will be pairs of cases that differ only insofar as some stipulated difference in practical interests requires that they differ, and where, as a result of this mere difference in practical interests, a difference in knowledge follows. If you go from a situation where you know that it doesn't matter much whether some proposition p is true to a situation where you know that it matters a lot whether p is true, this change in your practical interests increases both the strength of epistemic position required for you to know that p and the strength of epistemic position required for you to express a true proposition by uttering the sentence 'I know that p .' Pragmatism is a theory

¹ The label 'pragmatism' comes from Fantl and McGrath (2009). Though perhaps misleading, it does allow us to easily distinguish between the pragmatist thesis that knowledge depends at least in part on our practical interests and the *invariantist* thesis that the word 'knows' *does not* express different properties or relations in different contexts of utterance. (The more popular labels 'subject-sensitive invariantism' and 'interest-relative invariantism' both lack this desirable feature.) For representative examples of the view that I am calling 'pragmatism,' see Fantl and McGrath (2002), Hawthorne and Stanley (2008), Ganson (2008), Fantl and McGrath (2009), Ross and Schroeder (2012), Weatherson (forthcoming), and the portions of Hawthorne (2004) and Stanley (2005) explicitly about knowledge and practical interests.

about *knowledge*. Contextualism, in contrast, is a theory about the word ‘knows.’ According to contextualism, the word ‘knows’ can semantically express different properties or relations in different contexts of utterance, and (thus) a sentence containing the word ‘knows’ can semantically express different propositions in different contexts of utterance.² Contextualism entails that, if you and I occupy different contexts when we utter some sentence containing the word ‘knows,’ you might semantically express a true proposition while I semantically expresses a false one, even if we utter this sentence at exactly the same time. Given that contextualism is true, if you go from a situation where you that know that it doesn’t matter much whether *p* is true to a situation where you know that it matters a lot whether *p* is true, this change in your practical interests might increase the strength of epistemic position required you to express a true proposition by uttering the sentence ‘I know that *p*’ *without* increasing the strength of epistemic position required for you to know that *p*.³

What motivates contextualism and pragmatism, respectively? Contextualism evolved as both a solution to the so-called “skeptical paradox” and an explanation of our intuitions about the knowledge ascriptions and denials in cases like the following.⁴

Low Stakes: Rachel and Keith are driving home on Friday afternoon. They plan to stop at the bank to deposit their paychecks, but as they approach it they notice that the lines inside are very long. Although they generally like to deposit their paychecks as soon as possible, it’s not especially important in this case that they deposit their checks right away, so Keith suggests that they drive straight home and deposit their

² See, for example, Cohen (1988), DeRose (1995), and Lewis (1996).

³ On this point, see DeRose (2009), especially chapter 6.

⁴ See DeRose (1995), p. 39.

checks tomorrow morning. Rachel says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays.” Keith replies, “No, I know it will be open. I stop in at the bank every Saturday. It’s open until noon.”

High Stakes: Rachel and Keith are approaching their bank on Friday afternoon, as in Low Stakes, and they notice the long lines. Keith again suggests that they go home and return to deposit their paychecks tomorrow morning, explaining that he stops at the bank every Saturday morning and that it’s open until noon. But in this case, Rachel and Keith have just written a very large and important check. If their paychecks are not deposited before Monday morning, the important check they wrote will bounce, leaving them in a terrible situation. The bank isn’t open on Sunday. Rachel reminds Keith of these facts. She then says, “What if our bank discontinued its Saturday hours since your last visit? Do you know that it will be open?” Remaining just as confident as he was before that the bank will be open, Keith replies, “Well, no, I don’t know that the bank will be open. We had better stop and deposit the checks now.”⁵

While contextualism evolved as an explanation of our intuitions about the knowledge attributions and denials in cases like these, pragmatism evolved as both a solution to the lottery paradox and a *competing* explanation of our intuitions about cases like Low and High Stakes.⁶ So, while pragmatism and contextualism are logically consistent—pragmatists might

⁵ These cases are slight modifications of DeRose’s cases in DeRose (1992), p. 913.

⁶ See Hawthorne (2004), pp. 2-3, for the lottery paradox. Stanley (2005) contains the paradigm pragmatist treatment of High Stakes.

be right about *knowledge* while contextualists are right about the word ‘knows’—pragmatism and contextualism are standardly thought of as rivals.⁷

Neither contextualism nor pragmatism is universally accepted among epistemologists, however. Where ‘purism’ names the denial of pragmatism, and ‘invariantism’ names the denial of contextualism, many epistemologists are purist invariantists. Purist invariantism entails (a) that Keith knows that the bank will be open in Low Stakes *and* High Stakes if he knows that the bank will be open in either of them, and (b) that Keith would semantically express a true proposition by uttering ‘I know that the bank will be open’ in both cases if he would semantically express a true proposition by uttering this sentence in either case. But (a) and (b) are consistent with any level of skepticism, and few purist invariantists are sufficiently skeptical to reject the stipulation that Keith knows that the bank will be open in Low Stakes.⁸ Where ‘moderate purist invariantism’ names purist invariantism in conjunction with anti-skepticism, most epistemologists are moderate purist invariantists. Since moderate purist invariantists think we can get by in our epistemological theorizing with neither the contextualist thesis that ‘knows’ can express different properties in different contexts of utterance nor the pragmatist thesis that knowledge depends at least in part on our practical interests, I will call moderate purist invariantism, simply, ‘minimalism.’

This paper concerns a specific attempt to defend minimalism from the contextualist or pragmatist conclusions that we might naturally draw from Low and High Stakes, and other similar cases. In Low Stakes, Keith utters the sentence ‘I know that the bank will be open,’

⁷ They clearly *are* rivals in at least the sense that contextualists and pragmatists can’t both be right about all of the purported motivations for accepting their respective views. Low and High Stakes can’t motivate *both* contextualism *and* pragmatism, for example.

⁸ Wayne Davis may be an exception. See, for example, Davis (2007).

and then, in High Stakes, he utters the sentence ‘I don’t know that the bank will be open.’ Call these sentences ‘K’ and ‘ \neg K,’ respectively. Minimalists think that Keith says something false by uttering \neg K in High Stakes.⁹ Pragmatists and contextualists both think that he says something *true* by uttering \neg K in High Stakes. Everyone agrees that Keith *appears* to say something true by uttering \neg K in High Stakes, at least *prima facie*. The disagreement concerns the correct explanation of this appearance. As Jeff King and Jason Stanley note, a tempting minimalist explanation takes the form of the following “warranted assertability manoeuvre” (or ‘WAM’ for short):

Keith appears to say something true by uttering \neg K in High Stakes because (a) he *con conversationally implicates* something true by uttering \neg K in High Stakes and (b) we mistake the proposition that he conversationally implicates by uttering \neg K in High Stakes for what he *says* by uttering \neg K in High Stakes.¹⁰

In this paper, I argue that WAMs are insufficient for defending minimalism from cases like High Stakes. Thus, I argue, defending minimalism from cases like High Stakes will require adopting a different strategy. In §1, I outline the WAMs that Patrick Rysiew and Jessica Brown forward in response to High Stakes and, in §2, I show why their WAMs are ultimately unsuccessful. In §§3-4, I discuss Grice’s cancelability test and seek out a WAM

⁹ Like others, I take it for granted that the proposition one semantically expresses by uttering a sentence is identical to what one *says* by uttering that sentence.

¹⁰ See Stanley and King (2007), p. 134. Note that, according to the traditional form of a WAM (the form responsible for the moniker ‘warranted assertability maneuver’), we think Keith says something true by uttering \neg K in High Stakes because we mistake the *warrant* he has for asserting \neg K for the *truth* of the proposition that he asserts by uttering \neg K. Since it is enormously implausible that we are mistaking the warrant that Keith has for asserting \neg K for the truth of the proposition that he asserts that he asserts by uttering \neg K, I will ignore the more traditional form of a WAM. From here forward, I will assume that, if we are mistaking anything at all, we are mistaking a (true) conversationally implicated proposition for a (false) semantically expressed proposition, not a person’s *warrant* for an *assertion* for a *false* semantically expressed *proposition*.

that passes it. After finding a WAM that *may* pass Grice's cancelability test, I show that WAMs are really beside the point—that, for purposes of defending minimalism from cases like High Stakes, no WAM can succeed. WAMs are a species of Gricean response. In §6, after showing that minimalists cannot use WAMs to defend their view from High Stakes, I argue that minimalists cannot use *any* form of Gricean response to defend their view against the problems raised by High Stakes, or other high-stakes cases in the literature. The upshot is that, in order to defend their view in the face of High Stakes and similar cases, minimalists must prioritize the properly epistemological question whether Keith knows that the bank will be open over semantic questions about the sentence 'I don't know that the bank will be open.'

1 The Rysiew/Brown WAM

According to Patrick Rysiew, Keith semantically expresses the same proposition by uttering K in Low Stakes as he would semantically express by uttering K in *any context*: namely, the proposition that he knows that the bank will be open. Likewise for $\neg K$. By uttering $\neg K$ in any context, Keith would semantically express the proposition that he does *not* know that the bank will be open.¹¹ Rysiew endorses a relevant alternatives account of knowledge. On his view, S knows that p iff S can rule out all of the relevant $\neg p$ alternatives—where a relevant $\neg p$ alternative is a scenario where p is false that is likely to obtain.¹² In High Stakes, Keith can rule out all the likely scenarios where the bank will not be open. So, according to Rysiew, Keith knows the bank will be open in High Stakes, and he therefore says something false by

¹¹ Rysiew (2000), p. 487.

¹² Ibid., pp. 487-8.

uttering $\neg K$ in High Stakes. An irrelevant alternative may become salient in a context, however, and Rysiew thinks that an instance of ‘I know that p ’ uttered in such a context would implicate that the speaker can rule out this irrelevant alternative.¹³ When Rachel mentions the possibility that the bank has recently discontinued its Saturday hours, she makes this possibility salient. Given the evidence currently in Keith’s possession, Keith cannot rule this possibility out. But since it is *unlikely* that the bank has recently discontinued its Saturday hours, this alternative is irrelevant. Thus, argues Rysiew, it doesn’t matter that Keith can’t rule this possibility out; he still knows that the bank will be open on Saturday. The possibility that the bank recently discontinued its Saturday hours is still salient in High Stakes, however, so Keith cannot felicitously utter K in High Stakes. If he did, he would implicate that he *can* rule out the possibility that the bank has recently discontinued its Saturday hours. Since Keith knows that he cannot rule this possibility out, he utters $\neg K$ instead. And this way, says Rysiew, Keith implicates that he cannot rule out the possibility that the bank has recently discontinued its Saturday hours. Because the proposition that Keith cannot rule out the possibility that the bank has recently discontinued its Saturday hours is both true and easily mistaken for the proposition that Keith semantically expresses by uttering $\neg K$ —namely, the proposition that Keith does not know that the bank will be open on Saturday—it’s tempting to think that Keith says something true by uttering $\neg K$ in High Stakes. And this, according to Rysiew, is why Keith seems to say something true by uttering $\neg K$ in High Stakes.

¹³ Ibid., p. 490.

This is the gist of Rysiew's WAM. But how, *exactly*, is Keith supposed to implicate that he cannot rule out the possibility that the bank has recently discontinued its Saturday hours by uttering $\neg K$? The mere *salience* of this possibility leaves it mysterious how Keith would have implicated that he can rule out the possibility that the bank has recently discontinued its Saturday hours by simply uttering K . After all, Rachel could have made just about any possibility salient. She could have raised the possibility that Caesar had a lisp, or the possibility that the Browns will win the Superbowl, or the possibility that it is raining in Novosibirsk, or ... you name it. And surely, Keith would not have implicated that he could rule out any of *these* possibilities by uttering K . So salience *alone* does not adequately explain how Keith's uttering K would have implicated that he can rule out the possibility that their bank has recently stopped opening on Saturdays. As DeRose points out, without telling us how general conversational principles, the semantics of the word 'knows,' and the details of High Stakes combine to generate this implicature, Rysiew leaves it unclear that this implicature is really present.¹⁴

Rysiew does try to meet DeRose's challenge, but his response focuses solely on the possibility that the bank has recently discontinued its Saturday hours. It makes no mention of Rachel and Keith's *practical circumstances* in High Stakes. As Jessica Brown points out, however, these practical circumstances matter. They affect our intuitions about Keith's utterance of $\neg K$ in High Stakes.¹⁵ Brown develops Rysiew's account to capture this impact on our intuitions. To do this, she starts with Grice's Maxim of Relation and his well-known

¹⁴ DeRose (2002), p. 176.

¹⁵ Brown (2006).

example of the motorist who tells the pedestrian that she is running out of gas.¹⁶ In Grice's example, the pedestrian utters 'there is a garage nearby' and thereby implicates that there is an *open* garage nearby. The scenario that Brown imagines differs slightly from Grice's. In Brown's scenario, there is only one garage nearby and the pedestrian knows that this garage is *closed*, so she utters 'there is no garage nearby.' Brown thinks that, just as the pedestrian in Grice's petrol example implicates that there is an open garage nearby by uttering just 'there is a garage nearby,' the pedestrian in her example implicates that there *is not* an open garage nearby by uttering just, 'there is no garage nearby.' According to Brown, "[w]hile this utterance is literally false, it pragmatically conveys the true claim that there is no open garage nearby."¹⁷ Brown thinks that an utterance of K or $\neg K$ in High Stakes would, respectively, be analogous to an utterance of 'there is a garage nearby' or 'there is no garage nearby' in her petrol scenario.¹⁸ In High Stakes, Keith would have implicated a false proposition by uttering K . By uttering $\neg K$ instead, he thereby implicated a *true* proposition. Which proposition, exactly? Brown thinks that, had Keith uttered K in High Stakes, he would have implicated that his belief that the bank will be open matches the facts out to the nearest world in which the bank has recently changed its Saturday hours.¹⁹ Let ' m ' name the proposition that Keith's belief that the bank will be open matches the facts out to the nearest world in which the bank has recently changed its Saturday hours. Since Keith uttered $\neg K$ instead of K , says

¹⁶ Grice (1989), p. 32.

¹⁷ Brown (2006), p. 425.

¹⁸ Ibid., p. 426.

¹⁹ Brown is apparently taking it for granted that beliefs are not individuated by their contents. According to her DeRose-inspired account of *matching the facts*, Keith's belief that the bank will be open matches the facts out to the nearest world in which the bank has recently changed its Saturday hours only if, in the nearest world in which the bank has recently changed its Saturday hours, Keith *doesn't* believe that the bank will be open on Saturday. See *ibid.*, p. 424.

Brown, he implicated that $\neg m$ instead of m . And since $\neg m$ is *true*, he implicated a *true* proposition by uttering $\neg K$ in High Stakes. Finally, says Brown, when we read High Stakes and consider Keith's utterance of $\neg K$, his utterance rings true because it conversationally implicates that $\neg m$, and we mistake $\neg m$ for what Keith *says* by uttering $\neg K$.²⁰ But now that we've explained why Keith seems right to utter $\neg K$ in High Stakes, we are free to deny that Keith says something true by uttering $\neg K$ in High Stakes, says Brown.

2 A Worry about the Generality of Brown's WAM

Brown's development of Rysiew's WAM does seem like an improvement, since it makes some progress toward explaining how general conversational principles combine with the semantics of the word 'knows' and the details of High Stakes to generate the supposed implicature. There is an obvious problem with Brown's WAM, however. Brown is correct that Rachel and Keith's practical circumstances and concerns affect our intuitions about Keith's utterance of $\neg K$. Consider the following revision of High Stakes, which focuses *solely* on their practical circumstances and concerns. Unlike the original version of High Stakes, this version does not mention the possibility that their bank has changed its Saturday hours since Keith's last visit.

Practical High Stakes: Rachel and Keith are driving home on Friday afternoon.

They plan to stop at the bank to deposit their paychecks, but as they approach it they notice that the lines inside are very long. Keith suggests that they go home and return to deposit their paychecks on Saturday morning, explaining that he stops at

²⁰ Ibid.

the bank every Saturday morning and that it's open until noon. But in this case, Rachel and Keith have just written a very large and important check. If their paychecks are not deposited before Monday morning, the important check they wrote will bounce, leaving them in a terrible situation. The bank isn't open on Sunday, and Rachel reminds Keith of this fact. She then says, "It will be completely disastrous for us if we go home now and the bank isn't open tomorrow. Do you know that it will be open tomorrow?" Remaining just as confident as he was before that the bank will be open tomorrow, Keith replies, "Well, no, I don't know that the bank will be open tomorrow. We had better stop and deposit the checks now." (As it happens, the possibility that the bank has *changed its Saturday hours since Keith's last visit* never enters Keith or Rachel's mind.²¹)

Just as Keith appears to say something true by uttering $\neg K$ in the original version of High Stakes, he appears to say something true by uttering $\neg K$ in this version of High Stakes. But the explanation for this appearance is certainly not that, by uttering $\neg K$, Keith implicates *that his belief does not match the facts out to the nearest world in which the bank has changed its Saturday hours since his last visit*. After all, in *this* version of High Stakes, we have stipulated that the possibility that the bank has changed its Saturday hours since Keith's last visit is completely out of mind. So it is implausible that, by uttering $\neg K$ in *this* version High Stakes, Keith implicates that his belief does not match the facts out to the nearest world in which the bank has recently changed its Saturday hours. Since this version of High Stakes causes just as much

²¹ Notice that, just as Rachel and Keith can consider the possibility that the bank has recently changed its Saturday hours without having in mind any specific reason why the bank might have recently changed its Saturday hours, Rachel and Keith can consider the possibility that the bank will not be open on Saturday without considering any specific reason why the bank might not be open on Saturday.

trouble for minimalism as the *original* version of High Stakes, there is little point in responding to the original version of High Stakes with a WAM unless we can *also* respond to this version of High Stakes with a WAM.

So, is there any value of ' p ' where the following claim is plausible?

By uttering $\neg K$ in Practical High Stakes, Keith implicates that p . Keith seems to say something true by uttering $\neg K$ in Practical High Stakes because p is true and we mistake p for what Keith *says* by uttering $\neg K$ in Practical High Stakes.

We get our answer by paying close attention to the decision-theoretic nature of the *original* version of High Stakes—a feature that is preserved in Practical High Stakes, and which is present in *all* of the high-stakes cases in the literature (Cohen's airport case, Fantl and McGrath's train case, Weatherson's genie case, Ross and Schroeder's sandwich case, and so on).²²

3 A WAM with General Application

In the *original* version of High Stakes, Rachel and Keith are deciding whether they should wait in line to deposit their checks Friday evening, or go straight home and deposit them Saturday morning. Deciding whether they should wait in line or go straight home is the sole purpose of their conversation. They would both prefer to go straight home conditional on the bank's being open on Saturday, and they both know this. But they also both know that it would be too risky for them to go straight home without very strong evidence that the bank

²² See Cohen (1988), p. 58, Fantl and McGrath (2002), p. 67-8, Weatherson (forthcoming), pp. 10-11 of version available on Weatherson's website, and Ross and Schroeder (forthcoming), p. 3 of version available on Schroeder's website.

will be open on Saturday. They know that the bank has been open on many previous Saturdays, and, in this context, Rachel raises the possibility that the bank has just discontinued its Saturday hours. It is clear enough *why* Rachel raises this possibility: she thinks it would be too risky for them to go straight home unless they can rule it out. Rachel knows that *her* evidence that the bank will be open does not rule this possibility out, so she needs to know whether *Keith's* evidence rules it out. But instead of just asking Keith whether he has sufficient evidence to rule out the possibility that their bank has changed its Saturday hours since his last visit, she asks Keith *whether he knows that the bank will be open on Saturday*. Why does Rachel ask Keith *this* question? Why doesn't she just ask Keith whether he can rule out the possibility that their bank has recently changed its Saturday hours?

Plausibly, Rachel doesn't ask Keith whether he can rule out the possibility that their bank has recently changed its Saturday hours because she recognizes that an answer to this question would only give her *part* of the information that she needs. There are *many* possible worlds where the bank does not open this Saturday, and the world where the bank has recently discontinued its Saturday hours is only one of them. Rachel needs to know whether Keith can rule out *all* of these worlds—or, at least, whether he can rule out all such worlds *that are roughly as likely to be actual as the world where bank has recently changed its Saturday hours*. Plausibly, Rachel does not ask Keith whether he can rule out the possibility that their bank has recently changed its Saturday hours because she knows that, even if he *can* rule this possibility out, there might be other equally worrisome possibilities that he *cannot* rule out. So we can explain why Rachel *does not* ask Keith whether he can rule out the possibility that the bank has recently changed its Saturday hours.

But still, why does Rachel ask Keith what she *does* ask him: *whether Keith knows that the bank will be open?* Why does she ask *this* question? After all, if Rachel doesn't ask Keith whether he can rule out the possibility that their bank has recently changed its Saturday hours because she is thinking about other, equally worrisome possibilities, then why doesn't she use the possibility that bank has changed its Saturday hours since Keith's last visit *as an example* and then ask Keith whether he can rule out this possibility *and other equally worrisome possibilities?*

In context, the answer to this question is clear enough, and it has everything to do with the sole purpose of Rachel and Keith's conversation: to decide whether they should wait in line to deposit the checks Friday evening or, instead, go straight home and return to deposit them on Saturday morning. Compare High Stakes to the paradigm case of conversational implicature that we discussed above: the pedestrian's implicating that there is an open garage nearby by uttering just 'there is a garage nearby.' As we noted, Brown discusses a case slightly different from Grice's. In Brown's case, the motorist tells the pedestrian that she is running out of gas, and the pedestrian immediately responds by uttering 'there is no garage nearby' instead of 'there is a garage nearby.' So imagine a case that differs from *Brown's* case, as follows.

Motorist: I'm running out of gas. Is there a garage nearby?

Pedestrian: No, there isn't a garage nearby. The nearest one's in Sleetmute.

In this case, the motorist says that she is running out of gas and then asks, not whether there is an open garage nearby, but just whether there is a garage nearby. There is an obvious

symmetry between the motorist's question in this case and Rachel's question in High Stakes. The motorist's sole purpose in addressing the pedestrian is to find out where she can get gas, and it is transparent what the motorist means to accomplish by asking the pedestrian whether there is a garage nearby: she means to find out whether there is a garage nearby *where she can get gas*. Likewise, Rachel's sole purpose in questioning Keith is to determine the best course of action, and it is transparent what Rachel means to accomplish by asking Keith if he knows that the bank will be open: she means to find out whether Keith has knowledge *that they can reasonably act on* ('actionable knowledge,' as I will sometimes say).²³ But given this symmetry, the claim that Keith implicates that he lacks *actionable* knowledge by uttering $\neg K$ is just as plausible as Brown's claim that the pedestrian implicates that there is no *open* garage nearby by uttering 'there is no garage nearby.' If Brown is right that the pedestrian implicates that there is no open garage nearby by uttering 'there is no garage nearby,' then it would be surprising if Keith did not implicate that he lacks actionable knowledge by uttering $\neg K$. So Brown should think that, in the original version of High Stakes, Keith implicates that he lacks *actionable* knowledge by uttering $\neg K$.

Now we said that there was little point in responding to the original version of High Stakes with a WAM if there was no way to apply that WAM to Practical High Stakes. But, of course, the suggestion that Keith implicates that he lacks actionable knowledge by uttering $\neg K$ is only *more* plausible with respect to Practical High Stakes. So, noting the implausibility of the suggestion that Keith implicates that his belief does not match the facts out to the nearest world in which the bank has changed its Saturday hours since his last visit by uttering

²³ I owe this term to correspondence with Ernest Sosa. Note that, according to pragmatism, *all* knowledge is actionable knowledge.

$\neg K$ in Practical High Stakes, and noting the comparative plausibility of the suggestion that Keith implicates that he lacks actionable knowledge by uttering $\neg K$ in both the original version of High Stakes *and* Practical High Stakes, Brown should think that the following WAM will suffice in response to High Stakes if *any* WAM will suffice.

By uttering $\neg K$ in High Stakes, Keith implicates that he lacks actionable knowledge; Keith seems to say something true by uttering $\neg K$ in High Stakes because we mistake this true thing that he implicates by uttering $\neg K$ in High Stakes for what he *says* by uttering $\neg K$ in High Stakes.

4 Conversational Implicatures and Cancelability

As we noted above, when minimalists say that Keith implicates something true in High Stakes, and when they say that we mistake what Keith implicates in High Stakes for what he *says* in High Stakes, they are talking about *conversational* implicature. But conversational implicatures are *cancelable*, and this property of conversational implicatures will cause trouble for minimalists who respond to the limits of Brown's WAM with the actionable knowledge WAM above.²⁴

Again, consider the conversation Grice gives us in his well-known petrol example.²⁵

²⁴ Ibid., p. 39 As Grice puts it, "A putative conversational implicature that *p* is explicitly cancelable if, to the form of words the utterance of which putatively implicates that *p*, it is admissible to add 'but not *p*,' or 'I do not mean to imply that *p*,' and it is contextually cancelable if one can find situations in which the utterance of the form of words would simply not carry the implicature. Now I think that all conversational implicatures are cancelable." See Grice (1980), p. 44. Throughout, I will use 'implicate' as short of 'conversationally implicate,' and 'implicature' as short for 'conversational implicature.'

²⁵ Ibid., p. 32

Motorist: I'm running out of gas. Is there a garage nearby?

Pedestrian: Yeah, there is a garage nearby—right around the corner.

Here the pedestrian utters 'there is a garage nearby' and thereby implicates that there is an *open* garage nearby. Had the garage been *closed*, however, the pedestrian could have *cancelled* this implicature by adding 'but it is closed.' That is, had the pedestrian uttered 'yeah, there is a garage nearby—but it is closed,' she would have successfully conveyed to the motorist that there is a garage nearby *without* conveying to the motorist that there is an *open* garage nearby.

Now, as we saw above, Brown takes it for granted that, because the pedestrian implicates that there is an open garage nearby by uttering just 'there is a garage nearby,' the pedestrian would have implicated that there is *no* open garage nearby had she uttered just 'there is no garage nearby.' Grice would reject Brown's claim that the pedestrian would have implicated that there is no open garage nearby had she uttered just 'there is no garage nearby,' since the proposition that there is no open garage nearby fails the cancelability test. Consider the following answers that the pedestrian might have given to the motorist.

A1: No, there's no garage nearby, but there is an open garage nearby.

A2: No, there's no garage nearby, but I am not saying that there is no open garage nearby.

When the pedestrian utters 'there is no garage nearby,' she thereby communicates to the motorist that there is no *open* garage nearby. But by Grice's lights, the pedestrian does not *implicate* that there is no open garage nearby, since, downstream from an utterance of 'there is

no garage nearby,’ the proposition that there is no *open* garage nearby is not *cancellable*. As A1 and A2 show us, once the pedestrian has uttered ‘there is no garage nearby,’ it is not (in Grice’s words) admissible for the pedestrian to add that there *is* an open garage nearby, or that she is *not* saying that there is no open garage nearby, or anything like that.²⁶ By attempting to cancel the proposition that there is no open garage nearby, the pedestrian makes nonsense of her answer to the motorist.²⁷ A similar worry will threaten the actionable knowledge WAM above.

5 The Limits of WAMs

In High Stakes, Rachel raises the possibility that their bank has recently changed its Saturday hours and then, without pause, asks Keith whether he knows that the bank will be open. Keith responds by uttering $\neg K$, and, according to the WAM that we are forwarding as the most promising WAM available, Keith thereby implicates that he lacks knowledge *that they can act on* that the bank will be open. But the proposition that Keith lacks knowledge that they can act on that the bank will be open is no more a cancellable implicatum of Keith’s utterance of $\neg K$ than the proposition that there is no open garage nearby is a cancellable implicatum of the pedestrian’s utterance of ‘there is no garage nearby,’ or the proposition that somebody has long-jumped over 28 feet is a cancellable implicatum of your utterance of ‘Bob Beamon long-jumped over 29 feet.’ Consider the following way that High Stakes might have ended.

²⁶ Grice (1989), p. 44. Though I do not have space to defend the claim here, I believe that Brown’s claim that the pedestrian implicates that there is no open garage nearby by uttering ‘there is no garage nearby’ also fails Grice’s *calculability* test. (See *ibid.*, p. 39.)

²⁷ Some of Kent Bach’s comments in “The Top 10 Misconceptions about Implicature” suggest that Bach does not agree with Grice that conversational implicatures are always cancellable. See Bach (2006), p. 24.

Rachel: What if our bank discontinued its Saturday hours since your last visit? Do you know that it will be open?

Keith: No, I don't know that the bank will be open, but I do have knowledge that the bank will be open that we can act on.

What could Keith be telling Rachel here? Just as the pedestrian could not utter 'there is no garage nearby' and then admissibly add that there is an open garage nearby, and just as you could not utter 'Bob Beamon long-jumped over 29 feet' and then admissibly add that nobody long-jumped over 28 feet, Keith cannot utter $\neg K$ and then admissibly add that he has actionable knowledge that the bank will be open. With respect to Keith's utterance of $\neg K$, the proposition that Keith lacks actionable knowledge that the bank will be open performs miserably on Grice's cancelability test.

Now Brown anticipates worries about Grice's cancelability test in response to her claim that, by uttering $\neg K$ in High Stakes, Keith implicates that his belief that the bank will be open does not match the facts out to the nearest world in which the bank has changed its hours since his last visit. Here is what Brown says.

On the proposed account, '*S* knows that *p*' may pragmatically convey that *S* is in a very strong epistemic position, that her belief matches the facts across a wide range of worlds, including some so far away they are not normally taken to undermine knowledge. Contrary to the objection, it may be possible to cancel this implication.

To my ear, the claim that '*S* knows that *p*, but *S* is not in a really strong epistemic

position with respect to p or ' S knows that p , but her belief wouldn't match the facts in a really distant possible world' do not seem obviously inconsistent or uncomfortable.²⁸

I have some sympathies for what Brown says here. But nothing she says here it seems apposite to *our* worries about cancelability, since Brown only addresses *positive* utterances of the form 'I know that p .' She says nothing about *negative* utterances of the form 'I *do not* know that p ,' which are what we are worrying about. Even if Keith would implicate that p (for some value of ' p ') by uttering 'I know that the bank will be open' in High Stakes, it does not follow that Keith would implicate that $\neg p$ (for that same value of ' p ') by uttering 'I *do not* know that the bank will be open' in High Stakes. So we can easily agree with everything that Brown says above and continue to maintain that Keith *does not* implicate that he lacks actionable knowledge by uttering $\neg K$ in High Stakes. This would be no more difficult than agreeing with everything that Brown says above while also agreeing with Grice that the pedestrian in his scenario would not implicate that there is no open garage nearby by uttering 'there is no garage nearby.'

So we have good reason to doubt that, by uttering $\neg K$, Keith implicates that he lacks actionable knowledge. This means that the actionable knowledge WAM that we forwarded as an improvement over Brown's WAM will not work. There is one more WAM to consider, however. Again, return to Grice's petrol case. In that case, the pedestrian utters 'there is a garage nearby' and thereby implicates that there is an open garage nearby. What explains this implicature? According to Grice's Cooperative Principle, the pedestrian must make her

²⁸ Brown (2006), p. 428.

conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the conversation in which she is engaged.²⁹ According to Grice's Maxim of Relation, the pedestrian must make her conversational contribution *relevant*.³⁰ The pedestrian does not seem to be flouting the Maxim of Relation when she tells the motorist that there is a garage nearby, and we have no reason to think that she is being uncooperative by telling the motorist just that there is a garage nearby. Yet the pedestrian must regard her own utterance as irrelevant if she does not think that the nearby garage is *open*. After all, the motorist obviously needs to buy gas, and the pedestrian (we may presume) knows that a closed garage is useless with respect to these needs. So preserving the assumption that the pedestrian is observing the Cooperative Principle requires assuming that the pedestrian believes that the nearby garage is open. *This* (according to Grice) is why the pedestrian implicates that there is an open garage nearby by uttering just 'there is a garage nearby.'³¹

If this is why the pedestrian implicates that there is an open garage nearby by uttering 'there is a garage nearby,' however, then it seems that the pedestrian also implicates *that the motorist can get gas somewhere nearby* by uttering 'there is a garage nearby.' Just as the pedestrian must regard her utterance as irrelevant if she does not think that the nearby garage is open, she must regard her utterance as irrelevant if she does not think that the motorist can get gas somewhere nearby. After all, the pedestrian knows that the motorist is running out of gas. Preserving the assumption that the pedestrian is observing the Cooperative Principle requires assuming *both* that the pedestrian believes that the nearby garage is open and that

²⁹ Grice (1989), p. 26.

³⁰ *Ibid.*, p. 27.

³¹ *Ibid.*, p. 32.

the pedestrian believes that the motorist can get gas somewhere nearby. So, by uttering ‘there is a garage nearby,’ the pedestrian implicates both that there is an open garage nearby and that the motorist can get gas somewhere nearby.

Now take another look at the conversation where the pedestrian gives a *negative* answer to the motorist’s question. Here it is, again.

Motorist: I’m running out of gas. Is there a garage nearby?

Pedestrian: No, there’s no garage nearby. The nearest one’s in Sleetmute.

By uttering ‘there is no garage nearby,’ the pedestrian does not implicate that there is no open garage nearby, as we saw above. Plausibly, however, she *does* implicate that the motorist *cannot* get gas anywhere nearby. First, just as preserving the assumption that the pedestrian is observing the Cooperative Principle in Grice’s original petrol case requires assuming that the pedestrian believes that the nearby garage is open, preserving the assumption that the pedestrian is observing the Cooperative Principle requires assuming that the pedestrian believes that the motorist cannot get gas anywhere nearby. Second, the suggestion that the pedestrian implicates that the motorist cannot get gas anywhere nearby easily passes the cancelability test. Consider the following conversation.

Motorist: I’m running out of gas. Is there a garage nearby?

Pedestrian: No, there’s no garage nearby, but there *is* a place nearby where you can get gas. They sell it out at the airfield.

It would be perfectly admissible for the pedestrian to utter ‘there is no garage nearby’ and then add that there *is* a place nearby where the motorist can get gas. By uttering ‘there is no garage nearby, but there is a place nearby where you can get gas,’ the pedestrian successfully conveys to the motorist that there is no garage nearby *without* conveying to the motorist she cannot get gas anywhere nearby. So the proposition that the motorist cannot get gas anywhere nearby is a cancellable implicatum of the pedestrian’s utterance of ‘there is no garage nearby.’

Now consider the suggestion that, by uttering $\neg K$, Keith implicates *that Rachel and Keith cannot reasonably go straight home*. As we noted above, Rachel’s sole purpose in questioning Keith is to determine the best course of action, and it is transparent what Rachel means to accomplish by asking Keith if he knows that the bank will be open. She means to find out whether they can reasonably go straight home. So minimalists can point out that, just as preserving the assumption that the pedestrian is observing the Cooperative Principle in Grice’s original petrol case requires assuming that the pedestrian believes that the nearby garage is open, preserving the assumption that Keith is observing the Cooperative Principle in High Stakes requires assuming that Keith believes that he and Rachel cannot reasonably go straight home. This similarity motivates the claim that, by uttering $\neg K$ in High Stakes, Keith implicates that Rachel and Keith cannot reasonably go straight home.

Now, here is the question: will a WAM built on this supposed implicature pass the cancelability test? Consider the following alternative ending to High Stakes.

... If their paychecks are not deposited before Monday morning, the important check they wrote will bounce, leaving them in a terrible situation. The bank isn't open on Sunday. Rachel reminds Keith of these facts. She then says, "What if our bank discontinued its Saturday hours since your last visit? Do you know that it will be open?" Remaining just as confident as he was before that the bank will be open, Keith replies, "Well, no, I don't know that the bank will be open, but we can still reasonably go straight home."

Is it admissible for Keith to utter 'I don't know that the bank will be open' to Rachel and then add that they can still reasonably go straight home? By uttering 'I don't know that the bank will be open, but we can still reasonably go straight home,' does Keith successfully convey to Rachel that he does not know that the bank will be open *without* conveying to her that they *cannot* reasonably go straight home? I'm not sure. That Keith does not know that the bank will be open does not entail that Rachel and Keith cannot rationally go straight home, as it *does* entail that Keith lacks actionable knowledge that the bank will be open. So perhaps Keith *can* utter $\neg K$ and then cancel the proposition that Rachel and Keith cannot rationally go straight home. Then again, given the details of High Stakes, it is hard to imagine what reason Keith could have for thinking *both* that he does not know that the bank will be open *and* that Rachel and Keith can reasonably go straight home. I have a hard time seeing what Keith could be telling Rachel here. So I am inclined to say that, by attempting to cancel the proposition that Rachel and Keith cannot reasonably go straight home, Keith would make nonsense of his answer to Rachel's question. As a consequence, I am inclined to deny

that, by uttering $\neg K$, Keith implicates that he and Rachel cannot reasonably go straight home.

6 The Limits of Gricean Responses in General

Worries along these lines will plague any WAM that minimalists might offer in response to High Stakes. But don't minimalists have resources here for an adequate Gricean response to High Stakes, even if this response isn't exactly a WAM? For example, can't minimalists just shift their attention to *Rachel*, and argue that *she* is responsible for the implicatures that mislead us about the truth-value of Keith's utterance of $\neg K$? After all, the Cooperative Principle and Grice's Maxim of Relation apply to anything you might contribute to a conversation, including a *question*.³² The problem is, even if Rachel *is* responsible for the implicatures that mislead us about the truth-value of Keith's utterance of $\neg K$, this fact will ultimately be of little use to minimalists. In §2, we saw that there is little point in responding to High Stakes with a given WAM if that WAM cannot handle a slight revision of High Stakes. This lesson generalizes: there is little point in responding to High Stakes *in any particular way* if that way cannot handle a slight revision of High Stakes. So now consider the following revision of High Stakes.

Silent High Stakes: Keith is driving to the bank by himself on Friday afternoon when he notices the long lines. He knows that the bank has been open every previous Saturday, but he also knows that there will be disastrous consequences if he fails to deposit his paycheck before Sunday. Given the amount and quality of Keith's

³² See *ibid.*, pp. 26-7.

evidence that the bank will open, balanced against the severity of the consequences that Keith knows would follow if he failed to deposit his paycheck before Sunday, Keith cannot rationally go straight home—which is what he would prefer to do conditional on the proposition that the bank will be open.

Keith does not utter $\neg K$ in Silent High Stakes, so we cannot ask if he seems to say something true by uttering $\neg K$ in Silent High Stakes. But we *can* ask whether Keith is in position to know that the bank will be open in Silent High Stakes, and intuitively (at least *prima facie*) the answer is ‘no.’³³

Minimalists interested in WAMs or other Gricean responses to High Stakes now have a problem. Keith lacks knowledge in Silent High Stakes only if he also lacks knowledge in the original version of High Stakes, since his epistemic position in Silent High Stakes is just as strong as his epistemic position in the original version of High Stakes. So *Silent High Stakes* motivates the conclusion that Keith lacks knowledge in the original version of High Stakes. But there is no Gricean story to tell about Silent High Stakes. There is no way to apply the concepts, distinctions and principles familiar from Grice to this case. Keith does not utter K or $\neg K$ in Silent High Stakes, or even *mentally token* K or $\neg K$ in this version of the case, so Gricean considerations gain no purchase here. And the same goes for the sentences ‘Keith knows that the bank will be open’ and ‘Keith does not know that the bank will be open.’ These sentences make no appearance in Silent High Stakes. Of course, *we* might mentally token one or both of these sentences as we consider Silent High Stakes. But are we

³³ And even if we *lack* the intuition that Keith is not in position to know that the bank will be open in Silent High Stakes, principles like those defended in Fantl and McGrath (2002), Hawthorne and Stanley (2008), Ganson (2008), Fantl and McGrath (2009), Ross and Schroeder (2012), and Weatherson (forthcoming) entail or strongly suggest that Keith is not in position to know that the bank will be open in Silent High Stakes.

supposed to think that, as we sit here considering Silent High Stakes, we mentally token the sentence ‘Keith does not know that the bank will be open’ and thereby implicate *to ourselves* some true proposition that we then mistake for the false proposition that we would semantically express by uttering ‘Keith does not know that the bank will be open,’ and that *this* is why Keith seems to lack knowledge in Silent High Stakes? This suggestion seems wildly implausible.³⁴ So it seems doubtful that minimalists can give an adequate Gricean response to Silent High Stakes.

Since Silent High Stakes motivates the conclusion that Keith lacks knowledge in the original version of High Stakes, it seems doubtful that minimalists can give an adequate Gricean response to the original High Stakes case, whether this response takes the form of a WAM or not.

7 Conclusion

At root, the problem for minimalism is not that Keith seems to say something true by uttering $\neg K$ in High Stakes. The problem is that Keith seems to lack knowledge in High Stakes. Gricean considerations bear on our intuitions about Keith’s knowledge in High Stakes only *via* our intuition that Keith says something true by uttering $\neg K$ in High Stakes. But our intuition that Keith does not know that the bank is open in High Stakes does not depend on the intuition that Keith says something true by uttering $\neg K$ in High Stakes. As

³⁴ I do not mean to commit myself to the claim that, *in general*, it is impossible to implicate things to ourselves by mentally tokening sentences. As Jason Stanley has pointed out in conversation, if I know that it’s raining (for example) and it’s obvious to me that I know that it’s raining, it would be just as strange for me to *think to myself* ‘it might be raining’ as it would be to utter the sentence ‘it might be raining.’ I don’t want to rule out that possibility that the correct explanation here is at least *in part* that, by thinking to myself ‘it might be raining,’ I would implicate to myself that I don’t know whether it is raining.

Silent High Stakes shows us, things go the other way around. Our intuition that Keith says something true in High Stakes rests on our intuition that Keith lacks knowledge in High Stakes. So Gricean considerations are ultimately ineffectual for defending minimalism from the challenges raised by High Stakes. The upshot is that, instead of worrying about what is said and/or implicated by Keith in High Stakes, minimalists must direct their attention at the properly *epistemological* intuition that Keith lacks knowledge in High Stakes.

Of course, there is nothing special about High Stakes. The same lesson applies to *all* of the high-stakes cases in the literature.³⁵ To adequately respond to *any* high-stakes case, minimalists must prioritize the properly epistemological question whether the subject in that case knows over the linguistic question whether the relevant knowledge-ascribing sentences would semantically express true propositions. Questions about the semantics and pragmatics of knowledge-ascribing sentences are interesting in their own right, and they are also relevant to the debate between contextualists and invariantists over DeRoses's "skeptical paradox," so I don't want to suggest that one wastes one's time by asking how Gricean considerations bear on utterances like Keith's utterance of $\neg K$ in High Stakes. But I do want to insist that, at best, Gricean responses to cases like High Stakes will play a marginal role in any adequate defense of minimalism.

³⁵ Ross and Schroeder do not give minimalists any choice here. While their low-stakes sandwich case involves an utterance of a knowledge-ascribing sentence, their high-stakes sandwich case is a lot like Silent High Stakes in that it contains neither an utterance nor mental tokening of any knowledge-ascribing or knowledge-denying sentence. (See Ross and Schroeder (forthcoming), p. 3.) The same goes for Weatherson. His genie case does not involve any utterance of any knowledge-denying or knowledge-ascribing sentence. (See Weatherson (forthcoming), pp. 10-11.)

Chapter 2: Pragmatism and the Nature of Belief

Pragmatists and purists disagree about the relationship between knowledge and practical interests. Pragmatists think a difference in knowledge might follow from as little as a difference in practical interests, while purists deny this claim. According to purism, if you know that p and I do not, this difference traces to something beyond whatever differences there might be in our practical interests. Recent versions of pragmatism rely heavily on theories of *belief*, and this new emphasis has changed the shape of the pragmatic encroachment debate. Purists must now take a stand on the nature of belief if they want to resist pragmatism. A version of pragmatism is *doxastic* (I will say) just in case it forces purists to take a stand on the nature of belief, and *non-doxastic* just in case it allows purists to stay neutral on the nature of belief. If any version of doxastic pragmatism is correct, then pragmatism follows from the correct theory of belief with just the addition of assumptions that purists will happily accept. If some version of *non-doxastic* pragmatism is correct, however, then pragmatism does not follow from just the correct theory of belief and assumptions that purists will happily accept. Instead, pragmatism follows at least in part from controversial premises about truth, justification, or knowledge itself.³⁶ Dorit Ganson and Brian Weatherson have recently forwarded versions of doxastic pragmatism, and Jacob Ross, Mark Schroeder, Jeremy Fantl and Matthew McGrath have all articulated theories of belief that raise the possibility of defending doxastic pragmatism, even if these authors do

³⁶ By 'justification' in this context, I just mean whatever makes the difference between knowledge and true belief that falls short of knowledge.

not defend doxastic pragmatism themselves.³⁷ In this paper, I consider just the debate between purism and doxastic pragmatism.³⁸ I leave non-doxastic pragmatism to the side and argue that the theories of belief defended by Ganson and Weatherson are susceptible to counterexample, while the theories of belief defended by Ross, Schroeder, Fantl and McGrath are either susceptible to counterexample or too loosely connected to practical rationality to be of any use for defending doxastic pragmatism. The positive case for doxastic pragmatism is wanting, I argue. Moreover, in addition to finding a counterexample-free theory of belief that entails pragmatism with just the help of assumptions that purists will happily accept, doxastic pragmatists must convincingly argue against plausible, independently motivated theories of belief that are inconsistent with doxastic pragmatism. I give an example of such a theory and conclude that, everything considered, doxastic pragmatism is insufficiently motivated.

1. Ganson and Weatherson on Belief

According to Dorit Ganson and Brian Weatherson, believing that p is intimately related to conditionalizing on p . On Ganson's view,

(GB) S believes that p only if there is no difference between what S is actually willing to do and what S *would* be willing to do if she were to conditionalize on p .³⁹

³⁷ Ganson (2008), Weatherson (forthcoming), Fantl and McGrath (2009), and Ross and Schroeder (forthcoming).

³⁸ I consider the debate between purists and *non*-doxastic pragmatists in "A Paradox for Justification, Knowledge and Practical Rationality."

³⁹ Ganson (2008), p. 453.

And according to Weatherson,

(WB) S believes that p iff conditionalizing on p does not change S 's answer to any relevant question,⁴⁰

where the relevant questions take the following four forms: *How probable is it that q ? Is it more probable that q , or more probable that r ? How good of an idea is it to φ ? Is it better to φ , or better to ψ ?*⁴¹

Putting his view formally, Weatherson thinks BAP, BCP, BAU and BCU (below) give us necessary conditions on belief.

(BAP) For all relevant q , x , if p is believed, then $\Pr(q) = x$ iff $\Pr(q|p) = x$.

(BCP) For all relevant q , r , if p is believed, then $\Pr(q) \geq \Pr(r)$ iff $\Pr(q|p) \geq \Pr(r|p)$.

(BAU) For all relevant φ , x , if p is believed, then $U(\varphi) = x$ iff $U(\varphi|p) = x$.

(BCU) For all relevant φ , ψ , if p is believed, then $U(\varphi) \geq U(\psi)$ iff $U(\varphi|p) \geq U(\psi|p)$.⁴²

Ganson and Weatherson's theories of belief get plausible results in a wide range of cases, and they do entail pragmatism with just the help of assumptions that purists will happily accept. Pragmatism is true just in case a mere difference in practical interests can suffice for a difference in knowledge. So, for example, consider the following familiar cases, and assume that they differ only insofar as the stipulated difference in Hannah's practical interests requires that they differ.

⁴⁰ Weatherson (forthcoming), p. 6.

⁴¹ Ibid.

⁴² Ibid., p. 15. $U(\varphi)$ is the expected utility of φ -ing.

Low Stakes: Hannah is driving past the bank on Friday afternoon. She has her paycheck in hand and she plans to deposit it. But there are long lines, so she considers returning to deposit her paycheck on Saturday morning. She knows that it does not matter much when she deposits the paycheck.

High Stakes: Hannah is driving past the bank on Friday afternoon. She has her paycheck in hand and she plans to deposit it. But there are long lines, so she considers returning to deposit her paycheck on Saturday morning. She knows that she will incur an enormous fine if she does not deposit her check before Saturday afternoon.

Taking it for granted that Hannah knows that the bank will be open in Low Stakes, pragmatism is true if Hannah *does not* know that the bank will be open in High Stakes. This follows from our stipulation that Low Stakes differs from High Stakes only insofar as the difference in Hannah's practical interests requires that they differ. So consider the bearing of GB and WB on High Stakes. Both GB and WB make explicit reference to conditionalization, so consider the consequences that conditionalizing on the proposition that the bank will be open would have for Hannah's choice between stopping at the bank and going straight home.

We can represent Hannah's choice between stopping at the bank and going straight home with the following decision table.

Table 1

	The bank will be open	The bank will not be open
Stop at bank	<i>Wait in line</i>	<i>Avoid enormous fine</i>
Go straight home	<i>Avoid the lines</i>	<i>Incur enormous fine</i>

And we can suppose that the following value matrix accurately represents the values that Hannah places on waiting in line, avoiding the lines, avoiding the fine, and incurring the fine, respectively.⁴³

Value Matrix for Table 1

	The bank will be open	The bank will not be open
Stop at bank	-\$1	\$100,000
Go straight home	\$1	-\$100,000

Given this setup, if Hannah conditionalizes on the proposition that the bank will be open, her expected utility calculation will tell her that going straight home has higher expected utility than stopping at the bank. Specifically, if she conditionalizes on the proposition that the bank will be open, her expected utility calculation will tell her that the expected utility of stopping at the bank is -\$1, while the expected utility of going straight home is \$1. On the proper understanding of High Stakes, however, Hannah's expected utility calculation *would*

⁴³ Perhaps the fine is really this large, or perhaps Hannah knows that the consequences of the fine will ramify to a \$100,000 loss, or perhaps Hannah has a massive aversion to risk. Whatever the explanation, I have stipulated such extreme values to insure that there is no ambiguity about whether it would be rational for Hannah to go straight home.

not tell her that going straight home has higher expected utility than stopping at the bank. Rather, it would tell her that stopping at the bank has higher expected utility than going straight home, and by a large margin. High Stakes, after all, is meant to be a case in which Hannah cannot rationally do what she would prefer to do, conditional on the proposition that the bank will be open. So let's suppose that Hannah is proportioning her confidence that the bank will be open to her basic evidence that it will be open, and let's suppose that, *because* she is proportioning her confidence to her basic evidence, Hannah's confidence that the bank will be open is exactly 0.99.⁴⁴ Given these suppositions, we get the following result: the expected utility of stopping at the bank is \$999.01, while the expected utility of going straight home is -\$999.01.

So now, given all of this, what do Ganson and Weatherson's theories of belief tell us about Hannah in High Stakes? They both say that Hannah fails to know that the bank will be open for lack of *belief* that the bank will be open. According to Ganson's GB, Hannah believes that the bank will be open only if there is no difference between what Hannah is actually willing to do and what Hannah would be willing to do if she were to conditionalize on the proposition that the bank will be open. But there *is* a difference between what

⁴⁴ We should note two things here. First, I say "basic evidence" rather than just "evidence" to avoid complications resulting from Timothy Williamson's claim that one's evidence is all and only one's knowledge. I mean to stipulatively define 'basic evidence' so that one's basic evidence *that p* cannot include the proposition that *p*, regardless of whether one *knows* that *p*. Thus, the way I mean to stipulatively define 'basic evidence,' Hannah's basic evidence *that the bank will be open* does not include the proposition that the bank will be open, whether Hannah knows that the bank will be open or not. Second, some readers might object to the stipulation that the probability that the bank will be open conditional on Hannah's basic evidence is only 0.99. *Ex hypothesi*, Hannah has exactly the same basic evidence that the bank will be open in High Stakes as she has in *Low Stakes*, and, *ex hypothesi*, Hannah *knows* that the bank will be open in Low Stakes. But (the objection goes), knowledge requires better evidence than *this*. Readers who think that knowing that the bank will be open would require evidence that would justify some level of confidence higher than 0.99 are welcome to build the requisite level of evidence into both Low Stakes and High Stakes, and then (if necessary) adjust the consequences for failing to deposit the check in High Stakes accordingly.

Hannah is actually willing to do and what Hannah would be willing to do if she were to conditionalize on the proposition that the bank will be open: Hannah is not willing to go straight home, but she would be if she were to conditionalize on the proposition that the bank will be open. And according to Weatherson's WB, Hannah believes that the bank will be open only if conditionalizing on the proposition that the bank will be open does not change Hannah's answer to any relevant question. But conditionalizing on the proposition that the bank will be open *does* change Hannah's answer to some relevant questions: it changes her answer to the question, *which has higher expected utility: going straight home or stopping at the bank?*, for example. So Ganson and Weatherson's respective theories of belief both tell us that Hannah does not believe that the bank will be open. Since knowledge entails belief, both theories tell us that Hannah does not *know* that the bank will be open.⁴⁵

Given this consequence of Ganson and Weatherson's theories of belief, doxastic pragmatism is true if Ganson or Weatherson is right about belief. Both theories have a lot going for them. Weatherson's theory provides an elegant solution to the lottery paradox, for example, and fits very nicely with functionalism about belief.⁴⁶ If the broad functionalist picture is correct, we should not be surprised if something like Weatherson's view is true. Unfortunately, Weatherson and Ganson's theories are susceptible to the same

⁴⁵ Q: What happens if we just *stipulate* that Hannah believes that the bank will be open? A: Plausibly, Hannah fails to know that the bank will be open for lack of *justified* belief that it will be open. Take Weatherson's theory of belief. In order to satisfy the requirements of BAP, BCP, BAU and BCU, Hannah's confidence that the bank will be open must exceed 0.99999. This is just a consequence of the math. So, if Weatherson's theory of belief is correct and Hannah believes that the bank will be open, Hannah's confidence that the bank will be open exceeds 0.99999. But by hypothesis, Hannah's basic evidence that the bank will be open only justifies 0.99 confidence that it will be open. So Hannah is not justified in being 0.99999 confident that the bank will be open—much less *more than* 0.99999 confident that the bank will be open. Thus, plausibly, if Weatherson's theory of belief is correct, then Hannah is not justified in *believing* that the bank will be open. So, plausibly, if Weatherson's theory of belief is correct, then Hannah fails to know that the bank will be open even in a case where we stipulate that she *believes* that the bank will be open.

⁴⁶ See *ibid.*, p. 16.

counterexample. Consider the following three propositions: that two plus two equals four, that the world is more than five minutes old, and that the sun will rise tomorrow. These propositions are not equally probable. The probability that two plus two equals four exceeds the probability that the world is more than five minutes old, which exceeds the probability that the sun will rise tomorrow. Nevertheless, I *believe* all three of these propositions, and hence I believe that the sun will rise tomorrow. My belief that the sun will rise tomorrow is a counterexample to Ganson and Weatherson's theories of belief. (So is my belief that the world is more than five minutes old, but I will focus on my belief that the sun will rise tomorrow.) Here's why.

Take Ganson's theory of belief first. According to Ganson, I believe that the sun will rise tomorrow only if there is no difference between what I am actually willing to do and what I *would* be willing to do if I were to conditionalize on the proposition that the sun will rise tomorrow.⁴⁷ So let s be the proposition that the sun will rise tomorrow, and let w be the proposition that the world is more than five minutes old. I am actually willing to say that $\Pr(w)$ is greater than $\Pr(s)$, but I would *not* be willing to say this if I were to conditionalize on the proposition that the sun will rise tomorrow. After all, I know that $\Pr(s|s) = 1$. So there is a difference between what I am actually willing to do and what I would be willing to do if I were to conditionalize on the proposition that the sun will rise tomorrow. Thus, Ganson's theory of belief says (falsely) that I do not believe that the sun will rise tomorrow.

⁴⁷ Ganson *suggests* sufficient conditions for belief but stops short of actually endorsing any. (And note, since Ganson does not endorse any sufficient conditions for belief, some of Ross and Schroeder's objections to pragmatic credal reductivism in Ross and Schroeder (forthcoming) do not count against Ganson's account of belief.)

We get the same result from Weatherson's theory of belief. If not for their restrictions to *relevant* propositions and actions, BAP, BCP, BAU and BCU would tell us that S believes that p only if she is maximally confident that p —that she believes that p only if her confidence that p is literally 1. Even *with* their restrictions to relevant propositions and actions, however, these conditionals tell us that S believes that p only if *either* she is maximally confident that p *or* the difference between 1 and her confidence that p is irrelevant to whatever theoretical or practical question she is considering. So let Aw be the act of asserting that the world is more than five minutes old, and let As be the act of asserting that the sun will rise tomorrow. The probabilities of s and w , and the expected utilities of As and Aw , will be relevant to the plausibility of Weatherson's theory of belief. So how probable is s ? Which is more probable, s or w ? What is the expected utility of As ? And which has higher expected utility, As or Aw ? Because these questions are relevant to the plausibility of Weatherson's theory of belief—which is what we are presently considering—Weatherson's theory of belief says that I believe that the sun will rise tomorrow only if the following biconditionals all hold. (The following biconditionals are substitution instances of the consequences of BAP, BCP, BAU and BCU, respectively.)

$$(1) \quad \Pr(s) = x \equiv \Pr(s|s) = x$$

$$(2) \quad \Pr(w) \geq \Pr(s) \equiv \Pr(w|s) \geq \Pr(s|s)$$

$$(3) \quad U(As) = x \equiv U(As|s) = x$$

$$(4) \quad U(Aw) \geq U(As) \equiv U(Aw|s) \geq U(As|s)$$

Yet (1) through (4) are all false. First, (1) comes out true only when I assign the same value to $\Pr(s)$ as I assign to $\Pr(s|s)$. But right now, I am not assigning the same value to $\Pr(s)$ as I am assigning to $\Pr(s|s)$. I think $\Pr(s|s)$ is 1 while $\Pr(s)$ is less than 1. So (1) is false. Second, since I believe that $\Pr(w)$ is greater than $\Pr(s)$, I believe that $\Pr(w)$ is greater than *or equal to* $\Pr(s)$. But I do not believe that $\Pr(w|s)$ is greater than or equal to $\Pr(s|s)$. I know that $\Pr(s|s)$ is 1, after all. And since conditionalizing on s does not raise the probability of w , it is absurd that $\Pr(w|s)$ is as great as 1. So while I believe that $\Pr(w)$ is greater than or equal to $\Pr(s)$, I believe that $\Pr(w|s)$ is *less than* $\Pr(s|s)$ —in which case (2) is false. Both (1) and (2) are therefore false.

Now, in order to evaluate (3) and (4), we need to know what consequences might follow from my asserting that the sun will rise tomorrow and my asserting that the world is more than five minutes old, and what values I place on these consequences. So note that, among the possible consequences of my asserting these propositions, there are only two that I presently care about: asserting a truth and asserting a falsehood.⁴⁸ But given this, (3) and (4) come out false just like (1) and (2). $U(A_s)$ equals the probability that the sun will rise tomorrow multiplied by the value that I place on asserting a truth, plus the probability that the sun will *not* rise tomorrow multiplied by the value that I place on asserting a *falsehood*. Yet $U(A_s|s)$ just equals the value that I place on asserting a truth. Since I place a positive value on asserting a truth and a negative value on asserting a falsehood, and since the probability that I will assert a falsehood by asserting that the sun will rise tomorrow is higher than zero

⁴⁸ Here in my office by myself, what other consequences for asserting w or asserting s *are* there? I cannot assert w or s without making some sounds or marks on a page, of course, but I am utterly indifferent to these consequences. I simply do not care whether or not I produce them. So far as I can tell, as I sit here in my office by myself, asserting a truth and asserting a falsehood really are the only consequences of my asserting w or asserting s that I care about.

(it is higher than the probability that two plus two *does not* equal four, which is not *lower* than zero), $U(A_s)$ is lower than the value that I place on asserting a truth. Hence, $U(A_s)$ is lower than $U(A_s|s)$, and (3) is false. Finally, while $U(A_w)$ is higher than $U(A_s)$ —since the probability that the world is more than five minutes old is higher than the probability that the sun will rise tomorrow— $U(A_w)$ is lower than the value that I place on asserting a truth, and for the same reason that $U(A_s)$ is lower than the value that I place on asserting a truth: the probability that the world is *not* more than five minutes old is (just barely) higher than zero. And since the proposition that the sun will rise tomorrow makes no difference to the expected utility of A_w , $U(A_w|s)$ also comes out lower than the value that I place on asserting a truth. Thus, since $U(A_s|s)$ equals the value that I place on asserting a truth (whatever exactly it is), $U(A_w|s)$ is lower than $U(A_s|s)$, and (4) is false. So, according to WB, I fail at least four necessary conditions for believing that the sun will rise tomorrow.⁴⁹

But I assure you, I *do* believe that the sun will rise tomorrow. Even right now, while I am comparing the probability that the world is more than five minutes old to the probability that the sun will rise tomorrow, while I am considering the relative expected utilities of asserting that the sun will rise tomorrow and asserting that the world is more than five minutes old, and so on, I *do* believe that the sun will rise tomorrow. I believed that the sun will rise tomorrow before I sat down to write about Ganson and Weatherson's theories of belief, and I did not stop believing that the sun will rise tomorrow when I set out to answer such questions as, *which is more probable, that the world is more than five minutes old, or that the sun*

⁴⁹ In case this isn't all clear, see the Appendix II.

will rise tomorrow? So, Ganson and Weatherson's theories of belief both generate the false conclusion that I do not believe that the sun will rise tomorrow.⁵⁰

2. Ross and Schroeder on Belief

Conditionalizing on a proposition requires assigning that proposition probability 1, and Ganson and Weatherson's theories of belief both say that believing that p is intimately related to conditionalizing on p . But if I believe that the sun will rise tomorrow (as it seems transparent to me that I do), then the connection between believing a proposition and conditionalizing on that proposition cannot be as tight as Ganson and Weatherson think. Jacob Ross and Mark Schroeder also posit a connection between believing a proposition and conditionalization on it, but the connection that they posit is looser than the connections that Ganson and Weatherson posit. According to Ross and Schroeder,

⁵⁰ Objection: This counterexample is so obvious that it is hard to believe that Weatherson has not thought about it. But if he has, then he surely has something persuasive to say in reply. So it is hard to believe that this counterexample is really so decisive.

Reply: It is pretty clear from the article exactly what Weatherson would say. Consider the following passage from Weatherson (forthcoming), p. 15.

Charlie is trying to figure out exactly what the probability of p is. That is, for any $x \in [0, 1]$, whether $\Pr(p) = x$ is a relevant question. Now Charlie is well aware that $\Pr(p|p) = 1$. So unless $\Pr(p) = 1$, Charlie will give a different answer to the questions *How probable is p ?* and *Given p , how probable is p ?* So unless Charlie holds that $\Pr(p)$ is 1, she won't count as believing that p . One consequence of this is that Charlie can't reason, "The probability of p is exactly 0.978, so p ." That's all to the good, since that looks like bad reasoning.

Like Charlie, I am presently giving different answers to the question *How probable is it that the sun will rise tomorrow?* and the question *Given that the sun will rise tomorrow, how probable is it that the sun will rise tomorrow?* So Weatherson should say exactly the same thing about me as he says about Charlie. This means that, instead of denying that his view entails that I do not believe that the sun will rise tomorrow, Weatherson would argue that this consequence is actually a *virtue* of his view, since it prevents me from reasoning "the probability that the sun will rise tomorrow is n [for some value of ' n ' that is extremely close to 1], so the sun will rise tomorrow." For my own part, I fail to see that there would be any problem reasoning this way, and I suspect that Weatherson and I would be at loggerheads on this point.

(RSB) S believes that p only if she is defeasibly disposed to treat p as true in her reasoning,⁵¹

where S treats p as true in her reasoning just in case she evaluates her options by the same procedure as she would use to evaluate them conditional on p .⁵² Ross and Schroeder do not say exactly what it means to evaluate one's options by the same procedure as one would use to evaluate them conditional on p , but they do give clear examples. Suppose that you are deciding whether to φ or ψ , and compare the following ways that you might reason.

1. The consequences of φ -ing will be A . The consequences of ψ -ing will be B .
So, ...
2. If p , then consequences of φ -ing will be A and the consequences of ψ -ing will be B . But if $\neg p$, the consequences of φ -ing will be C and the consequences of ψ -ing will be D . The probability that p is n , and the probability that $\neg p$ is $(1 - n)$. So, ...

Ross and Schroeder make it clear that, if the procedure that you are using to evaluate your options is better represented by the second line of reasoning above (which explicitly mentions both p and its negation) than by the first line of reasoning above (which mentions neither p nor its negation), then you are not evaluating your options by the same procedure as you would use to evaluate them conditional on p .⁵³

⁵¹ Ross and Schroeder (forthcoming), p. 9.

⁵² Ibid., p. 6.

⁵³ Ibid., p. 7.

So now consider the bearing of RSB on Hannah in High Stakes. As we have filled in the details of High Stakes, the procedure that Hannah is using to evaluate her options is better represented by the second line of reasoning below than by the first line of reasoning below.

- 1a. If I go straight home, then I will avoid the long lines. If I stop, then I will have to wait in those lines. So, ...
- 2a. If I go straight home and the bank will be open tomorrow, then I will avoid the long lines. But if I go straight home and the bank will not be open tomorrow, disaster will ensue. If I stop and the bank will be open tomorrow, then I will have to wait in line, and for no reason. But if I stop and the bank will *not* be open tomorrow, simply waiting in line will allow me to avoid disaster. The probability that the bank will be open is 0.99, and the probability that the bank will *not* be open is 0.01. So, ...⁵⁴

Since the procedure that Hannah is using to evaluate her options is better represented by the second line of reasoning above than by the first line of reasoning above, Hannah is not evaluating her options by the same procedure as she would use to evaluate them conditional on the proposition that the bank will be open. Thus, by Ross and Schroeder's stipulative definition of 'treating p as true in one's reasoning,' Hannah is not treating the proposition that the bank will be open as true in her reasoning.

Does RSB now entail that Hannah does not believe that the bank will be open? It is tempting to answer 'yes,' since Hannah is not evaluating her options by the same procedure

⁵⁴ It is unrealistic that Hannah assigns such precise probabilities to the proposition that the bank will be open and its negation, of course, but the assumption that she does is harmless in this context.

as she would use to evaluate them conditional on the proposition that the bank will be open. Ross and Schroeder would answer ‘no,’ however. Hannah believes that the bank will be open in Low Stakes, and, because the disposition at issue in RSB is *defeasible*, Ross and Schroeder can allow that Hannah keeps her belief in High Stakes. An advantage of RSB over Ganson and Weatherson’s theories of belief, they say, is that RSB allows beliefs to be rationally stable *across* low- and high-stakes situations. As they put it,

it will be rational for an agent to change her beliefs as she moves between high- and low-stakes contexts only if such changing beliefs are licensed by procedures that strike an optimal balance between minimizing expected cognitive costs and maximizing the expected value of the agent’s deliberative conclusions. (Let’s call such procedures *optimific*.) But if [RSB] is true, then we should not expect the optimific procedures to license such changes in one’s beliefs, as such changes would require an unnecessary expenditure of cognitive resources. For recall that according to [RSB], the belief that *p* involves a *defeasible* disposition to treat *p* as true in reasoning, a disposition that is overridden when the costs of mistakenly acting as if *p* are salient. Hence, on this account, there will be no need to drop the belief that *p* when one enters a high-stakes context. For one can instead retain this belief while overriding the disposition to treat *p* as true, by attending to the costs of mistakenly acting as if *p*. And, having retained the belief that *p*, there will be no need to reacquire this belief when one reenters a low-stakes context. We should expect, therefore, that

the optimific rules would prescribe stability in one's beliefs as one moves between high- and low- stakes contexts.⁵⁵

So, on Ross and Schroeder's view, there is no reason why Hannah cannot believe that the bank will be open in High Stakes, just as rationality requires of her.

But if this is what Ross and Schroeder think, then how can they resist the conclusion that Hannah knows in High Stakes? And if they *cannot* resist the conclusion that Hannah knows in High Stakes, then how can they claim that pragmatism is true? Granted, cases like Low and High Stakes are not the only way to motivate pragmatism.⁵⁶ But Low and High Stakes give us a pair of cases where a difference in practical interests suffices for a difference in knowledge *if any cases do*, and pragmatism is false if there is *no* pair of cases where a difference in practical interests suffices for a difference in knowledge. After all, pragmatism is the view that a difference in knowledge might follow from nothing more than a difference in practical interests. So, whatever *motivates* pragmatism, pragmatism is almost certainly false unless Hannah lacks knowledge in High Stakes. The upshot is that Ross and Schroeder (and other pragmatists) must either abandon pragmatism or deny that Hannah knows in High Stakes.

As it happens, Ross and Schroeder have ample resources for denying that Hannah knows in High Stakes, but these resources do not sit well with doxastic pragmatism.⁵⁷ Ross and Schroeder see a tight connection between knowledge and rational action, and they argue for this connection as follows.⁵⁸ (I quote *verbatim*, except for the enumeration.)

⁵⁵ Ibid., pp. 21-2.

⁵⁶ Cf. Fantl and McGrath (2009).

⁵⁷ As we will see below, this needn't be a problem for Ross and Schroeder.

⁵⁸ Ross and Schroeder (forthcoming), pp. 13-14.

- (1) If, in C , it is relevant whether p , then believing that p essentially involves being disposed to treat p as true in one's reasoning in C . [from RSB]
- (2) If having attitude A essentially involves being disposed to φ under circumstance C , then an agent S is justified to occurrently have attitude A in C only if it is rationally permissible for S to φ in C . [assumption; what Ross and Schroeder call the "Justification Condition on Occurrent Attitudes"]
- \therefore (3) In C , if it is relevant whether p , and S is justified to occurrently believe that p , then it is rationally permissible for S to treat p as true in her reasoning. [from (1) and (2)]
- (4) In C , if S knows that p , then S is justified to occurrently believe that p . [assumption]
- (5) In C , if it is rationally permissible for S to treat p as true in her reasoning, then it is rationally permissible for S to act as if p . [assumption]
- \therefore (6) In C , if it is relevant whether p , and S knows that p , then it is rationally permissible for S to act as if p . [from (3) through (5)]
- (7) In C , if it is rationally *impermissible* for S to act as if p , then what is unconditionally optimal must differ from what is optimal conditional on p , and so it is relevant whether p . [assumption]
- \therefore (8) In C , if it is rationally impermissible for S to act as if p , then S does not know that p . [from (6) and (7)]

What about Hannah, then? According to (8), in High Stakes, if it is rationally impermissible for Hannah to act as if the bank will be open, then Hannah does not know that the bank will

be open. As Ross and Schroeder stipulatively define ‘acting as if p ,’ S acts as if p just in case she acts in the manner that would be rationally optimal conditional on p .⁵⁹ Since going straight home would be rationally optimal for Hannah conditional on the bank’s being open, Hannah acts as if the bank will be open only if she goes straight home. But since Hannah maximizes expected utility by stopping at the bank, and since this is perfectly clear to Hannah, Hannah cannot rationally go straight home. Thus, by Ross and Schroeder’s stipulative definition of ‘acting as if p ,’ it is rationally impermissible for Hannah to act as if the bank will be open. So according to (8), Hannah does not know that the bank will be open.

This way, Ross and Schroeder can get the same conclusion about Hannah as Ganson and Weatherson. But *unlike* Ganson and Weatherson, Ross and Schroeder cannot get this conclusion from their theory of belief *with just the addition of assumptions that purists would happily accept*. This is important relative to the purposes of this chapter, since we are presently evaluating doxastic pragmatism and, by stipulative definition, doxastic pragmatism is just the view that pragmatism follows from the correct theory of belief together with assumptions that purists will happily accept. Consider (4), for example. Ross and Schroeder endorse (4) on the basis that it seems plausible.⁶⁰ As Ross and Schroeder put it,

[i]t is generally agreed that knowledge requires justification, and, in particular, that it requires a level of justification that is at least high enough to justify believing that p .

We may distinguish, however, between *having the belief* that p and *occurrently believing* that p . Accordingly, we may distinguish between the level of justification required to

⁵⁹ Ibid., p. 4.

⁶⁰ Ibid., p. 13.

justify having the belief that p , and the level of justification required to justify occurrently believing that p . Plausibly, knowing that p requires a level of justification that is at least high enough to justify occurrently believing that p . Or, to put the same point in simpler (if less grammatical) terms, it is plausible that anyone who knows that p must be justified to occurrently believe that p .⁶¹

But purists who think that RSB is true may find (4) decidedly *implausible*. Indeed, purists who accept RSB might think that Hannah gives us a counterexample to (4). If RSB is true, then Hannah is occurrently believing that the bank will be open in High Stakes only if she is treating the proposition that the bank will be open as true in her reasoning in High Stakes: that is, only if she is reasoning in accord with 1a, above. But there is no reason why purists must agree that Hannah is justified in reasoning in accord with 1a, above. Purists who accept RSB can insist that Hannah knows that the bank will be open in virtue of *non*-occurrently believing that it will be open, and add that, since Hannah would not be justified in reasoning in accord with 1a, above, she lacks justification for occurrently believing that the bank will be open.⁶² This *sounds* strange; we are unaccustomed to the idea that one might be unjustified in occurrently believing but justified in *non*-occurrently believing. Given RSB, however, this is plausible. Since doxastic pragmatism is just the view that pragmatism follows from the correct theory of belief together with assumptions that purists will happily accept, and since purists who accept RSB will be unhappy with (4), Ross and Schroeder's argument for (8) will

⁶¹ Ibid., p. 13.

⁶² Of course, purists might also reject (2) and (5), which Ross and Schroeder similarly indorse on the basis of their *prima facie* plausibility. (See *ibid.*, pp. 13-14.)

not work as an argument for doxastic pragmatism.⁶³ This is no problem *for Ross and Schroeder*, of course. They can easily reply that, far from defending doxastic pragmatism, they are simply taking it for granted that non-doxastic pragmatism is true and then using RSB to explain *why* it is true.⁶⁴ But this *is* a problem for any doxastic pragmatist hoping to employ RSB or Ross and Schroeder's argument for (8) in defence of her position.

3. A Close Neighbour of Ross and Schroeder's Theory of Belief

But can't doxastic pragmatists employ a theory of belief very *similar* to RSB in defense of their view? According to RSB, *S* believes that *p* only if she is *defeasibly* disposed to treat *p* as true in her reasoning,⁶⁵ and this defeasibility condition is what prevents RSB from entailing pragmatism with just the addition of assumptions that purists will accept. So why can't doxastic pragmatists just drop this defeasibility condition and defend their view with NRSB, below?

(NRSB) *S* believes that *p* only if she is disposed to treat *p* as true in her reasoning.

Since *S* treats *p* as true in her reasoning just in case she evaluates her options by the same procedure as she would use to evaluate them conditional on *p*, NRSB says that *S* *believes* that *p* only if she evaluates her options by the same procedure as she would use to evaluate them conditional on *p*.⁶⁶ Since Hannah is not treating the proposition that the bank will be open as true in her reasoning in High Stakes, NRSB tells us that Hannah does not believe that the

⁶³ I consider the arguments available for (2), (4) and (5) in chapter 3, which addresses the debate between purists and *non*-doxastic pragmatists.

⁶⁴ This, in fact, is exactly what Ross and Schroeder say they are doing, minus the distinction between doxastic and non-doxastic pragmatism. (See *ibid.*, pp. 1-2.) Note, however, that RSB is at best an incomplete explanation, since, if (2), (4) or (5) is false, RSB is consistent with purism.

⁶⁵ *Ibid.*, p. 9.

⁶⁶ *Ibid.*, p. 6.

bank will be open in High Stakes. So, according to NRSB, Hannah does not *know* that the bank will be open in High Stakes. Unlike RSB, NRSB *does* entail pragmatism with just the help of assumptions that purists will happily accept.

Unfortunately, NRSB runs into the same counterexample as Ganson and Weatherson's theories of belief. In the section on Ganson and Weatherson's theories of belief, I evaluated both the act of asserting that the world is more than five minutes old and the act of asserting that the sun will rise tomorrow, and the way that I evaluated these acts is better represented by second line of reasoning, below, than by the first line of reasoning below.

- 1b. If I assert that the sun will rise tomorrow, then I will assert a truth. And if I assert that the world is more than five minutes old, I will also assert a truth.
So, ...
- 2b. If I assert that the sun will rise tomorrow and the sun *will* rise tomorrow, then I will assert a truth. But if I assert that the sun will rise tomorrow and the sun will *not* rise tomorrow, I will assert a falsehood. On the other hand, if I assert that the world is more than five minutes and the world *is* more than five minutes old, I will assert a truth. And if I assert that the world is more than five minutes old and the world is *not* more than five minutes old, I will assert a *falsehood*. The probability that the sun will rise tomorrow is extremely high, but not as high as the probability that the world is more than five minutes old. So, ...

But my reliance on the second line of reasoning above means that, while I was writing the section on Ganson and Weatherson's theories of belief, I was not treating the proposition that the sun will rise tomorrow as true in my reasoning. Since NRSB lacks RBS's defeasibility condition, NRSB says that, while I was writing that section, I did not believe that the sun will rise tomorrow. But of course, this is false; while I was writing the section on Ganson and Weatherson's theories of belief, I *did* believe that the sun will rise tomorrow. Perhaps I did not *occurently* believe that the sun will rise tomorrow, but NRSB does not restrict itself to occurrent belief. So NRSB is false.⁶⁷

Moreover, NRSB shares a problem with RSB—one that makes *both* of them look false. Consider the following platitude.

Maximal Confidence: S cannot be maximally confident that p without believing that p . That is, absolute certainty entails belief.

However exactly belief relates to confidence, the connection between belief and confidence is tight enough for Maximal Confidence to be plausible. But on both RSB and NRSB, Maximal Confidence comes out false. Just imagine that S always evaluates her options by reasoning the second way, below, rather than the first way, below.

1. The consequences of φ -ing will be A . The consequences of ψ -ing will be B .
So, ...

⁶⁷ Could a theory that *does* restrict itself to occurrent belief be useful to doxastic pragmatists? I cannot see how. From that fact that Hannah does not occurrently believe, it does not follow that she does not *know*; it only follows that she does not *occurently* know (where occurrently knowing that p is just knowing that p while occurrently believing that p). But purists will have no problem with the suggestion that, at those moments when Hannah is not occurrently believing that the bank will be open, she does not occurrently know that the bank will be open. So a theory that restricts itself to occurrent belief would be useless to pragmatists and, *a fortiori*, useless to *doxastic* pragmatists.

2. If p , then consequences of φ -ing will be A . And if $\neg p$, then the consequences of ψ -ing will be B . But if $\neg p$, the consequences of φ -ing will be C . And if $\neg p$, then the consequences of ψ -ing will be D . The probability that p is n and the probability that $\neg p$ is $(1 - n)$. So, ...

By hypothesis, as a result of her dispositions, S evaluates her options by reasoning the second way above rather than the first way above, *no matter what*—whether she is in a high stakes situation or a low stakes situation. But then, for any value of ' p ,' NRSB says that S does not believe that p . And for the same reason, since there are literally *no* circumstances where S would treat p as true in her reasoning, S is not even *defeasibly* disposed to treat p as true in her reasoning. Thus, for any value of ' p ,' RSB says that S does not believe that p . So according to both NRSB and RSB, S literally holds no beliefs. But compatibly with our stipulation that S always evaluates her options by reasoning the second way rather than the first way, above, there might be many values of ' p ' for which n equals 1. And in this case, there might be many values of ' p ' for which S is absolutely certain that p while she fails to satisfy the requirements imposed by RSB and NSRB. So RSB and NSRB deliver the implausible conclusion that absolute certainty does not suffice for belief. Thus NRSB conflicts with Maximal Confidence and seems susceptible to counterexample. Clearly enough, pragmatists cannot put NRSB to work in a successful argument for doxastic pragmatism.

4. Fantl and McGrath on Belief

The last view that we will consider is Jeremy Fantl and Matthew McGrath's. Fantl and McGrath argue that practical interests encroach on knowledge because they encroach on justified belief, and they worry that pragmatic encroachment on justified belief requires that belief be at least "weakly pragmatic."

Does [pragmatic encroachment on justified belief] require that belief be at least weakly pragmatic, where by this we will mean that that believing requires there being some φ such that one's credence in p is high enough for p to move one to φ ? ... [W]e know of no decisive argument. But we register this worry: if belief isn't at least weakly pragmatic, why is justified belief pragmatic? More carefully: if belief in p has no implications for p 's having sufficient credence to move you, why should justified belief have implications for p 's being warranted enough to justify you in anything let alone everything? If belief that p requires that p have sufficient credence to move you to φ for *some* φ or *to some restricted set of* φ , we can see how this could be, given appropriate assumptions about the normative demands 'reverberating' from one φ or a restricted class of φ to all φ . But if belief doesn't require even this limited pragmatic condition, we are left asking how justified belief could have the pragmatic condition it has.⁶⁸

Presumably, the restricted set that Fantl and McGrath mention includes only relevant values of ' φ .' But are they thinking of every relevant value of ' φ ' or some smaller percentage of

⁶⁸ Fantl and McGrath (2009), p. 140.

relevant values? Whatever exactly Fantl and McGrath have in mind, they seem committed to at least one of FM1, FM2, or FM3, below.

- (FM1) S believes that p only if S 's confidence that p is high enough for the proposition that p to move S to φ , for *some* relevant value of ' φ .'
- (FM2) S believes that p only if S 's confidence that p is high enough for the proposition that p to move S to φ , for some *high percentage* of relevant values of ' φ .'
- (FM3) S believes that p only if S 's confidence that p is high enough for the proposition that p to move S to φ , for *every* relevant value of ' φ .'

But FM1, FM2 and FM3 are all inadequate for defending doxastic pragmatism.

FM1 is clearly too weak. For example, suppose that Hannah has a friend with an idle curiosity about the bank's Saturday hours (she asks whether the bank will be open on Saturday just to make conversation), and suppose that Hannah's confidence that the bank will be open on Saturday is high enough for that proposition to be Hannah's motivating reason for telling her friend that the bank will be open on Saturday. In this case, Hannah satisfies the only necessary condition that FM1 imposes on belief, and she does so in a perfectly rational way. So FM1 lends no credibility to pragmatism. But on the other hand, FM2 and FM3 are too *strong*, since they are susceptible to the counterexample that we have already considered many times. Let s be the proposition that the sun will rise tomorrow, let m be the proposition that the sun will rise *at least one more time*, let w be the proposition that

the world is more than five minutes old, and ask, *which is more probable, m or w ?* This question presents us with exactly four options.

Option A: We can say that m is more probable than w .

Option B: We can say that m is less probable than w .

Option C: We can say that m and w are equally probable.

Option D: We can leave the question unanswered.

Option A is a relevant value of ‘ φ ,’ and, if I were maximally confident that the sun will rise tomorrow, then the proposition that the sun will rise tomorrow would move me to choose option A. After all, if I were maximally confident that the sun will rise tomorrow, I would reason as follows: *The probability that the sun will rise tomorrow is 1; so the probability that the sun will rise **at least one more time** is 1; so it is more probable that the sun will rise at least one more time than it is that the world is more than five minutes old.* But of course, since I am *not* maximally confident that the sun will rise tomorrow, my confidence that the sun will rise tomorrow is not high enough for that proposition to move me to choose option A. It follows that, right now, as I consider our options, there is a relevant value of ‘ φ ’ such that my confidence that the sun will rise tomorrow is not high enough for that proposition to move me to φ . Thus, according to FM3, I do not believe that the sun will rise tomorrow.

But FM2 *also* says that I do not believe that the sun will rise tomorrow. Again, as I consider our options, there are exactly four relevant values of ‘ φ ’—namely, options A through D, above. So consider the following claims.

(CA) My confidence that s is high enough for s to move me to choose option A.

- (CB) My confidence that s is high enough for s to move me to choose option B.
- (CC) My confidence that s is high enough for s to move me to choose option C.
- (CD) My confidence that s is high enough for s to move me to choose option D.

Is there a true claim in CA through CD? As it happens, there is not. CB is the closest, since I like option B. But CB is false. I do not choose option B because my confidence that s is high enough for s to move me to choose that option, as CB claims. Rather, I choose B because my confidence that w is high enough for w to move me to choose B, while my confidence that s is *low* enough for w to move me to choose B. And this means that my confidence that s is high enough for s to move me to φ for zero percent of the relevant values of ' φ .' So, according to FM2, I do not believe that s . Since s is the proposition that the sun will rise tomorrow, FM2 says that I do not believe that the sun will rise tomorrow. So FM2 gets the same result as FM3. But once again, I *do* believe that the sun will rise tomorrow. So FM2 and FM3 are both false.

Clearly enough, Fantl and McGrath do not give us a promising way to defend doxastic pragmatism. This is no strike against *Fantl and McGrath*, of course, since they give us every indication that they would rightly classify their view as a version of *non-doxastic* pragmatism. But these problems count decisively against any *doxastic* pragmatist who wants to employ FM1, FM2 or FM3 in defense of her view.

5. Rival Theories of Belief

We have considered seven theories of belief and found that each is either susceptible to a very mundane counterexample—my belief that the sun will rise tomorrow—or too loosely

connected to practical rationality to entail pragmatism with just the addition of assumptions that purists will happily accept. Since doxastic pragmatism is just the view that pragmatism follows from the correct theory of belief together with assumptions that purists will happily to accept, these seven theories of belief are all inadequate for defending doxastic pragmatism. How plausible is doxastic pragmatism, then? A discussion of *every* theory of belief that we might put to use in an argument for doxastic pragmatism is beyond the scope of this paper, but we have considered a representative sample and found all of the theories in that sample are inadequate for defending doxastic pragmatism. In addition, anyone hoping to defend doxastic pragmatism must confront this problem: there are plausible, independently motivated theories of belief in the literature that are clearly inconsistent with doxastic pragmatism. To give just one example, take Ernest Sosa's "affirmative" view of belief. On Sosa's view, mentally assenting to p suffices for occurrently believing that p .⁶⁹ One needn't conditionalize on p , or treat p as true in one's reasoning, or act as if p , or anything like that. In addition to finding a counterexample-free theory of belief that entails pragmatism with just the help of assumptions that purists will accept, defenders of doxastic pragmatism must also dispose of plausible, independently well-motivated theories of belief like Sosa's. This looks like a tall order. Everything considered, it seems doubtful that doxastic pragmatism is adequately motivated.

Appendices

1. Schroeder on Reasons to Withhold

⁶⁹ Sosa (2011), p. 41.

In “Stakes, Withholding, and Pragmatic Encroachment on Knowledge,” Schroeder defends a version of pragmatism that is clearly non-doxastic, and that therefore belongs with the versions of pragmatism that I discuss in “A Paradox for Justification, Knowledge and Practical Rationality.”⁷⁰ Schroeder’s argument seems to presuppose a theory of belief that is susceptible to the counterexample that we have considered many times, however, so I would like to address it here.

Schroeder distinguishes between reasons to believe and reasons to withhold, and then argues that reasons to withhold (whatever they are) cannot be evidence.

[T]he evidence is exhausted by evidence which supports p and evidence which supports $\neg p$. But the evidence which supports p is reason to believe p , and the evidence which supports $\neg p$ is reason to believe $\neg p$. Consequently the reasons to withhold must come from somewhere else. So they cannot be evidence.⁷¹

⁷⁰ Schroeder (forthcoming).

⁷¹ Ibid., p. 12. Of course, we might worry that Schroeder is mistaken that reasons to withhold cannot be evidence. Imagine three urns, each containing 100 marbles and nothing else. We know that all of the marbles in the black urn are black, we know that all of the marbles in the white urn are white, and we know that 50 marbles in the grey urn are black while the other 50 are white. We also know that one marble will be drawn at random from each urn and, consequently, we believe that the draw from the black urn will produce a black marble, we believe that the draw from the white urn will *not* produce a black marble, and we withhold with respect to the proposition that the draw from grey urn will produce a black marble. Now, *why* do we believe the first proposition, believe the negation of the second proposition, and withhold with respect to the third proposition? Speaking for myself, I believe that the draw from the black urn will produce a black marble because I know that 100% of the marbles in the black urn are black, I believe that the draw from the white urn will *not* produce a black marble because I know that 0% of the marbles in the white urn are black, and I *withhold* with respect to the proposition that the draw from the grey urn will produce a black marble because I know that 50% of the marbles in the grey urn are black. But now, why can’t I just argue as follows?

The fact that 100% of the marbles in the black urn are black is evidence with respect to the proposition that the draw from the black urn will produce a black marble, and the fact that 0% of the marbles in the white urn are black is evidence with respect to the proposition that the draw from the white urn will produce a black marble. *Vis-à-vis* the question *what is evidence with respect to what?*, there is no relevant difference between the fact/proposition pairs that I just mentioned and the following fact/proposition pair: the fact that 50% of the marbles in the grey urn are black and the proposition

What *are* the reasons to withhold? Schroeder does not offer a complete list, but he does say that costs of error can be an important reason to withhold.

[A] natural place to look for reasons to withhold is in the costs of error. When you form a belief, you take a risk of getting things wrong that you don't take by withholding. In contrast, when you withhold, you guarantee that you miss out on getting things right. So plausibly, one important source of reasons to withhold will come from the preponderance of the cost of having a false belief over the cost of missing out on having a true belief—or, as I will put it, the preponderance of the cost of *type-1 error* over *type-2 error*.⁷²

This all seems highly plausible. But so far, it lends no support to pragmatism, for the cost of a type-1 error might just be the possession of a false belief, and the cost of a type-2 error might just be the absence of a true belief. If these are the *only* costs of type-1 and type-2 errors, however, then *practical* interests seem beside the point. If these are the only costs, for then (for example) the costs of type-1 and type-2 errors will be the same in High Stakes as they are in Low Stakes. According to Schroeder, however, the costs of type-1 and type-2 errors will often go well *beyond* false belief and the absence of true belief, respectively.

that the draw from the grey urn will produce a black marble. So, the fact that 50% of the marbles in the grey urn are black is evidence with respect to the proposition that the draw from the grey urn will produce a black marble. But the fact that 50% of the marbles in the grey urn are black is a reason to withhold with respect to the proposition that the draw from the grey urn will produce a black marble. So, evidence with respect to a proposition might be a reason to withhold with respect to that proposition.

Reasoning along these lines casts doubt on Schroeder's claim that reasons to withhold cannot be evidence.

⁷² Ibid., p. 13.

[T]he most general sort of cost of type-1 error is simply mistakes that we make, when we act on a belief that turns out to be false. ... Correlatively, the most general sort of cost of type-2 error derives from the fact that sometimes we simply have to act, and ignorance doesn't help us. These two sorts of costs—of type-1 and type-2 errors—are clearly practical in nature. They derive from the connection between belief and action. ... Gratifyingly, in High Stakes the costs of type-1 error are extremely high, and the costs of type-2 error are very low, which on this picture supports the view that there are especially strong reasons to withhold in that case.⁷³

Purists should have no trouble agreeing with Schroeder that, if Hannah has decisive epistemic reason to withhold with respect to the proposition that the bank will be open, then she cannot know that the bank will be open. But why does Schroeder think that Hannah epistemically ought to withhold with respect to that proposition? Because, as indicated in the quotation above, he thinks that the high-costs of a type-1 error with respect to the proposition that the bank will be open give Hannah sufficient epistemic reason to withhold.

Now, purists may deny that the costs of a type-1 error with respect to the proposition that the bank will be open give Hannah any *epistemic* reason to withhold, so Schroeder's view cannot be properly classified as a version of doxastic pragmatism.⁷⁴ Still, other purists might agree that these reasons *are* genuinely epistemic and then resist Schroeder's argument by commenting on the nature of belief. In the quotation above, Schroeder says that the costs of a type-1 error are very high for Hannah, and he says that these costs "derive from the

⁷³ Ibid., pp. 13-14.

⁷⁴ Thanks to Schroeder for conversation on this point.

connection between belief and action.” What exactly *is* the connection between belief and action? Schroeder does not say, but his argument seems to presuppose SB, below.

(SB) S believes that p only if she is willing to act as if p (where S acts as if p just in case she acts in the manner that would be rationally optimal, conditional on p).

As we saw above, going straight home would be rationally optimal for Hannah conditional on the bank’s being open. So, according to SB, Hannah believes that the bank will be open only if she is willing to go straight home. Suppose that SB is *false*, then. In this case, Hannah can be unwilling to go straight home and yet still believe that the bank will be open. Yet if Hannah can believe that the bank will be open while she is unwilling to go straight home, then why should we think that, in High Stakes, the costs of type-1 error with respect to the proposition that the bank will be open are extremely high for Hannah? It seems that, if SB is false, Hannah can believe that the bank will be open and simultaneously take all of the precautions necessary for insuring that she does not incur the large fine. If Hannah can believe that the bank will be open and simultaneously take all of the precautions necessary for insuring that she does not incur the fine, however, then the cost of a type-1 error with respect to the proposition that the bank will be open will be no higher than forming a false belief. And this will be the case *in both Low Stakes and High Stakes*. So, if SB is false, the costs of type-1 error with respect to the proposition that the bank will be open seem to be the *same* in Low Stakes and High Stakes. It follows that, unless SB is *true*, the costs of type-1 error with respect to the proposition that the bank will be open should be no higher for

Hannah in High Stakes than they are in Low Stakes. Thus, Schroeder's argument that the costs of type-1 error increase as the stakes go up seems to presuppose SB. And since his argument for pragmatism depends on the claim that the costs of type-1 error increase as the stakes go up, his argument for pragmatism seems equally dependent on the truth of SB.

Fortunately for purists, SB looks false. First, my belief that the sun will rise tomorrow causes the same trouble for SB as it did for GB, WB, NRSB, FM2 and FM3. Again, let s be the proposition that the sun will rise tomorrow, let w be the proposition that the world is more than five minutes old, and consider the question, *which is more probable, w or s ?* We can imagine scenarios where (a), I am not willing to answer this question by asserting s , where (b), conditional on the proposition that the sun will rise tomorrow, I would maximize expected utility by asserting s , and where (c) I nevertheless believe that the sun will rise tomorrow. So SB is susceptible to the same counterexample that has plagued all but one of the views that we have considered so far. Second, in his paper with Ross defending RSB, *Schroeder* seems to give us a good reason to reject SB. Ross and Schroeder present a principle that they call "Stability."

Stability: A fully rational agent does not change her beliefs purely in virtue of an evidentially irrelevant change in her credences or preferences.⁷⁵

A component of Ross and Schroeder's argument for RSB is that RSB is consistent with Stability while theories of belief like Ganson's GB and Weatherson's WB are not.⁷⁶ Ross and

⁷⁵ Ross and Schroeder (forthcoming), p. 19 of version available on Schroeder's website.

⁷⁶ *Ibid.*, pp. 19-22.

Schroeder seem reasonable to question GB and WB on the basis of their conflict with Stability. Curiously, however, Stability also conflicts with Schroeder's SB.

By stipulative definition, one's credence in q is evidentially irrelevant to p just in case neither q nor one's credence in q provides any evidence for or against p .⁷⁷ So let d be the proposition that there will be disastrous consequences for Hannah if she does not deposit her check before Saturday afternoon. Because neither d nor Hannah's credence in d provides any evidence for or against the proposition that the bank will be open Saturday morning, d is evidentially irrelevant to that proposition. So now imagine Hannah in *Low Stakes* and notice what happens if we stipulate that Hannah is fully rational and then vary Hannah's credence in d . Manipulating just this variable, we change what rationality requires of Hannah. Once Hannah's credence in d gets sufficiently high, Hannah maximizes expected utility by *stopping at the bank*. Since Hannah is fully rational, we change what Hannah is willing to do; once her credence in d gets sufficiently high, she is not willing to go straight home. According to SB, however, Hannah believes that the bank will be open only if she *is* willing to go straight home. So, if SB is true, by varying just Hannah's credence in d , we change what Hannah *believes*. But again, d is evidentially irrelevant, and we are supposing that Hannah is fully rational. So, if SB is true, a fully rational agent may change her beliefs purely in virtue of an evidentially irrelevant change in her credences, contrary to Stability. But without SB, Schroeder's argument for pragmatism does not go through. Because SB seems to have counterexamples and conflict with Stability, we may doubt that Schroeder's argument for pragmatism succeeds.

⁷⁷ Ibid., p. 19.

2. *Math Relevant to §1*

Here is the math that I promised in footnote 13. First, we have the decision table and corresponding calculations for $U(A_s | s)$.

	s	
Decision Table for $U(A_s s)$:	Assert that	<i>Assert a truth</i>
	s	

	s	
Value Matrix for $U(A_s s)$:	Assert that	$a > 0$
	s	

	s	
Probability Matrix for $U(A_s s)$:	Assert that	1
	s	

We get $U(A_s | s)$ by multiplying the probability of s by the value of asserting a truth (which we are calling ' a '). Thus,

$$U(A_s | s) = (a \times 1) = a.$$

Next, we have the decision table and corresponding calculations for $U(A_s)$.

	s	$\neg s$		
Decision Table for $U(\mathcal{A}_s)$:	Assert that	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; padding: 5px;"><i>Assert a truth</i></td> <td style="width: 50%; padding: 5px;"><i>Assert a falsehood</i></td> </tr> </table>	<i>Assert a truth</i>	<i>Assert a falsehood</i>
<i>Assert a truth</i>	<i>Assert a falsehood</i>			
	s			

	s	$\neg s$		
Value Matrix for $U(\mathcal{A}_s)$:	Assert that	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; padding: 5px;">a</td> <td style="width: 50%; padding: 5px;">$b < 0$</td> </tr> </table>	a	$b < 0$
a	$b < 0$			
	s			

	s	$\neg s$		
Probability Matrix for $U(\mathcal{A}_s)$:	Assert that	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; padding: 5px;">$c < 1$</td> <td style="width: 50%; padding: 5px;">$1 - c$</td> </tr> </table>	$c < 1$	$1 - c$
$c < 1$	$1 - c$			
	s			

We get $U(\mathcal{A}_s)$ by multiplying the probability of s (which we are calling ' c ') by the value of asserting a truth, and adding this result to the probability of $\neg s$ multiplied by the value of asserting a falsehood (which we are calling ' b '). Thus,

$$U(\mathcal{A}_s) = ((a \times c) + (b \times (1 - c))) = d.$$

Since b is negative, d is less than a (whatever d and a happen to be). So $U(\mathcal{A}_s)$ is less than $U(\mathcal{A}_s | s)$. Thus,

$$U(\mathcal{A}_s) < U(\mathcal{A}_s | s).$$

Next, we have the decision table and corresponding calculations for $U(\mathcal{A}w)$.

		w	$\neg w$
Decision Table for $U(\mathcal{A}w)$:	Assert that	<i>Assert a truth</i>	<i>Assert a falsehood</i>
	w		

		w	$\neg w$
Value Matrix for $U(\mathcal{A}w)$:	Assert that	a	b
	w		

		w	$\neg w$
Probability Matrix for $U(\mathcal{A}w)$:	Assert that	$c < e < 1$	$1 - e$
	w		

We get $U(\mathcal{A}w)$ by multiplying the probability of w (which we are calling ' e ') by the value of asserting a truth, and adding this result to the probability of $\neg w$ multiplied by the value of asserting a falsehood. Thus,

$$U(\mathcal{A}w) = ((a \times e) + (b \times (1 - e))) = f.$$

Since b is negative, f is less than a (whatever f and a happen to be). But since e is greater than c (which is the probability of s), f is *greater* than d . Thus, $U(\mathcal{A}s)$ is less than $U(\mathcal{A}w)$, which is less than $U(\mathcal{A}s | s)$. Thus,

$$U(\mathcal{A}s) < U(\mathcal{A}w) < U(\mathcal{A}s|s).$$

Finally, we have the decision table and corresponding calculations for $U(\mathcal{A}w|s)$.

	$w \wedge s$	$\neg w \wedge s$
Decision Table for $U(\mathcal{A}w s)$:	Assert that	<i>Assert a truth</i>
	w	<i>Assert a falsehood</i>

	$w \wedge s$	$\neg w \wedge s$
Value Matrix for $U(\mathcal{A}w s)$:	Assert that	a
	w	b

	$w \wedge s$	$\neg w \wedge s$
Probability Matrix for	Assert that	$e \times c$
$U(\mathcal{A}w s)$:	w	$(1 - e) \times c$

We get $U(\mathcal{A}w|s)$ by multiplying the probability of the conjunction of w and s by the value of asserting a truth, and adding this result to the probability of the conjunction of $\neg w$ and s multiplied by the value of asserting a falsehood. Since w and s are probabilistically independent—since $\Pr(w) = \Pr(w|s)$ and $\Pr(s) = \Pr(s|w)$ —the probability of the conjunction of w and s is e multiplied by c , and the probability of the conjunction of $\neg w$ and s is $(1 - e)$ multiplied by c . Thus,

$$U(Aw|s) = ((a \times e \times c) + (b \times (1 - e) \times c)) = ((a \times e) + (b \times (1 - e))) = f.$$

Since conditionalizing on s does nothing to the probability of w , $\Pr(w|s)$ equals $\Pr(w)$, and $U(Aw|s)$ equals $U(Aw)$, which equals f . So, combining our four results, we get this:

$$U(A_s) < U(Aw) = U(Aw|s) < U(A_s|s).$$

And now, with this result, (3) and (4) are clearly false.

$$(3) \quad U(A_s) = x \equiv U(A_s|s) = x$$

$$(4) \quad U(Aw) \geq U(A_s) \equiv U(Aw|s) \geq U(A_s|s)$$

(3) is false because $U(A_s|s)$ is greater than $U(A_s)$, and so there is no value of 'x' such that $U(A_s|s)$ and $U(A_s)$ both equal x . And (4) is false because $U(Aw)$ is greater than $U(A_s)$ while $U(Aw|s)$ is *less* than $U(A_s|s)$.

Chapter 3: A Paradox for Justification, Knowledge and Practical Rationality

Purists and pragmatists disagree about the relationship between knowledge and practical interests. Pragmatists think that a difference in knowledge might follow from as little as a difference in practical interests, while purists deny this claim. The following familiar cases illustrate the difference. (Assume that they differ only insofar as the stipulated differences in Hannah's practical interests require that they differ.)

Low Stakes: Hannah is driving past the bank on Friday afternoon. She has her paycheck in hand and she plans to deposit it. But there are long lines so she considers returning to deposit her paycheck on Saturday morning. She knows that it does not matter much when she deposits the paycheck.

High Stakes: Hannah is driving past the bank on Friday afternoon. She has her paycheck in hand and she plans to deposit it. But there are long lines so she considers returning to deposit her paycheck on Saturday morning. She knows that there will be disastrous consequences if she does not deposit her check before Saturday afternoon.

By hypothesis, Hannah knows that the bank will be open on Saturday morning in Low Stakes. But many people have the intuition that she does *not* know that the bank will be open on Saturday morning in *High Stakes*. Since Low Stakes and High Stakes differ only insofar as the stipulated differences in Hannah's practical interests require that they differ, this intuition

causes trouble for purism. If it is correct, then a difference in knowledge *can* follow from a mere difference in practical interests, just as pragmatists claim. It is no surprise that High Stakes causes trouble for purism. Epistemologists are familiar with this point. What might surprise epistemologists, however, is that High Stakes causes trouble for *everyone*—pragmatists and purists alike. In this paper, I argue that High Stakes actually causes more trouble for *pragmatism* than purism. In §1, I fill in some details missing from Low and High Stakes, as they are presented above. In §2, I show how purists might defend against the pragmatist's claim that, in High Stakes, Hannah does not know that the bank will be open Saturday morning. In §3 and §4, I argue that High Stakes confronts us with a paradox on a par with the (so-called) skeptical paradox, and that the most plausible response to this paradox makes it hard for pragmatists to use High Stakes (and other cases familiar from the literature) to motivate their view. In §5, I ask how pragmatism might be motivated at all, if not by High Stakes or similar cases. I argue that High Stakes causes interesting problems for all of the most promising arguments for pragmatism, and I conclude that pragmatism is therefore insufficiently motivated.

1. Preliminaries

Our stipulation that Hannah knows that the bank will be open in Low Stakes has a number of obvious implications for Hannah's belief that the bank will be open in Low Stakes. It entails that Hannah has sufficient evidence that the bank will be open on Saturday morning, that her belief that the bank will be open is sufficiently reliably produced, that her belief is sufficiently safe and sensitive, and so on. It also entails that her belief that the bank will be open is *rational*, where this plausibly has implications for the level of *confidence* Hannah has

that the bank will be open. Plausibly, given that Hannah knows that the bank will be open in Low Stakes, Hannah must be proportioning her confidence that the bank will be open to her basic evidence that it will be open.⁷⁸ There is little controversy here. How should we understand *High Stakes*, however? *Ex hypothesi*, High Stakes differs from Low Stakes only insofar as the stipulated difference in Hannah's practical interests requires that they differ. When pragmatists say that a difference in knowledge might follow from as little as a difference in practical interests, they mean that a difference in knowledge might follow from as little as a difference *in how two people rationally ought to act*. Low Stakes and High Stakes are meant to exhibit exactly this kind of difference. So whatever exactly is at stake for Hannah in High Stakes, let's assume that Hannah has too much to lose and too little to gain for it to be rational for her to go straight home, whereas, in Low Stakes, she has enough to lose and too little to gain for it to be rational for her to *stop at the bank*. Now, given this assumption, theories of belief like Brian Weatherson's entail that Hannah believes that the bank will be open in High Stakes only if her confidence that the bank will be open is much higher than it should be, given her basic evidence.⁷⁹ So how should we flesh out our assumption that High

⁷⁸ I say "basic evidence" rather than just "evidence" to avoid complications resulting from Timothy Williamson's claim that one's evidence is all and only one's knowledge. I mean to stipulatively define 'basic evidence' so that one's basic evidence *that p* cannot include the proposition that *p*, regardless of whether one *knows* that *p*. Thus, the way I mean to stipulatively define 'basic evidence,' Hannah's basic evidence *that the bank will be open* does not include the proposition that the bank will be open, whether Hannah knows that the bank will be open or not.

⁷⁹ According to Weatherson, *S* believes that *p* iff conditionalizing on *p* does not change *S*'s answer to any relevant question. So ask which has higher expected utility for Hannah, stopping at the bank or going straight home. Since Hannah is presently deciding between stopping at the bank and going straight home, this is a relevant question. Conditional on the proposition that the bank will be open, going straight home has higher expected utility for Hannah than stopping at the bank. But unless Hannah's confidence is much higher than her evidence warrants, going straight home actually has *lower* expected utility for Hannah than stopping at the bank. So either Hannah's confidence is much higher than her evidence warrants or, on Weatherson's theory of belief, Hannah fails a necessary condition for believing that the bank will be open. For the details, see Weatherson (forthcoming), §2.

Stakes differs from Low Stakes only insofar as the stipulated difference in Hannah's practical interests requires that they differ?

Here we may distinguish between *doxastic* and *non-doxastic* versions of pragmatism. A version of pragmatism is *doxastic* (I will say) just in case it forces purists to take a stand on the nature of belief, and *non-doxastic* just in case it allows purists to stay neutral on the nature of belief. If any version of doxastic pragmatism is correct, then pragmatism follows from the correct theory of belief with just the addition of assumptions that purists will happily accept. If some version of *non-doxastic* pragmatism is correct, however, then pragmatism does not follow from just the correct theory of belief and assumptions that purists will happily accept. Instead, pragmatism follows at least in part from controversial premises about truth, justification, or knowledge itself.⁸⁰ I will focus solely on *non-doxastic* pragmatism in this paper, since the nature of belief deserves a paper of its own.⁸¹ In what follows, I will assume that doxastic pragmatism is false and use 'pragmatism' as short of 'non-doxastic pragmatism.'

Given that doxastic pragmatism is false, our stipulation that High Stakes differs from Low Stakes only insofar as the stipulated difference in Hannah's practical interests requires that they differ entails both of the following details. First, in both Low Stakes and High Stakes, Hannah believes that the bank will be open. Second, in both Low Stakes and High Stakes, Hannah is proportioning her confidence that the bank will be open to her basic

⁸⁰ By 'justification' in this context, I just mean whatever makes the difference between knowledge and true belief that falls short of knowledge.

⁸¹ I argue against doxastic pragmatism in "Pragmatism and the Nature of Belief," and develop my own theory of belief in "The Nature of Belief and the Value of Knowledge" (chapters 2 and 5 of my dissertation, respectively).

evidence that the bank will be open. This latter detail will help to insure that, in High Stakes, Hannah is rationally obliged to stop at the bank, which will be relevant below.

2. A Purist Response to High Stakes

2.1. Problems for K

Let ‘K’ name the proposition that, in High Stakes, Hannah knows that the bank will be open. The problems facing K are well known. First, we have stipulated that, in High Stakes, Hannah cannot rationally go straight home (as she would prefer to, conditional on the proposition that the bank will be open). But how can this be, given that K is true? The truth of K seems to conflict with our stipulation that Hannah cannot rationally go straight home. Second, knowledge-action principles like Hawthorne and Stanley’s “Reason-Knowledge Principle” (below) might cause trouble for K.

(RKP) Where S ’s choice is p -dependent, it is appropriate for S to treat the proposition that p as a reason for acting iff S knows that p .⁸²

According to Hawthorne and Stanley, “a choice between options $x^1 \dots x^n$ is p -dependent iff the most preferable of $x^1 \dots x^n$ conditional on the proposition that p is not the same as the most preferable of $x^1 \dots x^n$ conditional on the proposition that $\neg p$.”⁸³ Hannah is choosing between stopping at the bank and going straight home, and the most preferable of these options conditional on the proposition that the bank will be open is not the same as the most preferable of these options conditional on the proposition that the bank *will not* be

⁸² Hawthorne and Stanley (2008), p. 578.

⁸³ Ibid.

open. So, Hannah's choice depends on the proposition that the bank will be open tomorrow. Thus, RKP tells us that, if Hannah knows that the bank will be open tomorrow, then it is appropriate for her to treat the proposition that the bank will be open as a reason for acting. And since it is *not* appropriate for Hannah to treat the proposition that the bank will be open as a reason for acting (we may argue this, at least), RKP tells us that Hannah *does not* know that the bank will be open. Similar problems for K arise from Fantl and McGrath's KJ principle. Third, K conflicts with our sense that, if Hannah were to calculate the expected utilities of going straight home and stopping at the bank, she must include the possibility that the bank will not be open in her calculations. As Brian Weatherson argues, if *S* knows that some possibility is not actual, then it is legitimate for *S* to leave that possibility off of her decision table.⁸⁴ Since it is not legitimate for Hannah to leave the possibility that the bank will *not* be open off of her decision table, Hannah does not (on Weatherson's view) know that that possibility is not actual, and thus she does not know that the bank *will* be open. Fourth, K conflicts with defensible coherence requirements on knowledge. According to Weatherson, if *S* believes that *p*, if she prefers φ -ing to ψ -ing conditional on *p*, and if she *actually* prefers ψ -ing to φ -ing, then *S*'s belief that *p* does not cohere well enough with the rest of her cognitive system for her belief to count as an item of knowledge. But we can stipulate that Hannah prefers going straight home to stopping conditional on the bank's being open, but *actually* prefers stopping to going straight home.⁸⁵ And in this case, Weatherson's coherence principle tells us that Hannah's belief that the bank will be open does not cohere well enough with the rest of her cognitive system for her belief that the bank will be open to

⁸⁴ Weatherson (forthcoming), p. 6.

⁸⁵ *Ibid.*, p. 27.

count as an item of knowledge. Fifth, and finally, K conflicts with Schroeder's claim that Hannah's has sufficient epistemic reasons to withhold with respect to the proposition that the bank will be open in High Stakes that she lacks in Low Stakes.⁸⁶ According to Schroeder, epistemic reasons to withhold cannot come from *evidence*. Instead, they often come from the consequences of acting on a belief that turns out to be false. Since in High Stakes, it would be disastrous for Hannah to act on the belief that the bank will be open if that belief turns out to be false, Hannah has decisive epistemic reason to withhold in High Stakes.⁸⁷

To reasonably resist pragmatism, purists must resist all of these arguments against K. Here is how purists might do this.

2.2. Plausible Purist Responses to the Difficulties Facing K

The first obstacle to endorsing K is that, while we have stipulated that Hannah cannot rationally go straight home, it is hard to see how it could be irrational for Hannah to go straight home *if K is true*. We can unpack this obstacle in decision-theoretic terms. Suppose that Hannah has stopped to calculate the expected utilities of stopping at the bank and going straight home, and compare the following decision tables.

Table 1

	The bank will be open
t bank	<i>Wait in line</i>

⁸⁶ Schroeder (forthcoming), p. 12.

⁸⁷ Ibid. pp. 12-14.

Go straight home	<i>Avoid the lines</i>
------------------	------------------------

Table 2

	The bank will be open	The bank will not be open
Stop at bank	<i>Wait in line</i>	<i>Avoid huge fine</i>
Go straight home	<i>Avoid the lines</i>	<i>Incur huge fine</i>

On Table 1, the expected utilities of stopping at the bank and going straight home will be very easy for Hannah to calculate. They will just equal the values of waiting in line and avoiding the lines, respectively. Since it is better for Hannah to avoid the lines than wait in them, if Hannah were to run an expected utility calculation on Table 1, her expected utility calculation would say that going straight home has *higher* expected utility than stopping at the bank. So, if Hannah relied on Table 1 rather than Table 2 in her expected utility calculation, she would get the result that going straight home has higher expected utility than stopping at the bank. *But if Hannah knows that the bank will be open, then why shouldn't Hannah rely on Table 1 in her expected utility calculation?* And if it *is* appropriate for Hannah to rely on Table 1 in her expected utility calculation, then why is it irrational for Hannah to go straight home, since an expected utility calculation on Table 1 would tell Hannah that going straight home has the highest expected utility? Purists cannot simply say that Hannah knows (or is in position to know) that going straight home has *lower* expected utility than stopping at the bank. For if it

is appropriate for Hannah to rely on Table 1 in her expected utility calculation, then Hannah plausibly *does not* know that going straight home has lower expected utility than stopping at the bank. To the contrary, if it is appropriate for Hannah to rely on Table 1 in her expected utility calculation, then it seems that Hannah is in position to know that going straight home has *higher* expected utility than stopping at the bank. So how can purists acknowledge that it would be irrational for Hannah to go straight home while simultaneously endorsing K?

Here purists run into the third obstacle to endorsing K: Weatherson's claim that,

- (1) if S were to calculate the expected utility of some course of action, it would be legitimate for her to leave a state of affairs off of her decision table iff she knows that that state of affairs does not obtain.⁸⁸

Purists can address both the first obstacle to endorsing K and this obstacle at the same time. Weatherson clearly intends (1) to apply only to *relevant* states of affairs. It will prove instructive, however, to consider (1) apart from any restriction to relevant states of affairs. As it stands above, without any restriction to relevant states of affairs, (1) entails a conditional that is clearly mistaken: if S were to calculate the expected utility of some course of action, then it would be legitimate for her to leave a state of affairs off of her decision table *only if* she knows that that state of affairs does not obtain. To see why this conditional is mistaken, imagine Hannah in the car as she approaches the bank. She is choosing between stopping at the bank and going straight home. If Hannah were to perform an expected utility calculation on these options, should she rely on Table 2, above, or Table 3, below?

⁸⁸ Weatherson (forthcoming), p. 6.

Table 3

	Bank will be open; Obama is wearing a red shirt	Bank will be open; Obama is <i>not</i> wearing a red shirt	Bank will <i>not</i> be open; Obama is wearing a red shirt	Bank will <i>not</i> be open; Obama is <i>not</i> wearing a red shirt
Stop at bank	<i>Wait in line</i>	<i>Wait in line</i>	<i>Avoid huge fine</i>	<i>Avoid huge fine</i>
Go straight home	<i>Avoid the lines</i>	<i>Avoid the lines</i>	<i>Incur huge fine</i>	<i>Incur huge fine</i>

Even if Hannah has no idea whether Obama is wearing a red shirt, Hannah should clearly rely on Table 2 rather than Table 3. And the reason is obvious: whether Obama is wearing a red shirt makes no difference to Hannah's choice between stopping at the bank and going straight home, and (we may suppose) Hannah can see this up front. She does not need to perform an expected utility calculation on Table 2, perform an expected utility calculation on Table 3, and then compare the results to see that the color of Obama's shirt makes no difference to her choice between stopping at the bank and going straight home.

Here we may distinguish between two ways that a state of affair might fail to make a difference to an expected utility calculation. A state of affairs might fail to make *any difference at all* to an expected utility calculation, and a state of affairs might fail to make any *relevant* difference to an expected utility calculation. Let's stipulate that a state of affairs makes *no difference at all* to an expected utility calculation just in case the options under consideration

have exactly the same expected utilities on the table including that state of affairs as they have on the table excluding that state of affairs. Whether Obama is wearing a red shirt makes no difference at all to Hannah's expected utility calculation. In contrast, let's stipulate that a state of affairs makes no *relevant* difference to an expected utility calculation just in case the options under consideration have exactly the same *utility rankings* on the table including that state of affairs as they have on the table excluding that state of affairs. That is, a state of affairs makes no relevant difference to an expected utility calculation just in case the option with the highest expected utility on the table *including* that state of affairs has the highest expected utility on the table *excluding* that state of affairs (and, if relevant, the option with the second highest expected utility on the table *including* that state of affairs has the second highest expected utility on the table *excluding* that state of affairs, and so on). Given how we just defined *making no difference at all* and *making no relevant difference*, a state of affairs that makes no difference at all to an expected utility calculation will also make no relevant difference to that expected utility calculation.

Here is why it is important to distinguish between states of affairs that make no difference at all to an expected utility calculation and states of affairs that make no *relevant* difference to an expected utility calculation. Suppose that Hannah has a very slight interest in getting a free mint at the bank. In this case, we may compare Table 2, above, to Table 4, below.

Table 4

Bank will be open; gives out mints Bank will be open; bank does not give out mints Bank will *not* be open; bank gives out mints Bank will *not* be open; bank does *not* give out free mints

Stop at bank	<i>Wait in line but get free mint</i>	<i>Wait in line without getting free mint</i>	<i>Avoid huge fine and get free mint</i>	<i>Avoid huge fine but don't get free mint</i>
Go straight home	<i>Avoid the lines and get free mint</i>	<i>Avoid the lines but don't get free mint</i>	<i>Incur huge fine and don't get free mint</i>	<i>Incur huge fine and don't get free mint</i>

Since it is much more important to Hannah that she avoids the large fine than it is that she gets the free mint, stopping at the bank has higher expected utility for Hannah than going straight home on both Table 2 and Table 4, and *by roughly the same margin* on both Table 2 and Table 4. This means that, while the prospect of getting a free mint makes *some* difference to Hannah's expected utility calculation, it is transparent to Hannah that it makes no *relevant* difference. And intuitively, if Hannah can see that adding 'the bank gives out free mints' and 'the bank does not give out free mints' to her decision table would make no relevant difference, then Hannah can leave 'the bank gives out free mints' and 'the bank does not give out free mints' off of her decision table, regardless of whether she *knows* that the bank gives out free mints (or does not give out free mints).

Where are we going with all of this? We are taking issue with (1), as it is formulated above with no restriction to relevant states of affairs. While (1) says that S can legitimately leave a state of affairs off of her decision table iff she knows that that state of affairs does not obtain, we have arrived at the view that, if S knows that including a state of affairs on her decision table would make no relevant difference to her expected utility calculation, then S can legitimately leave that state of affairs off of her decision table. Weatherson's (1) is a biconditional while our principle is just a conditional. But the appropriately qualified reverse direction of our principle *also* looks plausible. That is, plausibly, it is legitimate for S to leave a state of affairs off of her decision table *only if* S knows that including that state of affairs on her decision table would make no relevant difference to her expected utility calculation. After all, in the normal case, we would *add* a state of affairs to a decision table precisely because, by our lights, it *might* make a relevant difference to our expected utility calculation. So the following biconditional is at least *prima facie* plausible.

- (2) If S is performing an expected utility calculation, it is legitimate for her to leave a state of affairs off of her decision table iff she knows that including that state of affairs on the decision table would make no relevant difference to her expected utility calculation.⁸⁹

⁸⁹ A natural objection to (2) says that *knowledge* is not the right relation to focus on; rather, (2) should be recast in terms of justified belief. I think purists should be happy to grant this objection. For (first) the arguments that follow can be recast to incorporate a version of (2) that focuses on justified belief rather than knowledge, and (second) whatever reason we have to think that (2) should be recast in terms of justified belief, we also have for thinking that (1) should be recast in terms of justified belief. (For reasons to think that knowledge *is* the right relation to build principles like (1) and (2) around, see Weatherson (forthcoming), pp. 4-10.)

With (2) in hand, we may return to the first obstacle to endorsing K: that it is difficult to see how it could be irrational for Hannah to go straight home, given that she knows that the bank will be open. Consider the bearing of (2) on Hannah's situation in High Stakes. If Hannah were to leave 'the bank will not be open' off of her decision table, she would get Table 1, above. And as we saw, an expected utility calculation on Table 1 would tell Hannah that stopping at the bank has *lower* expected utility than going straight home, since Hannah assigns negative value to waiting in line and positive value to skipping the lines. So according to (2), it is legitimate for Hannah to leave 'the bank will not be open' off of her decision table iff Hannah knows that an expected utility calculation on a table that *includes* 'the bank will not be open' would *also* say that stopping at the bank has lower expected utility than going straight home. But by including 'the bank will not be open' on her decision table, Hannah would get Table 2, above. So according to (2), it is legitimate for Hannah to leave 'the bank will not be open' off of her decision table iff Hannah knows that an expected utility calculation on Table 2 would say that stopping at the bank has lower expected utility than going straight home.

So, can Hannah know that an expected utility calculation on Table 2 would say that stopping at the bank has lower expected utility than going straight home? Presumably, she *cannot*, and purists can rely on this fact to defend K. To run an expected utility calculation on Table 2, Hannah must fill in the corresponding probability and value matrices, below.

Probability Matrix for Table 2

	The bank will be open	The bank will not be open
Stop at bank	<i>Probability that the bank will be open</i>	<i>Probability that the bank will not be open</i>
Go straight home	<i>Probability that the bank will be open</i>	<i>Probability that the bank will not be open</i>

Value Matrix for Table 2

	The bank will be open	The bank will not be open
Stop at bank	<i>Negative value of waiting in line</i>	<i>Positive value of avoiding the huge fine</i>
Go straight home	<i>Positive value of avoiding lines</i>	<i>Negative value of incurring the huge fine</i>

To fill in these probability and value matrices, Hannah must assign probabilities to the proposition that the bank will be open and its negation, and then assign values to the four consequences listed in Table 2. Crucially, Hannah's expected utility calculation on Table 2 will produce the result that stopping at the bank has lower expected utility than going straight home only if Hannah either assigns the proposition that the bank will be open and its negation *probabilities that she knows they do not have* (like 1 and 0, respectively) or assigns the four consequences listed in Table 2 *values that she knows that they do not have* (like -\$1,000 for incurring the fine and -\$10,000 for *avoiding* the fine). But given that an expected utility calculation on Table 2 would say that stopping at the bank has lower expected utility than

going straight home only if Hannah knowingly assigns *inaccurate* probabilities or values to the propositions and consequences in question, Hannah cannot *know* that an expected utility calculation on Table 2 would say that stopping at the bank has lower expected utility than going straight home. And in this case, (2) tells us that it is *not* legitimate for Hannah to leave the ‘the bank will not be open’ off of her decision table. So purists might reasonably maintain that, since K creates no problem for the stipulation that Hannah has a pretty good idea what the correct probabilities and values are, (2) will tell us that it is not legitimate for Hannah to leave ‘the bank will not be open’ off of her decision table *even if* K is true.

So, why is it irrational for Hannah to go straight home, even while she knows that the bank will be open? Because (purists might argue) an expected utility calculation on Table 2 would say that stopping at the bank has lower expected utility than going straight home only if Hannah fills in her probability and value matrices with probabilities and values that she knows that the relevant propositions and consequences do not have; thus Hannah cannot know that an expected utility calculation on Table 2 would say that stopping at the bank has lower expected utility than going straight home; and thus—according to (2)—it is not legitimate for Hannah to leave ‘the bank will not be open’ off of her decision table. But now, since Hannah cannot leave ‘the bank will not be open’ off of her decision table, a correctly executed expected utility calculation would tell Hannah that she had better stop at the bank. And since it is part of High Stakes that Hannah sees this (or is at least in good position to see this), Hannah cannot rationally go straight home. Since Hannah would have this barrier to rationally going straight home *even if she knew that the bank will be open* (purists can argue), it is no longer mysterious how it could be irrational for Hannah to go straight home even while

she knows that the bank will be open. The plausibility of this line of reasoning gives purists a way to overcome the first obstacle to endorsing K.⁹⁰

Now we may return to Weatherson's claim that,

- (1) if S were to calculate the expected utility of some course of action, it would be legitimate for her to leave a state of affairs off of her decision table iff she knows that that state of affairs does not obtain.⁹¹

Since it *would not* be legitimate for Hannah to leave 'the bank will not be open' off of her decision table, (1) tells us that Hannah does not know that this state of affairs will not obtain. And since Hannah cannot know that the bank will be open while she *fails* to know that that this state of affairs will not obtain, K is false if (1) is true. So (1) causes problems for K.

How should purists respond to (1)? The *left-to-right* direction of (1) is no threat to purism; only the *right-to-left* direction of (1) gives us reason to doubt K. So why does Weatherson

⁹⁰ The closest Weatherson comes to addressing (2) is at *ibid.*, p. 9, where he offers an argument against *this* principle:

- (A) It is legitimate for S to leave the possibility that $\neg p$ off of her decision table iff p is true and treating $\text{Pr}(p)$ as 1 rather than its actual value doesn't change what S should do.

Weatherson provides a scenario where p is true and treating $\text{Pr}(p)$ as 1 rather than its actual value doesn't change what S should do, and where (according to Weatherson) it would clearly be *illegitimate* for S to leave the possibility that $\neg p$ off of her decision table. It is hard to apply Weatherson's argument against (A) to (2), however, since Weatherson does not tell us whether, in his scenario, S knows that including the possibility that $\neg p$ on her decision table would make no relevant difference to her expected utility calculation. For my own part, if S *does* know that including the possibility that $\neg p$ on her decision table would make no relevant difference to her expected utility calculation, I see no reason why purists cannot insist that it *is* legitimate for S to leave the possibility that $\neg p$ off of her decision table—in which case Weatherson is mistaken that it is *not* legitimate for S to leave the possibility that $\neg p$ off of her decision table. But on the other hand, if S *does not* know that including the possibility that $\neg p$ on her decision table would make no relevant difference to her expected utility calculation, then Weatherson and (2) *agree*; they both say that it is *not* legitimate for S to leave the possibility that $\neg p$ off of her decision table. So, either way, purists can deny that Weatherson's argument against (A) causes any problems for (2).

⁹¹ Weatherson (forthcoming), p. 6.

think that the right-to-left direction of (1) is correct? His argument here is simple. He imagines a professor ('Dec') who is teaching her students about decision theory *via* the following decision table.

Table 5

	S1	S2
B1	<i>Lose \$200</i>	<i>Win \$1,000</i>
B2	<i>Lose \$100</i>	<i>Win \$1,500</i>

B1 and B2 are bets, and S1 and S2 are states of affairs. Dec wants her students to see that B2 is a better bet than B1. But one of her students ('Dom') thinks that Dec should put a third state of affairs on the table ('S3') in which everyone who chooses B1 will go to heaven while everyone who chooses B2 will go to hell. Because Dom thinks that the table should also include S3, and because his confidence that S3 will not obtain is not nearly as high as Dec's confidence that S3 will not obtain, Dom thinks that B1 is actually a better bet than B2. Now Weatherson takes it for granted that Dec *can* legitimately leave S3 off of the decision table, and he asks what explains this fact. Weatherson's answer is that Dec can leave S3 off of the table because she knows that S3 will not obtain. Weatherson thinks that Dec's knowledge of *the proposition that S3 will not obtain* suffices for the legitimacy of her leaving S3 off of the table, and he thinks that this explains why Dec can legitimately leave S3 off of the decision table.

Weatherson gets the right-to-left direction of (1) by generalizing from his explanation of the legitimacy of Dec's leaving S3 off of her decision table.⁹²

Weatherson's argument here seems plausible enough. But with (2) in hand, purists can reasonably give an *alternative* explanation of the legitimacy of Dec's leaving S3 off of her decision table, or just deny that it *is* legitimate for Dec to leave S3 off of her decision table. Either Dec knows that including S3 on her decision table would make no relevant difference to her expected utility calculation, or she does not. If Dec *does* know that including S3 on her decision table would make no relevant difference to her expected utility calculation, then purists can use *this* fact to explain why it is legitimate for Dec to leave S3 off of her decision table. They need not agree with Weatherson that Dec's knowledge of the proposition *that S3 will not obtain* suffices for the legitimacy of her leaving S3 off of the table. But on the other hand, if Dec *does not* know that including S3 on her decision table would make no relevant difference to her expected utility calculation, then purists can reasonably insist that it is *not* legitimate for Dec to leave S3 off of her decision table. After all, if Dec fails to know that including S3 on her decision table would make no relevant difference to her expected utility calculation, then, *by Dec's own lights*, Dom might be *right*, B1 might actually be a better bet than B2. But if purists can reasonably insist that, if Dec does not know that including S3 on her decision table would make no relevant difference to her expected utility calculation, then it is not legitimate for Dec to leave S3 off of her decision table, then purists need not agree with Weatherson that Dec's knowledge of the proposition *that S3 will not obtain* suffices for the legitimacy of her leaving S3 off of the table. So, either way, purists can resist

⁹² Ibid., pp. 5-7.

Weatherson's argument for the right-to-left direction of (1) by denying that Dec's knowledge of the proposition that S3 will not obtain suffices for the legitimacy of her leaving S3 off of the table. Thus purists have means of defending K against Weatherson's argument for (1).⁹³

The second obstacle to a satisfying endorsement of K was potential conflict between K and knowledge-action principles like Hawthorne and Stanley's RKP, below.

(RKP) Where S 's choice is p -dependent, it is appropriate for S to treat the proposition that p as a reason for acting iff S knows that p .⁹⁴

A first thing to note about RKP is that, as Hawthorne and Stanley have formulated it here, it seems clearly mistaken. According to Hawthorne and Stanley's definition of p -dependence, a choice between options $x^1 \dots x^n$ is p -dependent iff the most preferable of $x^1 \dots x^n$ conditional on p differs from the most preferable of $x^1 \dots x^n$ conditional on $\neg p$.⁹⁵ So RKP says that, if my choice between $x^1 \dots x^n$ is p -dependent and I know that p , then it is appropriate for me to treat the proposition that p as a reason for acting. Acting how, exactly? While Hawthorne and Stanley do not say, they do make it clear that, if my choice between $x^1 \dots x^n$ is p -dependent and I know that p , then it is appropriate for me to treat the proposition that p as a reason to choose whichever option I prefer most among $x^1 \dots x^n$, conditional on p .⁹⁶ Suitably modified, RKP is plausible. It runs into difficulties as Hawthorne and Stanley state it, however. To see why, suppose that I am a judge who prefers to punish the innocent

⁹³ Again, anyone who doubts that *knowledge* is the right relation to focus on can reject (1) on these grounds and then recast the above explanation for why it is (or is not) legitimate for Dec to leave S3 off of her decision table in terms of *justification* (or whatever other relation one favors).

⁹⁴ Hawthorne and Stanley (2008), p. 578.

⁹⁵ Ibid.

⁹⁶ Ibid.

and acquit the guilty, and consider my choice between punishing Smith and acquitting him. Whether or not Smith is innocent makes a difference to my preferences: conditional on the proposition that Smith is innocent, I prefer to punish him; and conditional on the proposition that he is *not* innocent, I prefer to acquit him. So it follows from Hawthorne and Stanley's definition of *p*-dependence that my choice between punishing Smith and acquitting him is dependent on the proposition that Smith is innocent. But now, because I prefer to punish Smith (conditional on the proposition that he is innocent), RKP tells us that, if I know that Smith is innocent, then it is appropriate for me to treat the proposition that he is innocent as a reason to *punish* him. Obviously, this is an unacceptable result. If I know that Smith is innocent, then it might be *epistemically* appropriate for me to treat the proposition that he is innocent as a reason to punish him, but it is surely not *unqualifiedly* appropriate for me to treat the proposition that he is innocent as a reason to punish him. So to get a plausible principle out of RKP, we must stipulate that the propriety in question is epistemic propriety. Call the principle that we get when we *do* stipulate that the propriety in question is epistemic propriety, 'RKP'.' Now, according to RKP', Hannah knows that the bank will be open only if it is *epistemically* appropriate for her to treat the proposition that the bank will be open as a reason for acting.

RKP' faces difficulties of its own, but here I want to suppose that it is true.⁹⁷ This supposition raises the question, *what does epistemic propriety amount to?* More to the point, *even if RKP' is true, aren't we in better position to determine whether Hannah knows that the bank will be open than we are in to determine whether it would be epistemically appropriate for her to treat the proposition that*

⁹⁷ See, for example, Brown (2008) and (2012), Cohen (2012), Neta (2007), and Reed (2010).

that the bank will be open as a reason for acting? Even if purists are not entirely persuaded of Timothy Williamson's "knowledge first" outlook on epistemology, they might reasonably maintain that we put the cart before the horse by first asking whether it is epistemically appropriate for Hannah to treat the proposition that the bank will be open as a reason for acting, and then asking whether Hannah knows that it will be open. Consider the following case.

Tina asks Mark and Hector to help her select beautiful music for an elegant party that she is throwing. Mark and Hector are considering a particular song; Hector thinks that it is beautiful but Mark disagrees. To persuade Hector that the song is not beautiful, Mark reasons out loud, "if a song is beautiful then it would be aesthetically appropriate to play it at an elegant party. It would not be aesthetically appropriate to play this song at Tina's party, so this song is not beautiful." Hector is unfamiliar with aesthetic propriety, so he asks Mark what it amounts to. Mark responds, "I do not know, exactly. I just heard Tina say that it is always aesthetically appropriate to play beautiful music at an elegant party." Hector responds, "In this case, since we are generally pretty good at telling whether a song is beautiful, it would make sense for us to decide whether this song is beautiful and then, on the basis of our decision, determine whether it would be aesthetically appropriate to play it at the party. At least, it would make a lot more sense for us to proceed this way than for us to decide whether it would be aesthetically appropriate to play the song at the party and then use that conclusion to determine whether the song is beautiful."

Since they are in much better position to determine whether the song is beautiful than they are in to determine whether it would be aesthetically appropriate to play it, Hector is right; Mark is getting things backwards. But the same points hold for the argument from RKP' to the conclusion that Hannah does not know that the bank will be open. Even if RKP' is true, purists can maintain that we are in better position to determine whether Hannah knows that the bank will be open than we are in to determine whether it would be epistemically appropriate for her to treat the proposition that the bank will be open as a reason for acting. So we cannot argue that, because it would be epistemically inappropriate for Hannah to treat the proposition that the bank will be open as a reason for acting, Hannah must not know that the bank will be open. So RKP is (technically speaking) false, and, when we make the obvious repair—when we add that the propriety in question is *epistemic* propriety—we get a principle that we cannot put to use in any decisive way against purists who want to endorse K.

Perhaps, however, other knowledge-action principles put us in better position to reject K. Consider Jeremy Fantl and Matthew McGrath's KJ, below.

(KJ) If *S* knows that *p*, then *p* is warranted enough to justify *S* in φ -ing, for any value of ' φ .'⁹⁸

As Fantl and McGrath make clear, KJ does not say that, if *S* knows that *p*, then *S* is justified in φ -ing, for any value of ' φ .' KJ is considerably more sensible than that. Suppose that we are leaving a party at a friend's house; you ask whether we need to stop for gas, and I say that we

⁹⁸ Fantl and McGrath (2009), p. 66.

have enough to get home. When I say this, I am not telling you that the amount of gas in our tank entails that we will get home. Instead, I am telling you that, if we do not get home, it will not be *for lack of gas*. Similarly, Fantl and McGrath are not saying that, if S knows that p , then the amount of warrant that she has for p entails that she is justified in φ -ing, for any value of ' φ .' They are saying that, if S knows that p and she is *not* justified in φ -ing, it is not *for lack of warrant with respect to p* . Fantl and McGrath cash out “warranted enough” in terms of epistemic position: for any value of ' φ ,' p is warranted enough to justify S in φ -ing just in case no weaknesses in S 's epistemic position with respect to p stand in the way of p 's justifying S in φ -ing.⁹⁹ And while Fantl and McGrath do not say exactly what *standing in the way* amounts to, they do tell us that, *if* we can vary whether p justifies S in φ -ing by simply raising the strength of her epistemic position with respect to p , then weaknesses in S 's epistemic position with respect to p *do* stand in the way of p 's justifying S in φ -ing.¹⁰⁰ Applying all this to Hannah, KJ tells us that, if Hannah is not justified in going straight home, but she would be if she were in a certain stronger epistemic position with respect to the proposition that the bank will be open, then Hannah does not know that the bank will be open.

KJ is *prima facie* plausible, and it definitely delivers a negative verdict on K. But purists might reasonably insist that KJ is susceptible to counterexample. Consider Hannah in *Low Stakes*, let ' b ' name the proposition that the bank will be open on Saturday, and suppose that, in Low Stakes, Hannah ratchets her confidence that the bank will *still exist* on Saturday all the

⁹⁹ Ibid, p. 64-6.

¹⁰⁰ Ibid., p. 67.

way up to 1, so that she is now *maximally* confident that the bank will still exist on Saturday.¹⁰¹ In this case, the following proposition is a substitution instance of KJ:

If Hannah knows that the bank will be open on Saturday, then *b* (the proposition that the bank will be open on Saturday) is warranted enough to justify Hannah in being maximally confident that the bank will *still exist* on Saturday.

Since, in Low Stakes, Hannah knows that the bank will be open on Saturday, KJ tells us that *b* is warranted enough to justify Hannah in being maximally confident that the bank will still exist on Saturday. But to all appearances, *b is not* warranted enough to justify Hannah in being maximally confident that the bank will still exist on Saturday. After all, we can vary whether *b* justifies Hannah in being maximally confident that the bank will still exist on Saturday by simply raising the strength of her epistemic position with respect to *b*. To do this, all we have to do is raise the strength of Hannah's epistemic position with respect to *b* until Hannah is justified in being maximally confident *that b*—until she is justified in being maximally confident that the bank will be open on Saturday. For once Hannah is justified in being maximally confident that the bank will be open on Saturday, noncontroversial assumptions guarantee that Hannah is also justified in being maximally confident that the bank will still exist on Saturday. So to all appearances, while Hannah knows that the bank will be open on Saturday (in Low Stakes), *b is not* warranted enough to justify Hannah in being maximally confident that the bank will still exist on Saturday. Since KJ tells us that, if

¹⁰¹ Note that Hannah is maximally confident that the bank will still exist on Saturday only if, no matter how finely we slice degrees of confidence, Hannah is at least as confident that the bank will still exist on Saturday as you and I are that $1 + 1 = 2$, that either Obama is President or he is not, and so on, for the most obvious of necessary truths.

Hannah knows that the bank will be open on Saturday, then *b* is warranted enough to justify Hannah in being maximally confident that the bank will still exist on Saturday, purists can insist that KJ is false.

Fantl and McGrath address an objection related to this one. According to the objection that they address, KJ entails infallibilism. Here is that objection: *If KJ is true, then, if you know that p , p is warranted enough to justify you in being certain that p . But p is warranted enough to justify you in being certain that p only if p does justify you in being certain that p . After all, if p were epistemically certain for you, then it would justify you in being certain that p . And if p is warranted enough to justify you in being certain that p , then p 's being short of epistemic certainty for you can't make a difference to whether you are justified in being certain that p . Therefore, if you know that p , you should be certain that p . Surely, though, you should only be certain that p if there is no chance that $\neg p$. Therefore, if you know that p , there is no chance that $\neg p$.* Fantl and McGrath find fault with one premise in this argument—the premise according to which, if p were epistemically certain for you, then p would justify you in being certain that p . Here is what they say about this premise.

The fallibilist can quite reasonably make the following response: p , regardless of its degree of certainty, never justifies you in being certain of it. Even if we bump up your epistemic position with respect to p so that there is a zero-chance that $\neg p$, p would not be a reason you have for believing there is a zero-chance that $\neg p$. What would the reasoning look like: p , so there is no chance that $\neg p$? This is clearly fallacious. To have a reason which justifies certainty, you need a reason to think there is no epistemic chance that $\neg p$, and p is not such a reason. If p is bumped up to certainty, you may well be justified in being certain that p . But it won't be p that's

justifying you, but rather some set of facts about the strength of your evidence for p .

In general, p does not justify assigning to p a credence of 1.¹⁰²

The reasoning in this passage looks quite plausible. It seems correct that p cannot justify you in being certain *of itself*. But note: the reasoning in this passage is also irrelevant to the argument that we are imagining purists wielding against KJ. These purists are not claiming that, in Low Stakes, even while Hannah knows that the bank will be open on Saturday, weaknesses in Hannah's epistemic position with respect to b stand in the way of b 's justifying Hannah in being maximally confident *that b* . These purists are claiming that, in Low Stakes, even while Hannah knows that the bank will be open on Saturday, weaknesses in Hannah's epistemic position with respect to b stand in the way of b 's justifying Hannah in being maximally confident that the bank will *still exist* on Saturday—which is not the same proposition as b . Fantl and McGrath's observation that p cannot justify anyone in being certain *that p* is not apposite to the claim that we are imagining purists forwarding. So it seems that the objection to KJ stands. And in this case, it seems doubtful that Hawthorne and Stanley's RKP, or Fantl and McGrath's KJ, causes decisive problems for purists who want to endorse K.

The fourth obstacle for endorsing K was that K conflicts with coherence requirements on knowledge that Weatherson has recently defended. Weatherson endorses the following coherence requirement.

¹⁰² Ibid., p. 225.

- (CR) If S believes that p , if she prefers φ -ing to ψ -ing conditional on p , and if she *actually* prefers ψ -ing to φ -ing, then S 's belief that p does not cohere well enough with the rest of her cognitive system for her belief to count as an item of knowledge.¹⁰³

We may stipulate that, while Hannah prefers going straight home to stopping at the bank, conditional on the bank's being open, she *actually* prefers stopping at the bank to going straight home. And in this case, CR tells us that Hannah's belief that the bank will be open does not cohere well enough with the rest of her cognitive system for her belief that the bank will be open to count as an item of knowledge.

Weatherson's CR gets right to the heart of the sensibilities that lead many people to adopt pragmatism. But is CR true? Incoherence can plausibly be a problem for knowledge. On many theories of defeat, for example, if S 's belief that p does not cohere well enough with some mental state M of S 's, then M gives S a defeater for her belief that p —in which case S is not justified in believing that p .¹⁰⁴ But in Hannah's case, purists will surely ask, *where exactly is the incoherence?* Reflecting on this question, it seems that purists can either deny that Hannah has incoherent mental states or reject the stipulation that Hannah prefers going straight home to stopping at the bank, conditional on the bank's being open.

¹⁰³ Though note two things. First, by using CR to argue that, in High Stakes, Hannah does not know that the bank will be open, we more or less hijack the principle. Weatherson uses CR to set the stage for an argument against Fantl and McGrath's explanation for why pragmatism is true. It is perfectly consistent with Weatherson's *intentions* for CR that CR gives purists no reason at all to think that Hannah lacks knowledge in High Stakes. Second, Weatherson presents CR within the context of his theory of belief, downstream from which CR is very plausible. Since Weatherson's theory of belief entails doxastic pragmatism, and since we are holding doxastic pragmatism to the side, we must consider CR apart from many of the considerations that Weatherson could produce in favor of CR. (See Weatherson (forthcoming), p. 27.)

¹⁰⁴ We can add 'knowledge-level' to this claim to make what we say here consistent with some comments that Weatherson makes about defeaters and justification at *ibid.*, p. 28.

What does Hannah's preferring to go straight home, conditional on the bank's being open, amount to? Conditionalizing on the proposition that the bank will be open amounts to assigning probability 1 to the proposition that the bank will be open, and this is what Hannah would do if she were to perform an expected utility calculation on Table 1, below.

Table 1

The bank will be open	
t bank	<i>Wait in line</i>
Go straight home	<i>Avoid the lines</i>

So consider the following subjunctive conditional.

- (3) If Hannah were to employ Table 1 in expected utility calculations for going straight home and stopping at the bank, then the result of her expected utility calculations would be that going straight home has higher expected utility than stopping at the bank.

And let 'STIP' name the stipulation that Hannah prefers going straight home to stopping at the bank, conditional on the bank's being open. Either STIP commits us to something *stronger* than (3), or it does not. If STIP does not commit us to anything stronger than (3), then it is unclear why purists should think that, if Hannah prefers going straight home to

stopping at the bank, conditional on the bank's being open, but *actually* prefers stopping at the bank to going straight home, then Hannah has incoherent mental states. Suppose that Hannah believes that the bank will be open without being maximally confident that the bank will be open; and suppose that, because Hannah is not maximally confident that the bank will be open, Hannah performs an expected utility calculation on Table 2, below, instead of Table 1.

Table 2

	The bank will be open	The bank will not be open
Stop at bank	<i>Wait in line</i>	<i>Avoid huge fine</i>
Go straight home	<i>Avoid the lines</i>	<i>Incur huge fine</i>

Additionally, suppose that Hannah assigns the proposition that the bank will be open the probability that matches her evidence that the bank will be open (*i.e.*, some probability less than 1); and suppose that, because Hannah sets up her expected utility calculation this way, Hannah concludes that stopping at the bank has higher expected utility than going straight home; and finally, suppose that, because Hannah has arrived at the conclusion that stopping at the bank has higher expected utility than going straight home, Hannah *actually* prefers stopping at the bank to going straight home.

Given these suppositions, Hannah believes that the bank will be open and she actually prefers stopping at the bank to going straight home. So suppose, finally, that (3) is true. In

this case, while Hannah believes that the bank will be open, and while she actually prefers stopping at the bank to going straight home, Hannah is such that, *if* she were to employ Table 1 in her expected utility calculations for going straight home and stopping at the bank, then her expected utility calculation would say that going straight home has higher expected utility than stopping at the bank. So now, given all this, must Hannah's cognitive system be incoherent? It is hard to see why purists cannot say, 'no.' If it is possible for Hannah to believe that the bank will be open without being maximally confident that the bank will be open, then there seems to be no reason why Hannah could not have a perfectly coherent cognitive system while believing that the bank will be open *and* actually preferring to stop at the bank *and* being such that (3) is true. But it *does* seem possible for Hannah to believe that the bank will be open without being maximally confident that the bank will be open,¹⁰⁵ so it seems that, if STIP commits us to nothing stronger than (3), then purists can reject Weatherson's CR as implausible. It seems that, if STIP commits us to nothing stronger than (3), then purists can insist that *Hannah* gives us a counterexample to CR.

Suppose that STIP *does* commit us to something stronger than (3), then. For example, suppose that it commits us to something like (4) or (5), below.

- (4) If Hannah were to perform expected utility calculations for going straight home and stopping at the bank *on the decision table that she thinks is best suited for determining the rational course of action*, then going straight home would come out as having higher expected utility for Hannah than stopping at the bank.

¹⁰⁵ After all, it seems that Hannah can believe that the bank will be open while she is less confident that the bank will be open than she is (for example) that $2 + 2 = 4$.

- (5) Hannah believes that, if the bank will be open, then she should prefer going straight home to stopping at the bank.

If STIP commits us to anything like (4) or (5), then it *does* seem that Hannah's cognitive system must be incoherent (assuming that she also believes that the bank will be open and actually prefers stopping at the bank to going straight home). But in this case, why should purists accept STIP? High Stakes loses relevance to the debate between pragmatists and purists if Hannah has some cognitive vice that will allow purists to *agree* with pragmatists that Hannah does not know that the bank will be open. But (4) and (5) both make Hannah look cognitively vicious in precisely this way. Purists may agree that, if (4) or (5) is true, then the relevant parts of Hannah's cognitive system lack sufficient coherence for Hannah to know that the bank will be open. But whether they agree or not, they will insist that (4) and (5) *need not* be true of someone in Hannah's situation, and the interesting question (they will point out) is whether someone in Hannah's situation for whom (4) and (5) *are not* true can know that the bank will be open. So if STIP *does* commit us to something stronger than (3)—like (4) or (5)—then it seems doubtful that purists should accept STIP. And in this case, purists can just deny that Weatherson's CR has any application to Hannah in High Stakes. So either STIP commits us to something stronger than (3), or it does not. If it does, purists can deny that CR applies to Hannah in High Stakes. If it does not, purists can deny that CR is true. Either way, CR does not prevent purists from endorsing K.

Fifth, and finally, we have Schroeder's claim that Hannah has reasons to withhold with respect to the proposition that the bank will be open in High Stakes that she lacks in Low

Stakes. Schroeder distinguishes between reasons to believe and reasons to withhold, and then argues that reasons to withhold (whatever they are) cannot be evidence.

[T]he evidence is exhausted by evidence which supports p and evidence which supports $\neg p$. But the evidence which supports p is reason to believe p , and the evidence which supports $\neg p$ is reason to believe $\neg p$. Consequently the reasons to withhold must come from somewhere else. So they cannot be evidence.¹⁰⁶

What *are* the reasons to withhold? Schroeder does not offer a complete list, but he does say that costs of error can be an important reason to withhold.

[A] natural place to look for reasons to withhold is in the costs of error. When you form a belief, you take a risk of getting things wrong that you don't take by withholding. In contrast, when you withhold, you guarantee that you miss out on getting things right. So plausibly, one important source of reasons to withhold will come from the preponderance of the cost of having a false belief over the cost of missing out on having a true belief—or, as I will put it, the preponderance of the cost of *type-1 error* over *type-2 error*.¹⁰⁷

This paragraph seems plausible enough. But so far, it lends no support to pragmatism, for the cost of a type-1 error might just be the possession of a false belief, and the cost of a type-2 error might just be the absence of a true belief. If these are the *only* costs of type-1 and type-2 errors, however, then *practical* circumstances and concerns seem beside the point.

¹⁰⁶ Ibid., p. 12.

¹⁰⁷ Ibid., p. 13.

For example, if these are the only costs, then the costs of type-1 and type-2 errors will be the same in High Stakes as they are in Low Stakes. According to Schroeder, however, the costs of type-1 and type-2 errors will often go well *beyond* false belief and the absence of true belief, respectively.

[T]he most general sort of cost of type-1 error is simply mistakes that we make, when we act on a belief that turns out to be false. ... Correlatively, the most general sort of cost of type-2 error derives from the fact that sometimes we simply have to act, and ignorance doesn't help us. These two sorts of costs—of type-1 and type-2 errors—are clearly practical in nature. They derive from the connection between belief and action. ... Gratifyingly, in High Stakes the costs of type-1 error are extremely high, and the costs of type-2 error are very low, which on this picture supports the view that there are especially strong reasons to withhold in that case.¹⁰⁸

Purists should have no trouble agreeing with Schroeder that, if Hannah has decisive epistemic reason to withhold with respect to the proposition that the bank will be open, then she cannot know that the bank will be open. But why does Schroeder think that Hannah epistemically ought to withhold with respect to that proposition? Because, as indicated in the quotation above, he thinks that the high-costs of a type-1 error with respect to the proposition that the bank will be open give Hannah sufficient epistemic reason to withhold.

Now, purists may deny that the costs of a type-1 error with respect to the proposition that the bank will be open give Hannah any *epistemic* reason to withhold.¹⁰⁹ A more promising

¹⁰⁸ Ibid., pp. 13-14.

¹⁰⁹ Thanks to Schroeder for conversation on this point.

response has purists commenting on the nature of belief, however. In the quotation above, Schroeder says that the costs of a type-1 error are very high for Hannah, and he says that these costs “derive from the connection between belief and action.” What exactly *is* the connection between belief and action? Schroeder does not say, but his argument seems to presuppose the following belief-action principle.

(BA) *S* believes that *p* only if she is willing to act as if *p* (where *S* acts as if *p* just in case she acts in the manner that would be rationally optimal, conditional on *p*).

As we saw above, going straight home would be rationally optimal for Hannah conditional on the bank’s being open. So, according to BA, Hannah believes that the bank will be open only if she is willing to go straight home. Suppose that BA is *false*, then. In this case, Hannah can be unwilling to go straight home and yet still believe that the bank will be open. Yet if Hannah can believe that the bank will be open while she is unwilling to go straight home, then why should we think that, in High Stakes, the costs of type-1 error with respect to the proposition that the bank will be open are extremely high for Hannah? It seems that, if BA is false, Hannah can believe that the bank will be open and simultaneously take all of the precautions necessary for insuring that she does not incur the large fine. Yet if Hannah can believe that the bank will be open and simultaneously take all of the precautions necessary for insuring that she does not incur the fine, then the cost of a type-1 error with respect to the proposition that the bank will be open will be no higher than forming a false belief. And this will be the case *in both Low Stakes and High Stakes*. So, if BA is false, the costs of type-1

error with respect to the proposition that the bank will be open seem to be the *same* in Low Stakes and High Stakes. It follows that, unless BA is *true*, the costs of a type-1 error with respect to the proposition that the bank will be open should be no higher for Hannah in High Stakes than they are in Low Stakes. Thus, Schroeder's argument that the costs of type-1 error increase as the stakes go up seems to presuppose BA. And since his argument for pragmatism depends on the claim that the costs of type-1 error increase as the stakes go up, his argument for pragmatism seems equally dependent on the truth of BA.

Fortunately for purists, BA looks false. First, BA seems susceptible to counterexamples. Consider the following propositions: that two plus two equals four, that the world is more than five minutes old, and that the sun will rise tomorrow. These propositions are not equally probable: the probability that two plus two equals four exceeds the probability that the world is more than five minutes old, which exceeds the probability that the sun will rise tomorrow. Nevertheless, I *believe* all three of these propositions, and hence I believe that the sun will rise tomorrow. My belief that the sun will rise tomorrow is a counterexample to BA. Just consider the question, *which is more probable, that the world is more than five minutes old, or that the sun will rise tomorrow?* I think it is more probable that the world is more than five minutes old than it is that the sun will rise tomorrow, and so I am not willing to answer this question by asserting that the sun will rise tomorrow. Yet conditional on the proposition that the sun will rise tomorrow, asserting that the sun will rise tomorrow would be the rationally optimal response to this question. (After all, conditional on the proposition that the sun will rise tomorrow, the probability that the sun will rise tomorrow is 1, while, conditional on the proposition that the sun will rise tomorrow, the probability that the world is more than five

minutes old is still less than 1, since it is still less than the probability that two plus two equals four.) So BA says (falsely) that I do not believe that the sun will rise tomorrow.

Second, in a paper with Jacob Ross, *Schroeder* gives us reason to reject BA. Ross and Schroeder wield the following principle against various theories of belief.

Stability: A fully rational agent does not change her beliefs purely in virtue of an evidentially irrelevant change in her credences or preferences.¹¹⁰

Since Stability is true, theories of belief that conflict with it are false, argue Ross and Schroeder.¹¹¹ Fortunately for purists, Stability is not consistent with BA.

By stipulative definition, one's credence in q is evidentially irrelevant to p just in case neither q nor one's credence in q provides any evidence for or against p .¹¹² So let d be the proposition that there will be disastrous consequences for Hannah if she does not deposit her check before Saturday afternoon. Because neither d nor Hannah's credence in d provides any evidence for or against the proposition that the bank will be open Saturday morning, d is evidentially irrelevant to the proposition that the bank will be open Saturday morning. So now imagine Hannah in *Low Stakes* and notice what happens if we stipulate that Hannah is fully rational, and then vary Hannah's credence in d . Manipulating just this variable, we change what rationality requires of Hannah. Once Hannah's credence in d gets sufficiently high, Hannah goes from maximizing expected utility by going straight home to maximizing expected utility by stopping at the bank. Since Hannah is fully rational, we thereby change what Hannah is willing to do; once her credence in d gets sufficiently high, she loses her

¹¹⁰ Ibid., p. 19.

¹¹¹ Ibid., pp. 19-22.

¹¹² Ibid., p. 19.

willingness to go straight home. According to BA, however, Hannah believes that the bank will be open only if she is willing to go straight home. So, if BA is true, by varying just Hannah's credence in *d*, we change what Hannah *believes*. But again, *d* is evidentially irrelevant, and we are supposing that Hannah is fully rational. So, if BA is true, a fully rational agent *may* change her beliefs purely in virtue of an evidentially irrelevant change in her credences, contrary to Stability. We therefore have a second reason to think that BA is false. But without BA, Schroeder's argument for pragmatism fails. Because BA seems to have counterexamples and conflict with Stability, purists can resist Schroeder's argument for pragmatism. And since Schroeder's argument is the final item in our list of obstacles to endorsing K, it seems that purists might overcome the most substantial obstacles to endorsing K.¹¹³

2.3. Obvious Considerations in Favor of K

Of course, in addition to resources for resisting the arguments against K, purists have positive reasons for endorsing K. The most obvious of these is that Hannah's belief that the bank will be open has a lot going for it, epistemically. First, since High Stakes differs from Low Stakes only insofar as the stipulated differences in Hannah's practical interests require that they differ, since, in Low Stakes, Hannah's belief that the bank will be open is true, and since the stipulated differences in Hannah's practical interests do not require that, in High Stakes, Hannah's belief that the bank will be open is *false*, Hannah's belief that the bank will be open is *true*. Second, by hypothesis, Hannah has very good evidence that the bank will be

¹¹³ Recall that we are restricting our attention to the debate between *non-doxastic* pragmatists and purists. Pragmatic theories of belief like that in Weatherson (forthcoming) also raise substantial obstacles for purists who want to endorse K. I deal with these obstacles to endorsing K in "Pragmatism and the Nature of Belief."

open and she is basing her belief that the bank will be open on this evidence. So Hannah's belief is well-supported by the evidence. Third, we have no reason to think that Hannah's belief that the bank will be open fails to cohere with *her other beliefs*, so purists can just stipulate that her beliefs are coherent, both deductively and probabilistically. Fourth, a process is reliable insofar as it produces true rather than false beliefs. The process whereby S considers the evidence that p , notes that she has fantastic evidence that p and no evidence that $\neg p$, and thereby comes to believe that p , will produce a high preponderance of true belief. But this is the process whereby Hannah formed her belief that the bank will be open. So Hannah's belief is highly reliably produced. Fifth, a belief is apt just in case it is accurate because adroit, and a belief is accurate because adroit just in case it is true and, in being true, it manifests cognitive virtues of the person who formed it.¹¹⁴ Hannah formed her belief by weighing the evidence for and against the proposition that the bank will be open. Hannah was not acting out of character when she formed her belief this way (purists do not beg any questions by stipulating this, at least), so her belief looks apt. Sixth, where p is a contingent proposition, S 's true belief that p is safe just in case most of the nearby worlds where S believes that p are also worlds where p is true. There are no nearby worlds where the bank has changed its hours without warning, where there is an impending white-out blizzard that has not been forecast, where all of the bank employees will fail to show up for no reason, *etc.* So most of the nearby worlds where Hannah believes that the bank will be open *are* worlds where the bank will be open. Thus, Hannah's belief is safe. Seventh, where p is a true contingent proposition, S 's belief that p is *sensitive* just in case S does not believe that p in any

¹¹⁴ See, for example, Sosa (2009), p. 134.

of the nearest worlds where p is false. The nearest worlds where the bank will not be open on Saturday are worlds where the bank has been closed on some previous Saturdays, or worlds where there is a blizzard forecast for Saturday, or worlds where the bank has announced that it has changed its hours, or worlds where the bank's employees are on strike, *etc.* In those worlds, Hannah has adequate reason to doubt that the bank will open on Saturday, and, consequently, she does not believe that the bank will open on Saturday. (Again, purists do not beg any questions by stipulating this.) So Hannah's belief is sensitive. Eighth, and finally, Hannah's belief is based on the long unbroken streak of Saturday openings that Hannah has witnessed over the years, in conjunction with very good evidence that this Saturday will be no different than any previous Saturday. Since there is no deviant causal chain involved in the reasoning that leads Hannah to believe that the bank will be open, and since there is nothing analogous to a fake barn in the vicinity, Hannah's belief manifests none of the luck that is characteristic of Gettiered beliefs. Insofar as there is a tight connection between epistemic luck and being Gettiered, Hannah's belief is not Gettiered.¹¹⁵

So, in favor of K, purists might point out that Hannah's belief that the bank will be open is true, based on good evidence, is deductively coherent with the rest of her beliefs, is rational (where rationality equals probabilistic coherence), is highly reliably produced, apt, safe, and sensitive, and is not Gettiered. Hannah's belief that the bank will be open has a lot going for it epistemically, and purists can bring these considerations to bear positively on K. And since purists would not beg the question against pragmatists by stipulating that Hannah is *aware* that her belief has all of these things going for it (in the sense of 'aware' familiar

¹¹⁵ See Sosa (2010) and Pritchard (2005) for attempts to say exactly what the luck characteristic of Gettiered beliefs amounts to, and Lackey (2008) for a response to Pritchard.

from internalist accounts of epistemic justification), we may assume that Hannah is aware of all the epistemic goods that her belief manifests. But given all this, purists might now challenge pragmatists to produce some property φ such that, clearly, K is true only if Hannah's belief that the bank will be open has φ and, clearly, Hannah's belief that the bank will be open *lacks* φ . And since (as we have just seen), purists have means of resisting the most formidable pragmatist objections to K , purists might insist that no such property is forthcoming.¹¹⁶ The upshot is that purists have considerable resources for defending K .

3. A Paradox for Justification, Knowledge and Practical Rationality

Now consider the following propositions, and assume that they describe Hannah in High Stakes, as she decides between going straight home and stopping at the bank.

- ($\neg J$) Hannah is not epistemically justified in believing that the bank will be open on Saturday.
- ($J \wedge \neg K$) Hannah *is* epistemically justified in believing that the bank will be open on Saturday, but she does not *know* that the bank will be open on Saturday.
- (K) Hannah *does* know that the bank will be open on Saturday.

Given the details of High Stakes, these propositions run epistemic justification up against knowledge in such a way that it is difficult to say how epistemic justification, knowledge, and practical rationality all fit together. At least one of $\neg J$ through K is true; the form of the

¹¹⁶ Remember, we are taking it for granted that doxastic pragmatism is false.

proposition that we get when we disjoin them guarantees this.¹¹⁷ And given that knowledge entails epistemically justified belief, *at most* one of $\neg J$ through K is true; each of $\neg J$ through K entails the negation of the other two.¹¹⁸ But as I will now argue, it is unclear what a successful argument for any particular option in $\neg J$, $(J \wedge \neg K)$ and K might look like.

Let's stipulate that S has *knowledge-level* justification for her belief that p just in case S cannot come to know that p by simply gaining more justification for believing that p (either because she *already* knows that p or because she lacks some property that is necessary for knowing that p and she cannot acquire this property by simply gaining more justification for believing that p).¹¹⁹ With this definition in hand, we may consider Fantl and McGrath's "equivalence thesis," below.

¹¹⁷ The disjunction of these propositions takes the following form: $(\neg p \vee (p \wedge \neg q) \vee q)$. As we can work out, a disjunction of this form cannot fail to have at least one true disjunct.

¹¹⁸ Of course, there are externalist conceptions of knowledge and deontological conceptions of justification that, in combination, entail that it is possible to know that p without being justified in believing that p . Readers who hold externalist conceptions of knowledge in combination with deontological conceptions of justification are invited to swap out every instance of 'justification' in what follows for 'warrant,' or whatever other term they prefer.

¹¹⁹ Note carefully: knowledge-level justification (so defined) is consistent with whatever theory of justification one might hold, so long as one's theory of justification allows for justification to come in degrees. On a *reliabilist* theory of justification (for example), S has knowledge-level justification for her belief that p just in case she cannot come to know that p by simply holding her belief that p in a more reliable way (either because she *already* knows that p or because she lacks some property that is necessary for knowing that p and she cannot acquire this property by simply holding her belief that p in a more reliable way). On an *evidentialist* theory of justification, the claim would be that S has knowledge-level justification for her belief that p just in case she cannot come to know that p by simply gaining more evidence that p (either because she already knows that p or because she lacks some property that is necessary for knowing that p and she cannot acquire this property by simply gaining more evidence that p). And so on.

I believe that Fantl and McGrath's definition of knowledge-level justification is equivalent to the definition of knowledge-level justification that we are presently discussing, though questions about Fantl and McGrath's use of the phrase 'stand in the way' make me somewhat uncertain. (See Fantl and McGrath (2009), pp. 97-8.)

(EQ) S is justified in believing that p just in case S is knowledge-level justified in believing that p .¹²⁰

Either EQ is true, or it is false. And if EQ is false, then justified beliefs are easier to acquire than knowledge-level justified beliefs.¹²¹

How plausible are $\neg J$ and $(J \wedge \neg K)$, if EQ is false? First, if EQ is false, $\neg J$ looks implausible, since Hannah's belief that the bank will be open possesses justification-relevant properties in abundance. Hannah is basing her belief that the bank will be open on good evidence, her belief that the bank will be open is the product of a highly reliable process, and so on. And as we just saw, we may suppose that Hannah is *aware* that her belief has all of these justification-relevant properties. But if Hannah's belief that the bank will be open has so much going for it by way of justification-relevant properties, then why would Hannah fail to be epistemically justified in believing that the bank will be open? If EQ is false (as we are supposing), then Hannah might have an epistemically justified belief that falls short of *knowledge-level* justification. So our reason for thinking that Hannah lacks an epistemically justified belief cannot be that Hannah must be more justified in order to *know*. Of course, Hannah's practical concerns might create problems for the claim that Hannah is *pragmatically* justified in believing that the bank will be open. But absent a reason to think that Hannah must be pragmatically justified in order to be *epistemically* justified, it would still be unclear why Hannah would fail to be epistemically justified in believing that the bank will be open. At the very least, given that EQ is false, $\neg J$ seems considerably less plausible than $(J \wedge \neg K)$,

¹²⁰ Ibid., p. 98.

¹²¹ Or warranted beliefs are easier to acquire than knowledge-level warranted beliefs, if the reader prefers, in line with fn 35.

which entails the negation of $\neg J$. So, given that EQ is false, it seems reasonable to reject $\neg J$.¹²²

So what should we make of $(J \wedge \neg K)$, given that EQ is false? I do not see how a successful argument for this option might go. If Hannah does not know that the bank will be open, it is not for lack of any justification-*irrelevant* property. Purists, of course, think that Hannah *does* know that the bank will be open. And pragmatists think that Hannah could gain enough justification for knowing that the bank will be open by (for example) going inside and hearing from the bank's president that, rain or shine, the bank *will* be open on Saturday—that the president will open the bank *herself* if she has to. So by all accounts—purist and pragmatist alike—if Hannah does not know that the bank will be open, this is because she lacks some property that is *relevant* to her justification for believing that the bank will be open. But according to $(J \wedge \neg K)$, Hannah *is justified* in believing that the bank will be open. So $(J \wedge \neg K)$ is true only if Hannah is justified in believing that the bank will be open, but not *knowledge-level* justified in believing that the bank will be open. In order to establish $(J \wedge \neg K)$, then, we must argue that, while Hannah is justified in believing that the bank will be open, she is not knowledge-level justified in believing that the bank will be open. But it seems doubtful that we can produce an argument for this conclusion that purists cannot reasonably resist. In order to establish that Hannah lacks knowledge-level justification, we would have to establish (first) that Hannah does not know that the bank will be open and (second) that Hannah could *come* to know that the bank will be open by simply gaining more justification. This just follows from our stipulative definition of knowledge-level justification.

¹²² And so far as I can tell, given our distinction between justification and knowledge-level justification, and our current assumption that EQ is false, pragmatists *themselves* would all reject $\neg J$.

But how are we supposed to establish, in the first place, that Hannah does not know that the bank will be open? As we saw in §2.2, purists have resources for resisting all of our best *arguments* for this conclusion.¹²³ So it seems that, in order to establish that Hannah lacks knowledge that the bank will be open, we would be forced rely on the bald *intuition* that Hannah does not know that the bank will be open. And this is a problem, since it seems doubtful that the bald intuition that Hannah lacks knowledge carries enough force. First, experimental philosophy raises questions about the reliability of intuitions in general. Purists might insist that, because it is doubtful that our intuitions are reliable, our choice between pragmatism and purism must be made by weighing the theoretical virtues of pragmatism and purism against each other. Second, even if the worries raised by experimental philosophy can be overcome (even if experimental philosophy does not undermine the use of intuitions in philosophy *in general*), Jennifer Nagel and others have raised worries about the *specific* intuition that Hannah lacks knowledge in High Stakes. The psychological literature on closed-mindedness casts doubt on the reliability of this intuition, argues Nagel.¹²⁴ And third, even if the psychological literature on closed-mindedness does not undermine the intuition that Hannah lacks knowledge in High Stakes,¹²⁵ the considerations adduced in §2.2 and §2.3 on behalf of *K counterbalance* (to at least some extent) whatever intuition purists might have that Hannah lacks knowledge in High Stakes. Even if they think the intuition that Hannah lacks knowledge in High Stakes survives the challenge from experimental philosophy and Nagel's undermining psychological explanation, purists might reasonably find this intuition

¹²³ Again, we are holding doxastic pragmatism to the side. I address doxastic pragmatism in my paper "Pragmatism and the Nature of Belief."

¹²⁴ See Nagel (2008) and, for example, Kruglanski (2004).

¹²⁵ Perhaps for the reasons discussed in Sripada and Stanley (forthcoming).

overwhelmed by the considerations adduced in §2.2 and §2.3 in favor of K. But given all of these worries, an argument for $(J \wedge \neg K)$ that bottoms out in the intuition that Hannah lacks knowledge in High Stakes seems unlikely to convince. So, given that EQ is false, the prospects for establishing $(J \wedge \neg K)$ look dim.

Let's assume that EQ is *true*, then. In this case, justified belief is just *knowledge-level* justified belief, and $\neg J$ and $(J \wedge \neg K)$ are both about knowledge-level justification. Given that EQ is true, $(J \wedge \neg K)$ seems indefensible. If EQ is true, then Hannah is justified in believing that the bank will be open only if she cannot come to know that the bank will be open by simply gaining more justification for believing that the bank will be open (either because she *already* knows that the bank will be open or because she lacks some property that is necessary for knowing that the bank will be open and she could not acquire this property by simply gaining more justification for believing that the bank will be open). This just follows from our stipulative definition of knowledge-level justification. Thus, since $(J \wedge \neg K)$ tells us that Hannah is justified in believing that the bank will be open, $(J \wedge \neg K)$ is true only if Hannah could not come to know that the bank will be open by simply gaining more justification for believing that the bank will be open, either because she already knows that the bank will be open or because she lacks some property that is necessary for knowing that the bank will be open and she cannot acquire this property by simply gaining more justification for believing that the bank will be open. But if Hannah already knows that the bank will be open, then K is true and $(J \wedge \neg K)$ is false. So, given that EQ is true, $(J \wedge \neg K)$ is true only if Hannah does not know that the bank will be open and Hannah could not come to know that the bank will be open by simply acquiring more justification for believing that the bank will be open. But

by all accounts—pragmatist and purist alike—either Hannah *does* know that the bank will be open or Hannah *could* come to know that the bank will be open by simply acquiring more justification for believing that the bank will be open. Again, purists think that Hannah *does* know that the bank will be open, and pragmatists think that Hannah could gain enough justification for knowing that the bank will be open by (for example) going inside and hearing from the bank’s president that the bank will be open on Saturday. Given that EQ is true, *nobody* has any interest in accepting $(J \wedge \neg K)$.

What should we make of $\neg J$, then? Here, we seem to run into the same problem as we encountered while considering $(J \wedge \neg K)$ under the assumption that EQ is *false*. Given that EQ is true, $\neg J$ is true just in case Hannah lacks *knowledge-level* justification for believing that the bank will be open. But how are we supposed to argue that Hannah lacks knowledge-level justification for believing that the bank will be open? As we have already seen, we cannot establish that Hannah lacks knowledge-level justification for believing that the bank will be open except by establishing (first) that Hannah does not know that the bank will be open and (second) that Hannah could *come* to know that the bank will be open by simply gaining more justification. But as we have already seen, we seem to lack compelling *arguments* that Hannah does not know that the bank will be open, it seems doubtful that we can rely on the bald *intuition* that Hannah does not know that the bank will be open, and it is not clear what other options are available to us. So given that EQ is true, it is not clear how we might establish $\neg J$.

The upshot is that either EQ is true or it is not, and, either way, it is unclear how a successful argument for $\neg J$ or $(J \wedge \neg K)$ might go. But at least one of $\neg J$, $(J \wedge \neg K)$ and K is

true, and K faces problems of its own. In §2.1, we considered many arguments against K , and, as we saw in §2.2, defending K from these arguments would be no easy task. So, while $\neg J$, $(J \wedge \neg K)$, and K cannot all be false, it is unclear what a successful argument for any particular option in $\neg J$, $(J \wedge \neg K)$, and K might look like. In this way, $\neg J$, $(J \wedge \neg K)$, and K confront us with a paradox.¹²⁶

4. Responding to the Paradox

So which should we prefer among $\neg J$, $(J \wedge \neg K)$, and K ? Whether EQ is true or not, knowledge entails justified belief, so $\neg J$ and $(J \wedge \neg K)$ both entail $\neg K$.¹²⁷ Which is more plausible, then: $\neg K$ or K ? Here, I think our leanings should be towards K . My reason for thinking this is simple. In Low Stakes, Hannah knows that the bank will be open. There are many epistemically relevant similarities between Hannah's belief that the bank will be open in Low Stakes and her belief that the bank will be open in High Stakes; this much is clear and uncontroversial. It is both unclear and highly controversial that there are any epistemically relevant *dissimilarities* between Hannah's belief that the bank will be open in Low Stakes and her belief that the bank will be open in High Stakes. So, our reasons for thinking that Hannah knows that the bank will be open in High Stakes outweigh our reasons

¹²⁶ Compare $\neg J$, $(J \wedge \neg K)$ and K to the skeptical paradox. The skeptical paradox is sometimes said to arise because its members are all attractive even while they are jointly inconsistent, and clearly so. I believe that this is the wrong way to characterize the skeptical paradox. For my own part, I find the proposition that I do not know that I am not a BIV wildly implausible, and yet I am still willing to grant that the skeptical paradox is, in fact, a paradox. The paradox, I think, arises from the difficulty of producing an *argument* against any option in the skeptical paradox that is not dialectically deficient. I want to say that $\neg J$, $(J \wedge \neg K)$ and K confront us with a paradox in much the same way, except now the challenge is to produce an argument *for* one of the options that does not seem dialectically deficient.

¹²⁷ Again, readers who hold a thoroughly externalist conception of knowledge in combination with a strictly deontological conception of justification—so that knowledge does not entail justified belief—are invited to swap out every instance of 'justification' in what follows for 'warrant,' or whatever other term they prefer.

for thinking that Hannah *does not* know that the bank will be open in High Stakes. In this case, K is more plausible than $\neg K$. And since $\neg J$ and $(J \wedge \neg K)$ both entail $\neg K$, K comes out more plausible than $\neg J$ and $(J \wedge \neg K)$. This makes K the most plausible option in $\neg J$, $(J \wedge \neg K)$ and K , so we should favor K .

5. Pragmatism Unmotivated?

But if K is the most plausible option in $\neg J$, $(J \wedge \neg K)$, and K , then Low Stakes and High Stakes cannot be used to motivate pragmatism, since K sits perfectly well with purism. And since there is no relevant difference between Low and High Stakes and the other low- and high-stakes cases in the literature, it looks like none of the cases familiar from the literature can be used to motivate pragmatism. What *does* motivate pragmatism, then?

One suggestion is that that, even if the cases familiar from the literature cannot motivate pragmatism, pragmatists can employ knowledge-action principles like RKP and KJ to motivate their view.¹²⁸ Is this suggestion correct? Certainly, pragmatists can employ knowledge-action principles like RKP and KJ to *exposit* and *articulate* their view. But the considerations adduced in §§2-4 cast doubt on the claim that pragmatists can employ knowledge-action principles like RKP and KJ to *motivate* their view.

We saw that RKP was (technically speaking) false, and that, once it was appropriately modified, it said that, where S 's choice is p -dependent, it is *epistemically* appropriate for S to treat the proposition that p as a reason for acting iff S knows that p . We called this resulting principle 'RKP'.¹²⁸ So consider (6), which we need in order to derive pragmatism from RKP'.

¹²⁸ The suggestion that KJ can motivate pragmatism is the centerpiece of Fantl and McGrath (2009).

- (6) Where S 's choice is p -dependent, whether it is epistemically appropriate for S to treat the proposition that p as a reason for acting varies with S 's practical interests.

Either we are accepting RKP', or we are not. If we are not, then we cannot employ RKP' in an argument for pragmatism. But if we are, we seem to prevent ourselves from accepting (6). We just saw that the most plausible option in the paradox presented by High Stakes says that Hannah knows that the bank will be open. But if Hannah knows that the bank will be open, and if RKP' is true, then, in High Stakes, it is epistemically appropriate for Hannah to treat the proposition that the bank will be open as a reason for acting. And by hypothesis, we are accepting RKP'. Yet whatever epistemic propriety might plausibly amount to, it must be epistemically appropriate for Hannah to treat the proposition that the bank will be open as a reason for acting in *Low Stakes*. So accepting RKP' puts pressure on us to say that, in both Low *and* High Stakes, it is epistemically appropriate for Hannah to treat the proposition that the bank will be open as a reason for acting. Now (6) entails that there will be pairs of cases that are identical except for some difference in practical interests, and where, as a result of this difference in practical interests, it is epistemically appropriate for S to treat the proposition that p as a reason for acting in one case but not the other. But Low and High Stakes should satisfy this description if any pair of cases satisfies it, and our acceptance of RKP' puts pressure on us to say that Low and High Stakes *do not* satisfy this description. So, if we accept RKP', we have reason to doubt that (6) is true. But without (6), we cannot employ RKP' in an argument for pragmatism. So either we accept RKP' or we do not, and,

either way, we cannot employ RKP' in a successful argument for pragmatism. So RKP' cannot motivate pragmatism. (And of course, RKP also cannot motivate pragmatism, since RKP appears to be false.)

But then, what *is* the motivation for pragmatism? We saw in §2.2 that KJ looks false. So it seems doubtful that KJ can motivate pragmatism. So what other options for motivating pragmatism are there? We considered two positions defended by Weatherson: (3), which says that you can leave a state of affairs off of a decision table if you know that it does not obtain, and CR, which says that, if you believe that p , if you prefer φ -ing to ψ -ing conditional on p , and if you *actually* prefer ψ -ing to φ -ing, then your belief that p does not cohere well enough with the rest of your cognitive system to count as an item of knowledge. But we found that (3) was implausible, and CR was never intended to motivate pragmatism in the first place. Weatherson uses CR to set the stage for an argument against Fantl and McGrath's *explanation* for why pragmatism is true, not to argue that pragmatism *is* true. (And of course, the problems with CR discussed in §2.2 would plague anyone who *did* want to use CR to motivate pragmatism.) Finally, Schroeder's argument for pragmatism is not compelling, since, as we saw, it seems to presuppose a false theory of belief. So again, what motivates pragmatism?

By 'pragmatism,' of course, we mean *non*-doxastic pragmatism. We have said nothing about the motivation for accepting *doxastic* pragmatism. The motivation for accepting doxastic pragmatism deserves a paper of its own. Given that K is the most plausible of the

options in $\neg J$, $(J \wedge \neg K)$ and K , however, it has proven doubtful that *non*-doxastic pragmatism is adequately motivated.¹²⁹

¹²⁹ Thanks to Lisa Miracchi and Ernest Sosa for helpful comments on an earlier draft of this paper.

Chapter 4: Problems for Pragmatism

Abstract: The pragmatic encroachment debate pits a view that I call ‘pragmatism’ against a view that I call ‘purism.’ Pragmatism is the view familiar from Jeremy Fantl, Matthew McGrath, John Hawthorne, Jason Stanley, and others, according to which knowledge depends at least in part on our practical interests. Purism is just the denial of pragmatism. According to purism, knowledge does not depend in any interesting way on our practical interests. A review of the literature reveals two kinds of arguments for pragmatism: principle-based arguments and case-based arguments. Principle-based arguments derive pragmatism from plausible principles that connect knowledge to practical interests. Case-based arguments rely on intuitions about cases that differ only with respect to practical interests. I argue that neither kind of argument succeeds, and that it is therefore unclear what reason there is to accept pragmatism.

1 Introduction

You are about to leave your house for work. Looking at the sky, it seems certain that it will not rain. Your hands are full, but your friend suggests that you grab your umbrella anyway. This is Seattle, after all, and you are wearing your best suit. What should you do? The answer depends, not just on the chances that it will rain, but on your *practical interests*—on the consequences that might follow from taking your umbrella or leaving it, and the values of these possible consequences.

Purists disagree with pragmatists about the relationship between knowledge and practical interests. According to pragmatism, knowledge depends at least in part on our practical interests, in the sense that a mere difference in practical interests can entail a difference in knowledge.¹³⁰ If pragmatism is true, there will be pairs of cases that differ only insofar as some difference in practical interests requires that they differ, and where, given *just* this difference in practical interests, a difference in knowledge follows. Purism is just the denial of pragmatism.¹³¹ If purism is true, knowledge *does not* depend in any interesting way on our practical interests; the only difference in knowledge that might follow from a *mere* difference in practical interests is a difference in knowledge *about those very practical interests*.¹³²

Why accept pragmatism or, alternatively, accept purism? Two reasons for accepting pragmatism are now familiar. First, pragmatism explains our intuitions about pairs of cases like the following.

Low Stakes: Hannah is driving home on Friday afternoon. She plans to stop at the bank to deposit her paycheck, but, as she approaches the bank, she notices that the lines inside are very long. She knows that it does not matter much when she deposits the check, so she considers driving straight home and depositing the check Saturday

¹³⁰ Paradigm statements of pragmatism include Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005), Hawthorne and Stanley (2008), Ganson (2008), Fantl and McGrath (2009), Weatherson (2012), Ross and Schroeder (forthcoming), and Schroeder (forthcoming).

¹³¹ Paradigm defenses of purism include Neta (2007), Brown (2008), Nagel (2008), DeRose (2009), Fumerton (2010), Reed (2010), Brown (2012), Reed (2012), Neta (2012), and Cohen (2012).

¹³² The labels ‘purism’ and ‘pragmatism’ come from Fantl and McGrath (2009). Though ‘pragmatism’ is perhaps misleading, it seems clearly preferable to the more popular labels ‘subject-sensitive invariantism’ and ‘interest-relative invariantism,’ since these labels wed the epistemological thesis that knowledge depends at least in part on our practical interests to the linguistic thesis that ‘knows’ is not a context-sensitive word. As Fantl and McGrath note, we should keep these theses apart (at least in our terminology), since the thesis that knowledge depends at least in part on our practical interests is perfectly consistent with the thesis ‘knows’ *is* a context-sensitive word.

morning, when she is out running errands. Sarah says, “Maybe the bank will not be open tomorrow. Lots of banks are closed on Saturdays.” Hannah replies, “No, I know that it will be open tomorrow. I stop in at the bank every Saturday and it is open until noon.”

High Stakes: Hannah is approaching her bank on Friday afternoon, as in Low Stakes. Again she notices the long lines, and again she considers driving straight home and depositing the check Saturday morning, when she is out running errands. But in this case, Hannah knows that she will incur an enormous fine if she does not deposit her check before noon on Saturday. Sarah mentions this fine to Hannah and says, “Sometimes banks change their hours. Do you know that our bank is still open on Saturdays?” Hannah replies, “I guess you’re right. I don’t know that the bank will be open tomorrow.”¹³³

Intuitively, Hannah knows that the bank will be open on Saturday in Low Stakes, but does not know that the bank will be open on Saturday in High Stakes. This is exactly what we would expect if pragmatism were true, however, since there is an obvious difference between Low Stakes and High Stakes with respect to Hannah’s practical interests. So, our intuitions about cases like Low Stakes and High provide at least *prima facie* motivation for accepting pragmatism. The second familiar motivation for accepting pragmatism is that it fits nicely with our habit of using knowledge attributions to defend and explain our actions.¹³⁴ Suppose we are making a midnight run to White Castle, and I am speeding towards a red light. As I barrel toward the intersection, you eventually say, “What are you doing? The light is red! You

¹³³ See DeRose (1992), p. 913, and Stanley (2005), pp. 2-3.

¹³⁴ See, for example, Hawthorne and Stanley (2008).

need to slow down!” In defense of my driving, I respond, “No, it’s okay. I know that the light will turn green.” Here, I am taking it for granted that my *knowing* that the light will turn green would suffice for the rationality of driving fast towards the intersection. But whether it *is* rational for me to drive this way depends on practical interests (on the consequences that might follow from my driving this way and the values of these possible consequences). So, if my knowing that the light will turn green would suffice for the rationality of my continuing to speed towards the intersection, then whether I *do* know that the light will turn green *also* depends on my practical interests, just as pragmatism says that it does. So our habit of using knowledge attributions to explain and defend our actions provides *prima facie* motivation for accepting pragmatism.

The primary motivation for *resisting* pragmatism is that it makes knowledge depend on “truth-irrelevant” factors. I scare-quote ‘truth-irrelevant’ here because purists are neither very clear what truth-irrelevant factors are, nor very clear why it would be a problem if knowledge depends on truth-irrelevant factors. In any case, the general idea is this. Knowledge entails things that come in degrees. To know that *p*, the cognitive faculties responsible for your belief that *p* must be *sufficiently* reliable, you must have *sufficient* evidence that *p*, your belief that *p* must be *sufficiently* safe and *sufficiently* sensitive, *et cetera*. In short, to know that *p*, your epistemic position *vis-à-vis* *p* must be *sufficiently* strong. If pragmatism is true, however, what *counts* as sufficiently strong changes as your practical interests change. If pragmatism is true, you could go from knowing that *p* to ignorance with respect to *p*, not because of any change in the strength of your epistemic position with respect to *p*, but because that strength of epistemic position is no longer strong enough for you to know that

p . That is, you could go from knowing that p to ignorance with respect to p without any change in the amount or quality of your *evidence* that p , without any change in the *reliability* of the cognitive faculties responsible for your belief that p , without any change in the *safety* or *sensitivity* of your belief that p , without any change in your level of *confidence* that p , without any change in the *rationality* of this level of confidence that p , without the addition of anything analogous to a fake barn in your environment, without the addition of any false lemma that you have relied on in your reasoning with respect to p (that is, without the addition of anything analogous to the salient features of Gettier scenarios)—and so on, for all of the conditions traditionally proposed as necessary for knowledge. But all of these conditions establish some connection between the believer and the fact that she arrived a *true* belief. So, if pragmatism is true, you could go from knowledge to ignorance without any change in the traditional truth-relevant factors. *This* (so far as I can tell) is what purists mean when they say that pragmatism would make knowledge depend on truth-irrelevant factors.

In this paper, I side with purists against pragmatists. The pragmatic encroachment literature contains two kinds of argument for pragmatism: principle-based arguments and case-based arguments. I argue that neither variety of argument provides sufficient motivation for accepting pragmatism, and that pragmatism is therefore unmotivated. In §2, I present a counterexample to a toy principle that entails pragmatism, and, in §3, I show how this counterexample creates problems for all of the most plausible principle-based arguments in the literature. In §4, I turn my attention to case-based arguments, and argue that, without the aid of principle-based arguments, case-based arguments have no chance of motivating pragmatism. In §5, I consider a challenge that pragmatist might pose to purists, and argue

that purists can easily meet this challenge. I conclude that, with all of the relevant considerations in view, it is unclear why anyone should accept pragmatism.

2 A Counterexample to the Knowledge-Preference Principle

As I noted in §1, pragmatism entails that there will be pairs of cases that differ only insofar as some difference in practical interests requires that they differ, and where, given *just* this difference in practical interests, a difference in knowledge follows. Low and High Stakes should satisfy this description if any pair of cases satisfies it, so consider how the following Knowledge-Preference principle bears on Low and High Stakes.

(KP) S knows that p in c only if she can rationally do what is most preferable conditional on p in c .¹³⁵

Assuming that Hannah knows that the bank will be open in Low Stakes, KP gets exactly the result that pragmatists need for High Stakes. In High Stakes, Hannah is choosing between stopping at the bank and going straight home, and going straight home is clearly preferable conditional on the proposition that the bank will be open on Saturday. Conditional on that proposition, Hannah wastes an insignificant amount of time by stopping at the bank on Friday, and saves that much time by going straight home on Friday. There is no serious chance that she will incur the fine if she goes straight home. Since it is better to save time

¹³⁵ How should we think about the notion of preferability at issue in KP—*subjectively*, in terms of the values that S personally places on the consequences that might follow from an action and S 's personal level of confidence in the relevant propositions, or more *objectively*, in terms of the values that S *ought* to place on the consequences that might follow from an action and the *epistemic probabilities* of the relevant propositions? Nothing I say will depend on our answer, though I think KP is considerably more plausible on the objective interpretation. Note also that, throughout, I will assume that an action or belief is *rational* just in case it is *not irrational*.

than to waste it, Hannah does what is most preferable conditional on the proposition that the bank will be open only if she decides to go straight home. But Hannah cannot *rationaly* decide to go straight home. Given the amount and quality of her evidence that the bank will be open on Saturday, and the severe consequences that she knows will follow if she fails to deposit her check before noon on Saturday, she can see that going straight home would be far too risky. This means that Hannah cannot rationally do what is most preferable conditional on the proposition that the bank will be open. So, according to KP, Hannah does not know that the bank will be open in High Stakes. This is exactly the result that pragmatists need.

KP resembles many of the principles actually endorsed by pragmatists, but it has clear counterexamples. Consider the following case.

Survey: Professor Clarke is walking toward the library after teaching her intro chemistry class. As she crosses the quad, a social scientist approaches her and asks her to participate in a survey. Clarke is curious, so she asks the social scientist how the survey works. The survey contains 100 questions, explains the social scientist. Each question contains a pair of propositions, at least one of which is true. Clarke's task will be to select a true proposition from each pair of propositions. There is a nice incentive for taking the survey, the social scientist explains. Each time Clarke selects exactly one true proposition from a pair of propositions, she will get \$1. For each question where she either selects a *false* proposition, selects *both* propositions, or selects *neither* proposition, however, she will get nothing. Clarke is in a hurry, but she sees immediately that, by randomly guessing, she can complete the survey in a matter

of minutes and expect to earn approximately \$50. She then realizes that, if she works through the survey and quickly selects the most *probable* of each pair of propositions, she can complete the survey almost as quickly and expect to earn even *more* than \$50. This looks like easy money, so Clarke agrees to complete the survey. As she works through each question, quickly selecting the most probable of each pair of propositions, she eventually comes to question 17, which presents her with a choice between *s*, the proposition that salt dissolves in water, and *w*, the proposition that the world is more than five minutes old. The probability that salt dissolves in water is very high. Nevertheless, the probability that the world is more than five minutes old is even *higher*. Clarke sees this right away, so she quickly selects *w* instead of *s* and moves on to the next question.

I want to say two things about this case. First, I want to insist that Clarke cannot *rationally* select *s* as she makes her choice between *s* and *w*. My reasons for thinking this are simple. Clarke's only options are *s* and *w*. She does not have any third option. So what are her reasons for selecting *s*, and what are her reasons for selecting *w*? We can imagine that she has some quirky or eccentric reasons for selecting one option rather than the other. Perhaps, for example, she prefers propositions about salt to propositions about the world, or perhaps she prefers the bottom of any pair of propositions to the top of any pair of propositions, if those propositions are listed vertically on a page, or perhaps her hand just happens to be closer to *w* than to *s*. Whatever we think about "reasons" like these, I want to stipulate that they do not apply in the survey case. I want to stipulate that the only reasons Clarke has for selecting one option rather than the other are the reasons explicitly mentioned in the description of

the case. So, again, we can ask what reasons Clarke has for selecting s , and what reasons she has for selecting w .

Perhaps Clarke has *this* reason for selecting s : it is *true*. But of course, if she has this reason for selecting s , then she has exactly this same reason for selecting w , since w is true too. Here is another reason that Clarke might have for selecting s : she will get \$1 if she selects s . But again, if Clarke has this reason for selecting s , then she has exactly this same reason for selecting w , since she will get \$1 if she selects w . Indeed, when we run through all of the reasons that she might have for selecting s , and all of the reasons she might have for selecting w , we see that each of her actual or potential reasons for selecting s is matched by an equally weighty reason for selecting w . But there is a glaring reason Clarke has for selecting w that is not counterbalanced by any reason she has for selecting s : namely, that the probability of w is higher than the probability of s . This means that, on balance, Clarke's reason for selecting w decisively outweigh her reasons for selecting s , and we can stipulate that this balance of reasons is transparent to Clarke. So Clarke knows that, all things considered, w is a better option than s . Given this, however, I want to insist that, if Clarke had selected s instead of w , this would have been a paradigm case of irrational behavior. After all, Clarke would have selected the option that she knows is literally the *worst* of her options, all-things-considered. So, in the survey case, Clarke cannot rationally select s .

The second thing I want to point out is that, as Clarke makes her choice between s and w , Clarke does what is most preferable conditional on s only if she *does* select s . After all, conditional on s , the probability of s is literally 1, while the probability of w is less than 1. The

upshot is that, as Clarke answers question 17, Clarke cannot *rationally* do what is most preferable conditional on s .

This is where KP runs into trouble. According to KP, Clarke knows that salt dissolves in water as she answers question 17 only if she can rationally do what is most preferable conditional on the proposition that salt dissolves in water as she answers question 17. Since Clarke cannot rationally select s as she answers question 17, and since she must select s in order to do what is most preferable conditional on s —which is the proposition that salt dissolves in water—Clarke *cannot* rationally do what is most preferable conditional on the proposition that salt dissolves in water. So KP tells us that Clarke does not *know* that salt dissolves in water as she answers question 17. If KP is true, as soon as Clarke turns her attention from question 16 to question 17, Clarke loses her knowledge that salt dissolves in water.

But this is enormously implausible. Take all of the evidence that you and I have that salt dissolves in water. We have actually seen salt dissolve in water innumerable many times over the course of our lives. But by hypothesis, Clarke's evidence that salt dissolves in water is much better than ours. And moreover, Clarke can give a detailed explanation in terms of physics and chemistry for *why* salt dissolves in water. The proposition that salt dissolves in water is both supported incredibly well by her empirical evidence and coheres incredibly well with her overall knowledge and understanding of the world. It seems just enormously implausible that Clarke could lose her knowledge that salt dissolves in water by simply finding herself in a situation where she compares the probability that salt dissolves in water to the probability of some *completely unrelated proposition*—namely, the proposition that the

world is more than five minutes old—and sees that the latter is higher than the former. But this is exactly what KP entails.

Moreover, the only thing even *potentially* preventing Clarke from knowing that salt dissolves in water as she answers question 17 is that she cannot rationally do what is most preferable conditional on the proposition that salt dissolves in water. But this potential barrier is absent (we can suppose) both immediately before she answers question 17 and immediately after she answers question 17. So KP predicts that, as Clarke goes from question 16 to question 17 to question 18, she goes from knowing to not knowing and back to knowing again, all in a matter of seconds, and as the result of nothing more than turning her attention to, and then away from, the proposition that salt dissolves in water. But it seems utterly implausible that Clarke could lose and then regain her knowledge that salt dissolves in water by simply turning her attention to, and then away from, that proposition. So I want to insist that, because KP entails that Clarke does not know that salt dissolves in water as she answers question 17, and because it predicts that she goes from knowing to not knowing and then back to knowing as she moves from question 16 to question 17 to question 18, we ought to reject KP as false.

3 A Dilemma for Principle-Based Arguments

3.1 A Dilemma for Hawthorne and Stanley's KR

No pragmatist actually defends KP, so the fact that KP is false does not count *directly* against any pragmatist. Together with the survey case, however, KP causes serious problems for all

of the most plausible principles in the pragmatic encroachment literature. Start with Hawthorne and Stanley's Knowledge-Reason Principle.

(KR) S knows that p in c iff it is appropriate for S to treat the proposition that p as a reason for acting in c .¹³⁶

KR has a lot going for it. As Hawthorne and Stanley point out, it fits very nicely with our habit of criticizing people for acting on what they do not know, and our habit of defending our actions by citing what we *do* know.¹³⁷ The problem is, KR is either false because it entails KP, or irrelevant to the plausibility of pragmatism if it does not entail KP. Consider the following argument.

KR: S knows that p in c iff it is appropriate for S to treat the proposition that p as a reason for acting in c .

Bridge 1: It is appropriate for S to treat the proposition that p as a reason for acting in c only if S can rationally do what is most preferable conditional on p in c .

∴ **KP:** S knows that p in c only if she can rationally do what is most preferable conditional on p in c . [from KR and Bridge 1]

¹³⁶ Hawthorne and Stanley (2008), p. 578. Before presenting KR, Hawthorne and Stanley define a notion of p -dependence according to which S 's choice between options o^1 through o^n is p -dependent just in case the most preferable of o^1 through o^n conditional on p differs from the most preferable of o^1 through o^n conditional on $\neg p$. Fully stated, KR says that, *where S 's choice is p -dependent*, S knows that p in c iff it is appropriate for S to treat the proposition that p as a reason for acting in c . By relativizing the principle to p -dependent choices, Hawthorne and Stanley avoid counterexamples from cases where S knows that p but cannot treat p as a reason for action because p is *irrelevant* to the action in question. Since I will only be concerned with cases where p is *relevant* to the action in question, I have omitted Hawthorne and Stanley's clause about p -dependence.

¹³⁷ Once the notion of treating a *proposition* as a reason for acting is properly clarified, that is. See the exchange between Stanley and Jussi Suikkanen at <http://peasoup.typepad.com/peasoup/2007/07/hawthorne-and-s.html>.

This argument is clearly valid, so KR and the bridge principle together entail KP. But as I just argued, KP is false. It follows that either KR is false or the bridge principle is false. Either way, KR cannot be employed in a successful argument for pragmatism.

Focus on the bridge principle. Either it is true, or it is false. If it is *true*, then KR entails KP. Since KP is false, KR is therefore false. Suppose that the bridge principle is *false*, then. In this case, it might be appropriate for *S* to treat the proposition that *p* as a reason for acting in *c* even while she *cannot* rationally do what is most preferable conditional on *p* in *c*. If it might be appropriate for *S* to treat the proposition that *p* as a reason for acting in *c* even while she cannot rationally do what is most preferable conditional on *p* in *c*, however, then how is KR supposed to support pragmatism?

Pragmatism, recall, is the view that knowledge depends at least in part on our practical interests. So, what exactly are *practical interests*? Pragmatists do not say exactly what practical interests are, but a survey of the literature makes it perfectly clear that, by ‘practical interests,’ pragmatists just mean the values and/or consequences that would figure into an expected utility calculation. This means that *S*’s practical interests *vis-à-vis* the proposition that *p* *just are* the sorts of things (other than the probabilities of *p* and its negation) that make a difference to whether *S* can rationally do what is most preferable conditional on *p*. If the bridge principle is false, however, whether it is appropriate for *S* to treat *p* as a reason for acting will not depend on whether she can rationally do what is most preferable conditional on *p*. So, if the bridge principle is false, whether it is appropriate for *S* to treat *p* as a reason for acting will not depend on her *practical interests*. If the propriety of treating *p* as a reason for acting does not depend on *S*’s practical interests, however, then why should we think that *S*’s

knowledge depends on her practical interests? If KR is true, whether *S* knows that *p* *does* depend on whether it is appropriate for *S* to treat *p* as a reason for acting. But if the bridge principle is false, whether it is appropriate for *S* to treat *p* as a reason for acting does not depend on her practical interests. Given that the bridge principle is false, KR gives us no reason to think that *knowledge* depends on our practical interests.

So, the problem for KR is this: Either the bridge principle is true, or it is false. If it is true, then KR entails KP; since KP is false, KR is false too. But on the other hand, if the bridge principle is false, then it is unclear how KR lends any support to pragmatism. Since we cannot successfully argue for pragmatism on the basis of a false principle, and since we cannot successfully argue for pragmatism on the basis of a principle that does not seem to support pragmatism, KR seems ill-suited for defending pragmatism.¹³⁸

3.2 Dilemmas for Fantl and McGrath's KJ, and Ross and Schroeder's KA

Exactly the same problem afflicts Fantl and McGrath's Knowledge-Justification Principle, and Ross and Schroeder's Knowledge-Action Principle, below.

(KJ) *S* knows that *p* in *c* only if *p* is warranted enough to justify *S* in φ -ing in *c*, for any value of ' φ .'¹³⁹

¹³⁸ The fact that KR and the bridge principle jointly entail KP is no problem for the *sufficiency* direction of KP, of course. But even if the sufficiency direction is true and knowing that *p* suffices for the propriety of treating *p* as a reason for acting, this fact is of little use to the pragmatist. Even if whether *S* can treat *p* as a reason for acting in *c* depends on whether *S* knows that *p* in *c*, it does not follow that whether *S* knows that *p* in *c* depends on whether *S* can treat *p* as a reason for acting in *c*. If the sufficiency direction of KR is true, purists can happily accept it.

¹³⁹ Fantl and McGrath (2009), p. 66. As Fantl and McGrath make clear, KJ does not say that, if *S* knows that *p*, then the amount of warrant that she has for *p* entails that she *is* justified in φ -ing, for any value of ' φ .' Rather, the way Fantl and McGrath explicitly cash out "warranted enough," KJ says that, if *S* knows that *p* and she is *not* justified in φ -ing, it is not *for lack of warrant with respect to p*. So KJ is considerably more plausible than it initially strikes many people.

(KA) S knows that p in c only if S can rationally act as if p in c .¹⁴⁰

Fantl and McGrath provide an illuminating argument for KJ, and Ross and Schroeder show how KA follows from their account of the nature of belief together with independently plausible principles. Unfortunately, like KR, each principle is either false because it entails KP, or apparently irrelevant to the plausibility of pragmatism if it does not entail KP. Consider the following valid arguments.

KJ: S knows that p in c only if p is warranted enough to justify S in φ -ing in c , for any value of ' φ .'

\therefore **KJ':** S knows that p in c only if p is warranted enough to justify S in doing what is most preferable conditional on p in c . [from KJ]

Bridge 2: If p is warranted enough to justify S in doing what is most preferable conditional on p in c , then S can rationally do what is most preferable conditional on p in c .

\therefore **KP:** S knows that p in c only if she can rationally do what is most preferable conditional on p in c . [from KJ' and Bridge 2]

KA: S knows that p in c only if she can rationally act as if p in c .

Bridge 3: S can rationally act as if p in c only if she can rationally do what is most preferable conditional on p in c .

¹⁴⁰ Ross and Schroeder (forthcoming), p. 4.

∴ **KP**: S knows that p in c only if she can rationally do what is most preferable conditional on p in c . [from KA and Bridge 3]

These arguments are valid, and they both conclude with KP. Since KP is false, KJ and KA are both false unless the relevant bridge principles are false. But if the bridge principles are false, it will be unclear how KJ or KA support pragmatism, just as it is unclear how KR supports pragmatism if its bridge principle is false. If Bridge 2 is false, whether p is warranted enough to justify S in doing what is most preferable conditional on p will not depend on whether S can rationally do what is most preferable conditional on p . So, if Bridge 2 is false, whether p is warranted enough to justify S in doing what is most preferable conditional on p will not depend on her *practical interests*. But in this case, even if KJ is true, KJ will give us no reason to think that *knowledge* depends on practical interests. If Bridge 2 is false, it seems that KJ could be true even while knowledge does *not* depend on practical interests, since (again) practical interests *just are* the sorts of things (other than the relevant probabilities) that make a difference to whether we can rationally do what is most preferable conditional on a given proposition. Likewise for KA. If Bridge 3 is false, whether S can rationally act as if p will not depend on whether S can rationally do what is most preferable conditional on p . So, if Bridge 3 is false, whether S can rationally act as if p will not depend on her practical interests. In this case, even if KA is true, KA will give us no reason to think that S 's *knowledge* depends on her practical interests. So KJ and KA are either false because they entail KP, or ill-suited for defending pragmatism if they do not entail KP. The upshot is that we cannot employ KR, KJ or KA in a successful argument for pragmatism.

3.3 *A Counterexample to Weatherson's BQ and Ganson's BC*

Where Hawthorne and Stanley's KR, Fantl and McGrath's KJ, and Ross and Schroeder's KA all focus directly on the connection between knowledge and practical interests, Brian Weatherson and Dorit Ganson focus on the connection between rational or justified *belief* and practical interests. Weatherson defends the Belief-Question Principle, below, while Ganson defends the Belief-Credence Principle.

- (BQ) S believes that p in c only if conditionalizing on p in c does not change S 's answer to any relevant question in c .¹⁴¹
- (BC) S believes that p in c only if her credence that p in c is high enough to insure that she is willing to act as if p in c .¹⁴²

BQ and BC both make justified or rational belief depend on practical interests. Consider High Stakes, for example. BQ and BC both entail that, unless Hannah's credence that the bank will be open is much higher than it epistemically ought to be, Hannah does not believe that the bank will be open in High Stakes. So, according to both principles, Hannah believes that the bank will be open in High Stakes only if she is epistemically irrational in her credence that the bank will be open. On the assumption that Hannah cannot know that the bank will be open while being epistemically irrational in her credence that the bank will be open, BQ and BC both entail that Hannah does not know that the bank will be open in High Stakes.

¹⁴¹ Weatherson (2012), p. 14 of version available on Weatherson's website.

¹⁴² Ganson (2008), p. 443.

BQ and BC both have a lot going for them. They provide an elegant solution to the lottery paradox, for example, and they fit very nicely with functionalism about belief.¹⁴³ If the broad functionalist picture is correct, we should not be surprised if something like BQ or BC is true. Unfortunately, BQ and BC suffer exactly the same problem as KP. Take BQ first, and return to the survey case. As Clarke is answering question 17, she is making a choice between selecting s , the proposition that salt dissolves in water, and selecting w , the proposition that the world is more than five minutes old. So, here is a question that is relevant as Clarke answers question 17: which is more probable, s or w ? Clarke's answer to this question is that w is more probable than s . But conditional on the proposition that salt dissolves in water, the probability of s is 1, while the probability of w is lower than 1. So, conditional on the proposition that salt dissolves in water, the probability of s is higher than the probability of w . This means that conditionalizing on the proposition that salt dissolves in water *does* change Clarke's answer to a relevant question as she is answering question 17. So, according to BQ, Clarke does not believe that salt dissolves in water as she answers question 17. But knowledge entails belief. So, according to BQ, Clarke does not *know* that salt dissolves in water as she answers question 17. BQ has exactly the same false consequence as KP.

The same goes for Ganson's BC. By "act as if p ," Ganson means that S does what is best conditional on p .¹⁴⁴ But conditional on the proposition that salt dissolves in water, selecting s is a better option than selecting w . Since s and w are Clarke's only options as she answers question 17, Clarke is willing to act as if salt dissolves in water only if she is willing to select s

¹⁴³ See (2012), p. 16 of version available on Weatherson's website.

¹⁴⁴ Ganson (2008), p. 443.

instead of w . But of course, as Clarke answers question 17, her credence that salt dissolves in water *is not* high enough to insure that she is willing to select s instead of w . So, according to BC, Clark does not *believe* that salt dissolves in water as she answers 17. Since knowledge entails belief, BC says that Clarke does not *know* that salt dissolves in water as she answers question 17. BC therefore has exactly the same false consequence as KP and BQ, and we ought to reject all three principles for the same reason.¹⁴⁵

3.4 Pragmatism and the Apparent Irrelevance of Schroeder's RB

While Hawthorne and Stanley's KR, Fantl and McGrath's KJ, and Ross and Schroeder's KA are all either false or apparently irrelevant to the plausibility of pragmatism, Weatherson's BQ and Ganson's BC seem clearly false. Before turning to *case-based* arguments for pragmatism, I want to consider a final principle that pragmatists might rely on to motivate their view: Mark Schroeder's Rational Belief Principle, below.

- (RB) It is epistemically rational for S to believe that p in c just in case, in c , S has at least as much epistemic reason to believe that p as to believe that $\neg p$, and S has at least as much epistemic reason to believe that p as to withhold with respect to p .¹⁴⁶

¹⁴⁵ In Weatherson (2012), Weatherson also defends a principle according to which S 's knowing that p in c would entail that, if S were to calculate the expected utility of some course of action in c , she could legitimately leave the possibility that $\neg p$ off of her decision table. This principle seems false for the same reason as KP, BQ and BC are false. After all, if Clarke were to calculate the expected utility of selecting s as she answers question 17, it would *not* be legitimate for her to leave the possibility that salt does not dissolve in water off of her decision table. So, like KP, BQ and BC, Weatherson's principle says that, as Clarke answers questions 17, she does not know that salt dissolves in water.

¹⁴⁶ Schroeder (2012), p. 274. Schroeder calls this principle "Bf-Sufficiency." I have changed its name to avoid abbreviating it "BS."

Two things should strike us immediately about RB, I think. First, it seems *trivially true*. It amounts to the claim that S cannot rationally believe that p if her epistemic reasons for *not* believing that p outweigh her epistemic reasons *for* believing that p . The second thing that should strike us, I think, is that RB seems perfectly consistent with *purism*. After all, why can't purists agree that S cannot rationally believe that p if her epistemic reasons for *not* believing that p outweigh her epistemic reasons *for* believing that p ?

That RB is *not* consistent with purism is the main thing Schroeder means to establish, and his argument from RB to pragmatism goes as follows. First, Schroeder distinguishes between reasons to believe a proposition and reasons to withhold with respect to that proposition, and then argues that reasons to withhold (whatever they are) cannot be *evidence*.

[T]he evidence is exhausted by evidence which supports p and evidence which supports $\neg p$. But the evidence which supports p is reason to believe p , and the evidence which supports $\neg p$ is reason to believe $\neg p$. Consequently the reasons to withhold must come from somewhere else. So they cannot be evidence.¹⁴⁷

What *are* the reasons to withhold, then? Schroeder does not offer a complete list, but he does say that costs of error can be an important reason to withhold.

[A] natural place to look for reasons to withhold is in the costs of error. When you form a belief, you take a risk of getting things wrong that you do not take by withholding. In contrast, when you withhold, you guarantee that you miss out on getting things right. So plausibly, one important source of reasons to withhold will

¹⁴⁷ Ibid., p. 276.

come from the preponderance of the cost of having a false belief over the cost of missing out on having a true belief—or, as I will put it, the preponderance of the cost of *type-1 error* over *type-2 error*.¹⁴⁸

This paragraph seems plausible enough. But so far, it lends no support to pragmatism, for the cost of type-1 errors might just be the possession of false beliefs, and the cost of type-2 errors might just be the absence of true beliefs. If these are the *only* costs of type-1 and type-2 errors, however, then *practical interests* seem beside the point. For example, if these are the only costs, then the costs of type-1 and type-2 errors will be the same in High Stakes as they are in Low Stakes. According to Schroeder, however, the costs of type-1 and type-2 errors will often go well *beyond* false belief and the absence of true belief, respectively.

[T]he most general sort of cost of type-1 error is simply mistakes that we make, when we act on a belief that turns out to be false. Correlatively, the most general sort of cost of type-2 error derives from the fact that sometimes we simply have to act, and ignorance does not help us. These two sorts of costs—of type-1 and type-2 errors—are clearly practical in nature. They derive from the connection between belief and action. Gratifyingly, in High Stakes the costs of type-1 error are extremely high, and the costs of type-2 error are very low, which on this picture supports the view that there are especially strong reasons to withhold in that case.¹⁴⁹

So, on Schroeder's view, Hannah's epistemic reasons to withhold with respect to the proposition that the bank will be open in High Stakes outweigh her epistemic reasons to

¹⁴⁸ Ibid., p. 277.

¹⁴⁹ Ibid.

believe that the bank will be open in High Stakes. According to RB, then, Hannah cannot rationally believe that the bank will be open in High Stakes. Since knowledge entails rational belief, Hannah cannot *know* that the bank will be open in High Stakes.¹⁵⁰ And of course, if Hannah fails to know that the bank will be open in High Stakes, then pragmatism is true, since the only relevant difference between Low Stakes and High Stakes (we are supposing) is the stipulated difference in Hannah's practical interests. This is how Schroeder derives pragmatism from RB.

What should we make of Schroeder's argument? Purists might balk at Schroeder's claim that reasons to withhold cannot be evidence, or deny that the costs of a type-1 error with respect to the proposition that the bank will be open give Hannah any *epistemic* reason to withhold, but the response I favor has purists commenting on the nature of *belief*. In the quotation above, Schroeder says that the costs of a type-1 error are very high for Hannah in High Stakes, and he says that these costs "derive from the connection between belief and action." What exactly *is* the connection between belief and action, then? Schroeder does not say, but his argument seems to presuppose the following Belief-Preference Principle.

(BP) *S* believes that *p* in *c* only if she is willing to do what is most preferable conditional on *p* in *c*.

As we saw above, going straight home is the most preferable of Hannah's options conditional on the proposition that the bank will be open. So, according to BP, Hannah believes that the bank will be open only if she is willing to go straight home. Suppose that

¹⁵⁰ Again, I am taking it for granted that a belief or action is rational just in case it is *not* irrational.

BP is *false*, then. In this case, it is possible for *S* to believe that *p* in *c* without being willing to do what is most preferable conditional on *p* in *c*. So, if BP is false, it should be possible for Hannah to believe that the bank will be open in High Stakes without being willing to go straight home. If Hannah can believe that the bank will be open while she is unwilling to go straight home, however, then why should we think that, in High Stakes, the costs of type-1 error with respect to the proposition that the bank will be open are extremely high for Hannah? If BP is false, there seems no reason why Hannah cannot believe that the bank will be open and simultaneously take all of the precautions necessary for insuring that she does not incur the large fine. But if Hannah can do *this*—if she can believe that the bank will be open and simultaneously take all of the precautions necessary for insuring that she does not incur the fine—then the cost of a type-1 error with respect to the proposition that the bank will be open will be no higher for Hannah than forming a false belief. And this will be the case in both High Stakes *and* Low Stakes. So, if BP is false, then costs of type-1 error with respect to the proposition that the bank will be open seem to be the same for Hannah in Low Stakes and High Stakes. It follows that, unless BP is *true*, the costs of a type-1 error with respect to the proposition that the bank will be open should be no higher for Hannah in High Stakes than they are in Low Stakes. Thus, Schroeder's argument that the costs of type-1 error increase as the stakes go up seems to presuppose that BP is true. And since Schroeder's argument for pragmatism depends on the claim that the costs of type-1 error increase as the stakes go up, his argument for pragmatism seems equally dependent on the truth of BP.

Fortunately for purists, BP looks false. Again, consider the survey case. As Clarke is answering question 17, she is choosing between *s* and *w*. As we have already seen, selecting *s* is the most preferable of Clarke's options conditional on the proposition that salt dissolves in water. So, according to BP, Clarke *believes* that salt dissolves in water as she answers question 17 only if she selects *s*. But Clarke *does not* select *s*. She selects *w*. So, according to BP, Clarke does not believe that salt dissolves in water as she answers question 17. Since knowledge entails belief, BP tells us that Clarke does not *know* that salt dissolves in water as she answers question 17. BP therefore has the same false consequence as KP, BQ and BC, and we ought to reject all four of them for the same reason. Since Schroeder's argument from RB to pragmatism seems to presuppose BP, it is therefore unclear how RB supports pragmatism.

3.5 *Taking Stock*

Hawthorne and Stanley's KR, Fantl and McGrath's KJ, Ross and Schroeder's KA, Weatherson's BQ, Ganson's BC, and Schroeder's RB are the most plausible principles in the pragmatic encroachment literature. If any principle-based argument for pragmatism will succeed, an argument based on one of *these* principles should succeed. But these principles are all either false or apparently irrelevant to the plausibility of pragmatism. It therefore seems doubtful that any principle-based argument for pragmatism will succeed.

4 **Case-Based Arguments**

If principle-based arguments fail, then pragmatism is ultimately unmotivated unless *case-based* arguments succeed, since principle-based arguments and case-based arguments are all the arguments there are. Unfortunately for pragmatism, without the help of principle-based

arguments, case-based arguments seem clearly incapable of motivation pragmatism. (I will move rather quickly here, since, to my knowledge, there is only *one* case-based argument in the literature: Stanley’s argument in *Knowledge and Practical Interests*.)

The relevant cases are the bank cases we have been considering all along, plus Fantl and McGrath’s train cases, Cohen’s plain cases, Weatherson’s genie cases, and Ross and Schroeder sandwich cases.¹⁵¹ These cases differ from Low and High Stakes only with respect to irrelevant details, so problems that arise from Low and High Stakes will afflict any argument for pragmatism that we might base on an alternative pair of cases. Now, as we noted in §1, pragmatism is the view that knowledge depends at least in part on our practical interests, in the sense that a *mere* difference in practical interests might entail a difference in knowledge. As we also noted in §1, there is intuitively a difference in knowledge between Low Stakes and High Stakes. Intuitively, Hannah knows that the bank will be open in Low Stakes, but does not know that the bank will be open in High Stakes. Low and High Stakes do not differ *just* with respect to Hannah’s practical interests, however. Low and High Stakes are originally due to Keith DeRose, who relied on these cases to argue for *contextualism*—the thesis that the word ‘know’ expresses different properties or relations in different contexts of utterance.¹⁵² This is why Low and High Stakes both contain utterances of the word ‘know.’ In Low Stakes, Hannah says to Sarah, “I know that the bank will be open tomorrow,” and then, in High Stakes, she says to Sarah, “I *do not* know that the bank will be open tomorrow.” An obvious worry about any argument for pragmatism based on High Stakes is that our

¹⁵¹ See Fantl and McGrath (2002), Cohen (1999), Weatherson (2012), and Ross and Schroeder (forthcoming). Note that Fantl and McGrath, Weatherson, and Ross and Schroeder respectively use these cases to illustrate their views, not to *motivate* them.

¹⁵² See DeRose (1992).

intuitions about Hannah might track what Hannah tells us about her knowledge rather than her *practical interests*. This becomes especially apparent when we consider cases like Revised High Stakes, below.

Revised High Stakes: Hannah is approaching her bank on Friday afternoon, as in Low Stakes. Again she notices the long lines, and again she considers driving straight home and depositing the check Saturday morning, when she is out running errands. But in this case, Hannah knows that she will incur an enormous fine if she does not deposit her check before noon on Saturday, so she decides that she had better not risk it. As Hannah pulls into the parking lot, however, Sarah complains about the long lines and says “Why are you stopping? The bank will be open tomorrow, so you can just deposit your check when you’re out running errands.” Hannah replies, “Yeah, I know. But I figure I ought to play it safe and deposit the check now.”

Without the aid of principles like Hawthorne and Stanley’s KR or Fantl and McGrath’s KJ, I do not have any intuition that Hannah lacks knowledge in Revised High Stakes, and many purists have the intuition that in fact Hannah *does* know that the bank will be open in Silent High Stakes. Low Stakes and Revised High Stakes exhibit exactly the same difference in practical interests as Low Stakes and High Stakes, however. So, it seems doubtful that *pragmatism* is the best explanation of our intuitions about cases like Low Stakes and High Stakes. Our intuitions seem to track what Hannah tells us about her knowledge in High Stakes, not her practical interests.

Now, when I say “our intuitions,” I am talking about *my* intuitions and the intuitions of other *philosophers*. This might be a problem, for the intuitions of the *folk* might be the intuitions that really matter. So, perhaps, to adequately evaluate the prospects for case-based arguments for pragmatism, we need to do some experimental philosophy and see if the *folk* have pragmatist intuitions.

The relevant experiments have been conducted, and the results are not promising for case-based arguments. Josh May, Walter Sinnott-Armstrong, Jay Hull, Aaron Zimmerman, Adam Feltz, Chris Zarpentine, Wesley Buckwalter, Jonathan Schaffer, and Joshua Knobe have all conducted experiments to see whether the folk have pragmatist intuitions, and they claim to show that in fact the folk *do not* have pragmatist intuitions.¹⁵³ Not everyone agrees with them. Angel Pinillos, Chandra Sripada, and Jason Stanley have conducted their own experiments, and they claim to show that, in fact, the folk *do* have pragmatist intuitions.¹⁵⁴ In the most recent instalment in this literature, however, Buckwalter and Schaffer undermine the methodology employed in the experiments of Pinillos, Sripada, and Stanley, and argue that, everything considered, the data suggest that in fact the folk do not have pragmatist intuitions.¹⁵⁵ This literature is subtle and complicated, and I lack room for the details here. But suffice it to say, given the data currently in our possession, it is certainly not clear that the folk *do* have pragmatist intuitions. So, shifting attention away from the intuitions of philosophers and toward the intuitions of the folk does not seem to help.

¹⁵³ See Buckwalter (2010), May, Sinnott-Armstrong, Hull, and Zimmerman (2010), Feltz and Zarpentine (2010), and Knobe and Schaffer (2012).

¹⁵⁴ See Pinillos (2012), and Sripada and Stanley (2012).

¹⁵⁵ Buckwalter and Schaffer (forthcoming).

Moreover, even if it eventually becomes clear that the folk *do* have pragmatist intuitions, it will still not follow that we therefore ought to accept pragmatism. After all, it is not clear that the point of theorizing about knowledge is *simply* to capture the intuitions of the folk. Consider the following claim.

Balance: The *best* theory of knowledge, all things considered, will strike the right balance between the intuitions of the folk, on the one hand, and the subtle theoretical considerations that epistemologists worry about, on the other.

Balance is plausible. But if it is right, we should be unsurprised if, at the end of the day, the best theory of knowledge does not perfectly capture the intuitions of the folk. So, even if turns out that the folk *do* have pragmatist intuitions, it will not follow that we therefore ought to accept pragmatism.

Case-based arguments fail for the following three reasons, then. First, while most philosophers have the requisite intuition about the *original* high-stakes case (the one due to Keith DeRose), it is not clear that the correct explanation of this intuition has much to do with Hannah's *practical interests*, as becomes clear when we compare the original High Stakes to Revised High Stakes. Second, even if we should be worrying about the intuitions of the *folk* instead of the intuitions of *philosophers*, it is not clear that the *folk* have pragmatist intuitions. The data currently in our possession suggest that in fact they do not. Third, even if it eventually becomes clear that the folk *do* have pragmatist intuitions, it still will not follow that we therefore ought to accept pragmatism, since principles like Balance are plausible, and since, if any principle like Balance is true, the mere fact that the folk have pragmatist

intuitions would not count decisively in favor of pragmatism. Everything considered, then, it looks like case-based arguments for pragmatism fail.

5 A Challenge for Purism

But if case-based arguments fail, then what is supposed to motivate pragmatism? *All* of the arguments in the literature are either case-based arguments or principle-based arguments, and we have already seen that principle-based arguments fail. So what reason do we have for accepting pragmatism?

I see one final consideration that pragmatists might rely on to motivate their view. It consists in a challenge that pragmatists might pose to purists. Return to High Stakes. If purism is true, then, presumably, Hannah knows *both* that the bank will be open on Saturday, *and* that, conditional on the bank's being open on Saturday, going straight home is the most preferable of her options. But purists agree with pragmatists that Hannah cannot *rationally* go straight home. So, how can this be? How can it be that Hannah cannot rationally go straight home in High Stakes, even though, in High Stakes, she knows that the bank will be open on Saturday, and also knows that, *conditional on the bank's being open on Saturday*, going straight home is the most preferable of her options? Purists owe us an explanation. From the perspective *of their own view*, how can it be that Hannah cannot rationally go straight home?

I think purists can meet this challenge. Consider Hannah's response to Sarah in Revised High Stakes. In Revised High Stakes, Sarah tells Hannah that, because the bank will be open on Saturday, Hannah can deposit her check when she is out running errands on Saturday morning. Hannah's reply is that, while she knows that the bank will be open on Saturday

morning (and presumably also knows that, conditional on the bank's being open on Saturday morning, going straight home is the most preferable of her options), going straight home would be too *risky*. From the perspective of purism, Hannah's response to Sarah makes perfect sense. If her response makes sense, however, then the reason that Hannah cannot rationally go straight home in Revised High Stakes is simply that she knows that going straight home would be too risky. If purism is true, Hannah's knowing that going straight home would be too risky is perfectly consistent with her knowing both that the bank will be open on Saturday morning and that, conditional on the bank's being open on Saturday morning, going straight home is the most preferable of her options. Since this is exactly what Hannah communicates to Sarah, and since Hannah's response to Sarah makes perfect sense, given that purism is true, it is not mysterious from the perspective of purism why Hannah cannot rationally go straight home in Revised High Stakes.

But of course, it is no more mysterious why Hannah cannot rationally go straight home in *High Stakes*. High Stakes is identical to Revised High Stakes in terms of the relevant items of knowledge, after all. So, if purism is true, the explanation for why Hannah cannot rationally go straight home in High Stakes is exactly the same as the explanation for why she cannot rationally go straight home in Revised High Stakes: she knows that going straight home would be too risky. Again, if purism is true, this is perfectly consistent with Hannah's knowing both that the bank will be open on Saturday morning and that, conditional on the bank's being open on Saturday morning, going straight home is the most preferable of her options. So it is *not* mysterious, from the perspective of purism, why Hannah cannot

rationally go straight home in High Stakes. Purists can meet this final challenge that pragmatists might pose.

6 Conclusion

In §2 and §3, I argued that no principle-based argument for pragmatism succeeds, and, in §4, I argued that no case-based argument for pragmatism succeeds. Principle-based arguments and case-based arguments are all the arguments there are, however, so it looks like there is no successful argument for pragmatism. In §5, I considered the possibility of motivating pragmatism by challenging purists to explain why, in High Stakes, Hannah cannot rationally go straight home. But as I just argued, purists *can* explain why Hannah cannot rationally go straight home in High Stakes. What motivates pragmatism, then? I see no satisfactory answer to this question, so I conclude that pragmatism is ultimately unmotivated.¹⁵⁶

¹⁵⁶ Thanks to Lisa Miracchi, Ernest Sosa, and audiences at Rutgers, Georgetown, and Notre Dame for helpful comments and conversation about this paper.

Chapter 5: Advice to Minimalists

Contextualists, pragmatists, and non-skeptical purist invariantists offer competing analyses of the following familiar case.

High Stakes: It's Friday afternoon. Rachel and Keith must deposit their paychecks before Sunday. As they approach the bank they notice long lines, so Keith suggests that they deposit their checks Saturday morning. "I stop at the bank every Saturday and it's open until noon," he tells Rachel. Rachel reminds Keith that it's extremely important that they deposit their checks before Sunday. "What if our bank has changed its Saturday hours since you last visited," she asks. "Do you *know* that it will be open tomorrow?" Still believing that the bank *will* be open tomorrow, Keith answers, "I guess I don't know that the bank will be open tomorrow. I'd better go in and make sure."¹⁵⁷

Where $\neg K$ is the sentence 'I don't know that the bank will be open,' contextualists and pragmatists agree that Keith says something true by uttering $\neg K$. Pragmatists think Keith says something true because he lacks *knowledge*. Keith would be irrational to *act* on the proposition that the bank will be open, so he does not *know* that it will be open. Contextualists think Keith says something true because the word 'know' is context sensitive: in the context of Keith's utterance of $\neg K$ in High Stakes, 'know' expresses a property that is

¹⁵⁷ For the original high-stakes case and its low-stakes counterpart, see DeRose (1992), p. 913. Contextualists, pragmatists and non-skeptical purist invariantists all agree about the correct interpretation of the low-stakes case, so I leave it out of this paper entirely. The low-stakes case will play no role in the arguments that follow.

not instantiated by Keith's belief that the bank will be open, so Keith expresses a true proposition by uttering $\neg K$. Non-skeptical purist invariantists reject all of these claims. According to non-skeptical purist invariantism, Keith *does* know that the bank will be open, so he expresses a *false* proposition by uttering $\neg K$. Since non-skeptical purist invariantists think they can get by in their epistemological theorizing without the contextualist thesis that 'know' expresses different properties in different contexts, or the pragmatist thesis that practical circumstances and concerns encroach on knowledge, I will call non-skeptical purist invariantism, simply, 'minimalism.'

High Stakes causes trouble for minimalism, but not the trouble that we are accustomed to hearing about. The standard argument against minimalism says that we must adopt either contextualism or pragmatism in order to explain why Keith seems to say something true by uttering $\neg K$ in High Stakes. Minimalists do not find the standard argument compelling. *Ultima facie* seemings are what matter here, and minimalists think that, with all of the relevant considerations in view, Keith seems to say something *false* by uttering $\neg K$ in High Stakes. An interesting question, however, is whether Keith's utterance of $\neg K$ is *appropriate*. Many minimalists will concede that it is. This admission, however, puts contextualists and pragmatists in position to run an argument against minimalism that takes the *propriety* of Keith's utterance as a datum rather than its apparent *truth*. Call this argument 'the propriety argument.'

The Propriety Argument

- (1) Keith's utterance of $\neg K$ is perfectly appropriate.

- (2) The correct explanation of (1) is either that Keith says something true by uttering $\neg K$ or that Keith implicates something true by uttering $\neg K$.
- (3) Keith does not implicate anything true by uttering $\neg K$.
- ∴ (4) The correct explanation of (1) is that Keith says something true by uttering $\neg K$.
- (5) Keith says something true by uttering $\neg K$ only if minimalism is false.
- ∴ (6) Minimalism is false.

Now minimalists will object, not just to the *conclusion* of this argument, but to my claim that it constitutes an interesting objection to minimalism *in its own right*. In response to *the standard argument*, many minimalists have already adopted the following “warranted assertability manoeuvre” (‘WAM’ for short).

Keith seems to *say* something true by uttering $\neg K$ because he *implicates* something true by uttering $\neg K$ and we mistake this true thing that he implicates by uttering $\neg K$ for what he says by uttering $\neg K$.¹⁵⁸

This WAM rejects premise (3) of the propriety argument, so minimalists who have already adopted this WAM in response to the *standard argument* will deny that the propriety argument even gets off the ground. These minimalists will think that they can deal with propriety argument exactly the way they have already dealt with the standard argument: with the above WAM.

¹⁵⁸ See, for example, Brown (2006).

I agree with minimalists that, everything considered, Keith seems to say something *false* by uttering $\neg K$ in High Stakes, so I find the standard argument unconvincing.¹⁵⁹ But Keith's utterance of $\neg K$ does seem appropriate, and I doubt that Keith *implicates* anything of interest with this utterance. So I accept premises (1) and (3) of the propriety argument and thereby reject the above WAM. Since (5) is beyond question, I am forced to choose between minimalism and premise (2). I think premise (2) is mistaken, and I want to encourage minimalists to join me in responding to the propriety argument by rejecting just premise (2). In §§1-5, I argue that Keith does not implicate anything of interest by uttering $\neg K$, so (3) is true. And in §2, I give an explanation of the propriety of Keith's utterance that puts minimalists in excellent position to accept (1) and reject (2). An upshot of my argument is that minimalists are mistaken to rely on WAMs in response to High Stakes.

1 The Rysiew/Brown WAM

According to Patrick Rysiew, Keith semantically expresses the same proposition by uttering K in Low Stakes as he would semantically express by uttering K in *any context*: namely, the proposition that he knows that the bank will be open. Likewise for $\neg K$. By uttering $\neg K$ in any context, Keith would semantically express the proposition that he does *not* know that the bank will be open.¹⁶⁰ Rysiew endorses a relevant alternatives account of knowledge. On his view, S knows that p iff S can rule out all of the relevant $\neg p$ alternatives—where a relevant $\neg p$ alternative is a scenario where p is false that is likely to obtain.¹⁶¹ In High Stakes, Keith

¹⁵⁹ I give the full battery of considerations in “Pragmatism and the Nature of Belief” and “A Paradox for Justification, Knowledge and Practical Rationality.”

¹⁶⁰ Rysiew (2000), p. 487.

¹⁶¹ *Ibid.*, pp. 487-8.

can rule out all the likely scenarios where the bank will not be open. So, according to Rysiew, Keith knows the bank will be open in High Stakes, and he therefore says something false by uttering $\neg K$ in High Stakes. An irrelevant alternative may become salient in a context, however, and Rysiew thinks that an instance of ‘I know that p ’ uttered in such a context would implicate that the speaker can rule out this irrelevant alternative.¹⁶² When Rachel mentions the possibility that the bank has recently discontinued its Saturday hours, she makes this possibility salient. Given the evidence currently in Keith’s possession, Keith cannot rule this possibility out. But since it is *unlikely* that the bank has recently discontinued its Saturday hours, this alternative is irrelevant. Thus, argues Rysiew, it doesn’t matter that Keith can’t rule this possibility out; he still knows that the bank will be open on Saturday. The possibility that the bank recently discontinued its Saturday hours is still salient in High Stakes, however, so Keith cannot felicitously utter K in High Stakes. If he did, he would implicate that he *can* rule out the possibility that the bank has recently discontinued its Saturday hours. Since Keith knows that he cannot rule this possibility out, he utters $\neg K$ instead. And this way, says Rysiew, Keith implicates that he cannot rule out the possibility that the bank has recently discontinued its Saturday hours. Because the proposition that Keith cannot rule out the possibility that the bank has recently discontinued its Saturday hours is both true and easily mistaken for the proposition that Keith semantically expresses by uttering $\neg K$ —namely, the proposition that Keith does not know that the bank will be open on Saturday—it’s tempting to think that Keith says something true by uttering $\neg K$ in

¹⁶² Ibid., p. 490.

High Stakes. And this, according to Rysiew, is why Keith seems to say something true by uttering $\neg K$ in High Stakes.

This is the gist of Rysiew's WAM. But how, *exactly*, is Keith supposed to implicate that he cannot rule out the possibility that the bank has recently discontinued its Saturday hours by uttering $\neg K$? The mere *salience* of this possibility leaves it mysterious how Keith would have implicated that he can rule out the possibility that the bank has recently discontinued its Saturday hours by simply uttering K . After all, Rachel could have made just about any possibility salient. She could have raised the possibility that Caesar had a lisp, or the possibility that the Browns will win the Superbowl, or the possibility that it is raining in Novosibirsk, or ... you name it. And surely, Keith would not have implicated that he could rule out any of *these* possibilities by uttering K . So salience *alone* does not adequately explain how Keith's uttering K would have implicated that he can rule out the possibility that their bank has recently stopped opening on Saturdays. As DeRose points out, without telling us how general conversational principles, the semantics of the word 'knows,' and the details of High Stakes combine to generate this implicature, Rysiew leaves it unclear that this implicature is really present.¹⁶³

Rysiew does try to meet DeRose's challenge, but his response focuses solely on the possibility that the bank has recently discontinued its Saturday hours. It makes no mention of Rachel and Keith's *practical circumstances* in High Stakes. As Jessica Brown points out, however, these practical circumstances matter. They affect our intuitions about Keith's

¹⁶³ DeRose (2002), p. 176.

utterance of $\neg K$ in High Stakes.¹⁶⁴ Brown develops Rysiew's account to capture this impact on our intuitions. To do this, she starts with Grice's Maxim of Relation and his well-known example of the motorist who tells the pedestrian that she is running out of gas.¹⁶⁵ In Grice's example, the pedestrian utters 'there is a garage nearby' and thereby implicates that there is an *open* garage nearby. The scenario that Brown imagines differs slightly from Grice's. In Brown's scenario, there is only one garage nearby and the pedestrian knows that this garage is *closed*, so she utters 'there is no garage nearby.' Brown thinks that, just as the pedestrian in Grice's petrol example implicates that there is an open garage nearby by uttering just 'there is a garage nearby,' the pedestrian in her example implicates that there *is not* an open garage nearby by uttering just, 'there is no garage nearby.' According to Brown, "[w]hile this utterance is literally false, it pragmatically conveys the true claim that there is no open garage nearby."¹⁶⁶ Brown thinks that an utterance of K or $\neg K$ in High Stakes would, respectively, be analogous to an utterance of 'there is a garage nearby' or 'there is no garage nearby' in her petrol scenario.¹⁶⁷ In High Stakes, Keith would have implicated a false proposition by uttering K . By uttering $\neg K$ instead, he thereby implicated a *true* proposition. Which proposition, exactly? Brown thinks that, had Keith uttered K in High Stakes, he would have implicated that his belief that the bank will be open matches the facts out to the nearest world in which the bank has recently changed its Saturday hours.¹⁶⁸ Let '*m*' name the

¹⁶⁴ Brown (2006).

¹⁶⁵ Grice (1989), p. 32.

¹⁶⁶ Brown (2006), p. 425.

¹⁶⁷ *Ibid.*, p. 426.

¹⁶⁸ Brown is apparently taking it for granted that beliefs are not individuated by their contents. According to her DeRose-inspired account of *matching the facts*, Keith's belief that the bank will be open matches the facts out to the nearest world in which the bank has recently changed its Saturday hours only if, in the nearest world in which the bank has recently changed its Saturday hours, Keith *doesn't* believe that the bank will be open on Saturday. See *ibid.*, p. 424.

proposition that Keith's belief that the bank will be open matches the facts out to the nearest world in which the bank has recently changed its Saturday hours. Since Keith uttered $\neg K$ instead of K , says Brown, he implicated that $\neg m$ instead of m . And since $\neg m$ is *true*, he implicated a *true* proposition by uttering $\neg K$ in High Stakes. Finally, says Brown, when we read High Stakes and consider Keith's utterance of $\neg K$, his utterance rings true because it conversationally implicates that $\neg m$, and we mistake $\neg m$ for what Keith *says* by uttering $\neg K$.¹⁶⁹ But now that we've explained why Keith seems right to utter $\neg K$ in High Stakes, we are free to deny that Keith says something true by uttering $\neg K$ in High Stakes, says Brown.

2 A Worry about the Generality of Brown's WAM

Brown's development of Rysiew's WAM does seem like an improvement, since it makes some progress toward explaining how general conversational principles combine with the semantics of the word 'knows' and the details of High Stakes to generate the supposed implicature. There is an obvious problem with Brown's WAM, however. Brown is correct that Rachel and Keith's practical circumstances and concerns affect our intuitions about Keith's utterance of $\neg K$. Consider the following revision of High Stakes, which focuses *solely* on their practical circumstances and concerns. Unlike the original version of High Stakes, this version does not mention the possibility that their bank has changed its Saturday hours since Keith's last visit.

Practical High Stakes: Rachel and Keith are driving home on Friday afternoon.

They plan to stop at the bank to deposit their paychecks, but as they approach it they

¹⁶⁹ Ibid.

notice that the lines inside are very long. Keith suggests that they go home and return to deposit their paychecks on Saturday morning, explaining that he stops at the bank every Saturday morning and that it's open until noon. But in this case, Rachel and Keith have just written a very large and important check. If their paychecks are not deposited before Monday morning, the important check they wrote will bounce, leaving them in a terrible situation. The bank isn't open on Sunday, and Rachel reminds Keith of this fact. She then says, "It will be completely disastrous for us if we go home now and the bank isn't open tomorrow. Do you know that it will be open tomorrow?" Remaining just as confident as he was before that the bank will be open tomorrow, Keith replies, "Well, no, I don't know that the bank will be open tomorrow. We had better stop and deposit the checks now." (As it happens, the possibility that the bank has *changed its Saturday hours since Keith's last visit* never enters Keith or Rachel's mind.¹⁷⁰)

Just as Keith appears to say something true by uttering $\neg K$ in the original version of High Stakes, he appears to say something true by uttering $\neg K$ in this version of High Stakes. But the explanation for this appearance is certainly not that, by uttering $\neg K$, Keith implicates *that his belief does not match the facts out to the nearest world in which the bank has changed its Saturday hours since his last visit*. After all, in *this* version of High Stakes, we have stipulated that the possibility that the bank has changed its Saturday hours since Keith's last visit is completely out of mind. So it is implausible that, by uttering $\neg K$ in *this* version High Stakes, Keith implicates

¹⁷⁰ Notice that, just as Rachel and Keith can consider the possibility that the bank has recently changed its Saturday hours without having in mind any specific reason why the bank might have recently changed its Saturday hours, Rachel and Keith can consider the possibility that the bank will not be open on Saturday without considering any specific reason why the bank might not be open on Saturday.

that his belief does not match the facts out to the nearest world in which the bank has recently changed its Saturday hours. Since this version of High Stakes causes just as much trouble for minimalism as the *original* version of High Stakes, there is little point in responding to the original version of High Stakes with a WAM unless we can *also* respond to this version of High Stakes with a WAM.

So, is there any value of ' p ' where the following claim is plausible?

By uttering $\neg K$ in Practical High Stakes, Keith implicates that p . Keith seems to say something true by uttering $\neg K$ in Practical High Stakes because p is true and we mistake p for what Keith *says* by uttering $\neg K$ in Practical High Stakes.

We get our answer by paying close attention to the decision-theoretic nature of the *original* version of High Stakes—a feature that is preserved in Practical High Stakes, and which is present in *all* of the high-stakes cases in the literature (Cohen's airport case, Fantl and McGrath's train case, Weatherson's genie case, Ross and Schroeder's sandwich case, and so on).¹⁷¹

3 A WAM with General Application

In the *original* version of High Stakes, Rachel and Keith are deciding whether they should wait in line to deposit their checks Friday evening, or go straight home and deposit them Saturday morning. Deciding whether they should wait in line or go straight home is the sole purpose of their conversation. They would both prefer to go straight home conditional on

¹⁷¹ See Cohen (1988), p. 58, Fantl and McGrath (2002), p. 67-8, Weatherson (forthcoming), pp. 10-11 of version available on Weatherson's website, and Ross and Schroeder (forthcoming), p. 3 of version available on Schroeder's website.

the bank's being open on Saturday, and they both know this. But they also both know that it would be too risky for them to go straight home without very strong evidence that the bank *will* be open on Saturday. They know that the bank has been open on many previous Saturdays, and, in this context, Rachel raises the possibility that the bank has just discontinued its Saturday hours. It is clear enough *why* Rachel raises this possibility: she thinks it would be too risky for them to go straight home unless they can rule it out. Rachel knows that *her* evidence that the bank will be open does not rule this possibility out, so she needs to know whether *Keith's* evidence rules it out. But instead of just asking Keith whether he has sufficient evidence to rule out the possibility that their bank has changed its Saturday hours since his last visit, she asks Keith *whether he knows that the bank will be open on Saturday*. Why does Rachel ask Keith *this* question? Why doesn't she just ask Keith whether he can rule out the possibility that their bank has recently changed its Saturday hours?

Plausibly, Rachel doesn't ask Keith whether he can rule out the possibility that their bank has recently changed its Saturday hours because she recognizes that an answer to this question would only give her *part* of the information that she needs. There are *many* possible worlds where the bank does not open this Saturday, and the world where the bank has recently discontinued its Saturday hours is only one of them. Rachel needs to know whether Keith can rule out *all* of these worlds—or, at least, whether he can rule out all such worlds *that are roughly as likely to be actual as the world where bank has recently changed its Saturday hours*. Plausibly, Rachel does not ask Keith whether he can rule out the possibility that their bank has recently changed its Saturday hours because she knows that, even if he *can* rule this possibility out, there might be other equally worrisome possibilities that he *cannot* rule out. So

we can explain why Rachel *does not* ask Keith whether he can rule out the possibility that the bank has recently changed its Saturday hours.

But still, why does Rachel ask Keith what she *does* ask him: *whether Keith knows that the bank will be open?* Why does she ask *this* question? After all, if Rachel doesn't ask Keith whether he can rule out the possibility that their bank has recently changed its Saturday hours because she is thinking about other, equally worrisome possibilities, then why doesn't she use the possibility that bank has changed its Saturday hours since Keith's last visit *as an example* and then ask Keith whether he can rule out this possibility *and other equally worrisome possibilities?*

In context, the answer to this question is clear enough, and it has everything to do with the sole purpose of Rachel and Keith's conversation: to decide whether they should wait in line to deposit the checks Friday evening or, instead, go straight home and return to deposit them on Saturday morning. Compare High Stakes to the paradigm case of conversational implicature that we discussed above: the pedestrian's implicating that there is an open garage nearby by uttering just 'there is a garage nearby.' As we noted, Brown discusses a case slightly different from Grice's. In Brown's case, the motorist tells the pedestrian that she is running out of gas, and the pedestrian immediately responds by uttering 'there is no garage nearby' instead of 'there is a garage nearby.' So imagine a case that differs from *Brown's* case, as follows.

Motorist: I'm running out of gas. Is there a garage nearby?

Pedestrian: No, there isn't a garage nearby. The nearest one's in Sleetmute.

In this case, the motorist says that she is running out of gas and then asks, not whether there is an open garage nearby, but just whether there is a garage nearby. There is an obvious symmetry between the motorist's question in this case and Rachel's question in High Stakes. The motorist's sole purpose in addressing the pedestrian is to find out where she can get gas, and it is transparent what the motorist means to accomplish by asking the pedestrian whether there is a garage nearby: she means to find out whether there is a garage nearby *where she can get gas*. Likewise, Rachel's sole purpose in questioning Keith is to determine the best course of action, and it is transparent what Rachel means to accomplish by asking Keith if he knows that the bank will be open: she means to find out whether Keith has knowledge *that they can reasonably act on* ('actionable knowledge,' as I will sometimes say).¹⁷² But given this symmetry, the claim that Keith implicates that he lacks *actionable* knowledge by uttering $\neg K$ is just as plausible as Brown's claim that the pedestrian implicates that there is no *open* garage nearby by uttering 'there is no garage nearby.' If Brown is right that the pedestrian implicates that there is no open garage nearby by uttering 'there is no garage nearby,' then it would be surprising if Keith did not implicate that he lacks actionable knowledge by uttering $\neg K$. So Brown should think that, in the original version of High Stakes, Keith implicates that he lacks *actionable* knowledge by uttering $\neg K$.

Now we said that there was little point in responding to the original version of High Stakes with a WAM if there was no way to apply that WAM to Practical High Stakes. But, of course, the suggestion that Keith implicates that he lacks actionable knowledge by uttering $\neg K$ is only *more* plausible with respect to Practical High Stakes. So, noting the implausibility

¹⁷² I owe this term to correspondence with Ernest Sosa. Note that, according to pragmatism, *all* knowledge is actionable knowledge.

of the suggestion that Keith implicates that his belief does not match the facts out to the nearest world in which the bank has changed its Saturday hours since his last visit by uttering $\neg K$ in Practical High Stakes, and noting the comparative plausibility of the suggestion that Keith implicates that he lacks actionable knowledge by uttering $\neg K$ in both the original version of High Stakes *and* Practical High Stakes, Brown should think that the following WAM will suffice in response to High Stakes if *any* WAM will suffice.

By uttering $\neg K$ in High Stakes, Keith implicates that he lacks actionable knowledge; Keith seems to say something true by uttering $\neg K$ in High Stakes because we mistake this true thing that he implicates by uttering $\neg K$ in High Stakes for what he *says* by uttering $\neg K$ in High Stakes.

4 Conversational Implicatures and Cancelability

As we noted above, when minimalists say that Keith implicates something true in High Stakes, and when they say that we mistake what Keith implicates in High Stakes for what he *says* in High Stakes, they are talking about *conversational* implicature. But conversational implicatures are *cancelable*, and this property of conversational implicatures will cause trouble for minimalists who respond to the limits of Brown's WAM with the actionable knowledge WAM above.¹⁷³

¹⁷³ Ibid., p. 39 As Grice puts it, "A putative conversational implicature that *p* is explicitly cancelable if, to the form of words the utterance of which putatively implicates that *p*, it is admissible to add 'but not *p*,' or 'I do not mean to imply that *p*,' and it is contextually cancelable if one can find situations in which the utterance of the form of words would simply not carry the implicature. Now I think that all conversational implicatures are cancelable." See Grice (1980), p. 44. Throughout, I will use 'implicate' as short of 'conversationally implicature,' and 'implicature' as short for 'conversational implicature.'

Again, consider the conversation Grice gives us in his well-known petrol example.¹⁷⁴

Motorist: I'm running out of gas. Is there a garage nearby?

Pedestrian: Yeah, there is a garage nearby—right around the corner.

Here the pedestrian utters 'there is a garage nearby' and thereby implicates that there is an *open* garage nearby. Had the garage been *closed*, however, the pedestrian could have *cancelled* this implicature by adding 'but it is closed.' That is, had the pedestrian uttered 'yeah, there is a garage nearby—but it is closed,' she would have successfully conveyed to the motorist that there is a garage nearby *without* conveying to the motorist that there is an *open* garage nearby.

Now, as we saw above, Brown takes it for granted that, because the pedestrian implicates that there is an open garage nearby by uttering just 'there is a garage nearby,' the pedestrian would have implicated that there is *no* open garage nearby had she uttered just 'there is no garage nearby.' Grice would reject Brown's claim that the pedestrian would have implicated that there is no open garage nearby had she uttered just 'there is no garage nearby,' since the proposition that there is no open garage nearby fails the cancelability test. Consider the following answers that the pedestrian might have given to the motorist.

A1: No, there's no garage nearby, but there is an open garage nearby.

A2: No, there's no garage nearby, but I am not saying that there is no open garage nearby.

¹⁷⁴ Ibid., p. 32

When the pedestrian utters ‘there is no garage nearby,’ she thereby communicates to the motorist that there is no *open* garage nearby. But by Grice’s lights, the pedestrian does not *implicate* that there is no open garage nearby, since, downstream from an utterance of ‘there is no garage nearby,’ the proposition that there is no *open* garage nearby is not *cancellable*. As A1 and A2 show us, once the pedestrian has uttered ‘there is no garage nearby,’ it is not (in Grice’s words) admissible for the pedestrian to add that there *is* an open garage nearby, or that she is *not* saying that there is no open garage nearby, or anything like that.¹⁷⁵ By attempting to cancel the proposition that there is no open garage nearby, the pedestrian makes nonsense of her answer to the motorist.¹⁷⁶ A similar worry will threaten the actionable knowledge WAM above.

5 The Limits of WAMs

In High Stakes, Rachel raises the possibility that their bank has recently changed its Saturday hours and then, without pause, asks Keith whether he knows that the bank will be open. Keith responds by uttering $\neg K$, and, according to the WAM that we are forwarding as the most promising WAM available, Keith thereby implicates that he lacks knowledge *that they can act on* that the bank will be open. But the proposition that Keith lacks knowledge that they can act on that the bank will be open is no more a cancellable implicatum of Keith’s utterance of $\neg K$ than the proposition that there is no open garage nearby is a cancellable implicatum of the pedestrian’s utterance of ‘there is no garage nearby,’ or the proposition

¹⁷⁵ Grice (1989), p. 44. Though I do not have space to defend the claim here, I believe that Brown’s claim that the pedestrian implicates that there is no open garage nearby by uttering ‘there is no garage nearby’ also fails Grice’s *calculability* test. (See *ibid.*, p. 39.)

¹⁷⁶ Some of Kent Bach’s comments in “The Top 10 Misconceptions about Implicature” suggest that Bach does not agree with Grice that conversational implicatures are always cancellable. See Bach (2006), p. 24.

that somebody has long-jumped over 28 feet is a cancellable implicatum of your utterance of ‘Bob Beamon long-jumped over 29 feet.’ Consider the following way that High Stakes might have ended.

Rachel: What if our bank discontinued its Saturday hours since your last visit? Do you know that it will be open?

Keith: No, I don’t know that the bank will be open, but I do have knowledge that the bank will be open that we can act on.

What could Keith be telling Rachel here? Just as the pedestrian could not utter ‘there is no garage nearby’ and then admissibly add that there is an open garage nearby, and just as you could not utter ‘Bob Beamon long-jumped over 29 feet’ and then admissibly add that nobody long-jumped over 28 feet, Keith cannot utter $\neg K$ and then admissibly add that he has actionable knowledge that the bank will be open. With respect to Keith’s utterance of $\neg K$, the proposition that Keith lacks actionable knowledge that the bank will be open performs miserably on Grice’s cancelability test.

Now Brown anticipates worries about Grice’s cancelability test in response to her claim that, by uttering $\neg K$ in High Stakes, Keith implicates that his belief that the bank will be open does not match the facts out to the nearest world in which the bank has changed its hours since his last visit. Here is what Brown says.

On the proposed account, ‘*S* knows that *p*’ may pragmatically convey that *S* is in a very strong epistemic position, that her belief matches the facts across a wide range

of worlds, including some so far away they are not normally taken to undermine knowledge. Contrary to the objection, it may be possible to cancel this implication. To my ear, the claim that ‘*S* knows that *p*, but *S* is not in a really strong epistemic position with respect to *p*’ or ‘*S* knows that *p*, but her belief wouldn’t match the facts in a really distant possible world’ do not seem obviously inconsistent or uncomfortable.¹⁷⁷

I have some sympathies for what Brown says here. But nothing she says here it seems apposite to *our* worries about cancelability, since Brown only addresses *positive* utterances of the form ‘I know that *p*.’ She says nothing about *negative* utterances of the form ‘I *do not* know that *p*,’ which are what we are worrying about. Even if Keith would implicate that *p* (for some value of ‘*p*’) by uttering ‘I know that the bank will be open’ in High Stakes, it does not follow that Keith would implicate that $\neg p$ (for that same value of ‘*p*’) by uttering ‘I *do not* know that the bank will be open’ in High Stakes. So we can easily agree with everything that Brown says above and continue to maintain that Keith *does not* implicate that he lacks actionable knowledge by uttering $\neg K$ in High Stakes. This would be no more difficult than agreeing with everything that Brown says above while also agreeing with Grice that the pedestrian in his scenario would not implicate that there is no open garage nearby by uttering ‘there is no garage nearby.’

So we have good reason to doubt that, by uttering $\neg K$, Keith implicates that he lacks actionable knowledge. This means that the actionable knowledge WAM that we forwarded as an improvement over Brown’s WAM will not work. There is one more WAM to consider,

¹⁷⁷ Brown (2006), p. 428.

however. Again, return to Grice's petrol case. In that case, the pedestrian utters 'there is a garage nearby' and thereby implicates that there is an open garage nearby. What explains this implicature? According to Grice's Cooperative Principle, the pedestrian must make her conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the conversation in which she is engaged.¹⁷⁸ According to Grice's Maxim of Relation, the pedestrian must make her conversational contribution *relevant*.¹⁷⁹ The pedestrian does not seem to be flouting the Maxim of Relation when she tells the motorist that there is a garage nearby, and we have no reason to think that she is being uncooperative by telling the motorist just that there is a garage nearby. Yet the pedestrian must regard her own utterance as irrelevant if she does not think that the nearby garage is *open*. After all, the motorist obviously needs to buy gas, and the pedestrian (we may presume) knows that a closed garage is useless with respect to these needs. So preserving the assumption that the pedestrian is observing the Cooperative Principle requires assuming that the pedestrian believes that the nearby garage is open. *This* (according to Grice) is why the pedestrian implicates that there is an open garage nearby by uttering just 'there is a garage nearby.'¹⁸⁰

If this is why the pedestrian implicates that there is an open garage nearby by uttering 'there is a garage nearby,' however, then it seems that the pedestrian also implicates *that the motorist can get gas somewhere nearby* by uttering 'there is a garage nearby.' Just as the pedestrian must regard her utterance as irrelevant if she does not think that the nearby garage is open, she must regard her utterance as irrelevant if she does not think that the motorist can get gas

¹⁷⁸ Grice (1989), p. 26.

¹⁷⁹ *Ibid.*, p. 27.

¹⁸⁰ *Ibid.*, p. 32.

somewhere nearby. After all, the pedestrian knows that the motorist is running out of gas. Preserving the assumption that the pedestrian is observing the Cooperative Principle requires assuming *both* that the pedestrian believes that the nearby garage is open and that the pedestrian believes that the motorist can get gas somewhere nearby. So, by uttering ‘there is a garage nearby,’ the pedestrian implicates both that there is an open garage nearby and that the motorist can get gas somewhere nearby.

Now take another look at the conversation where the pedestrian gives a *negative* answer to the motorist’s question. Here it is, again.

Motorist: I’m running out of gas. Is there a garage nearby?

Pedestrian: No, there’s no garage nearby. The nearest one’s in Sleetmute.

By uttering ‘there is no garage nearby,’ the pedestrian does not implicate that there is no open garage nearby, as we saw above. Plausibly, however, she *does* implicate that the motorist *cannot* get gas anywhere nearby. First, just as preserving the assumption that the pedestrian is observing the Cooperative Principle in Grice’s original petrol case requires assuming that the pedestrian believes that the nearby garage is open, preserving the assumption that the pedestrian is observing the Cooperative Principle requires assuming that the pedestrian believes that the motorist cannot get gas anywhere nearby. Second, the suggestion that the pedestrian implicates that the motorist cannot get gas anywhere nearby easily passes the cancelability test. Consider the following conversation.

Motorist: I'm running out of gas. Is there a garage nearby?

Pedestrian: No, there's no garage nearby, but there *is* a place nearby where you can get gas. They sell it out at the airfield.

It would be perfectly admissible for the pedestrian to utter 'there is no garage nearby' and then add that there *is* a place nearby where the motorist can get gas. By uttering 'there is no garage nearby, but there is a place nearby where you can get gas,' the pedestrian successfully conveys to the motorist that there is no garage nearby *without* conveying to the motorist she cannot get gas anywhere nearby. So the proposition that the motorist cannot get gas anywhere nearby is a cancellable implicatum of the pedestrian's utterance of 'there is no garage nearby.'

Now consider the suggestion that, by uttering $\neg K$, Keith implicates *that Rachel and Keith cannot reasonably go straight home*. As we noted above, Rachel's sole purpose in questioning Keith is to determine the best course of action, and it is transparent what Rachel means to accomplish by asking Keith if he knows that the bank will be open. She means to find out whether they can reasonably go straight home. So minimalists can point out that, just as preserving the assumption that the pedestrian is observing the Cooperative Principle in Grice's original petrol case requires assuming that the pedestrian believes that the nearby garage is open, preserving the assumption that Keith is observing the Cooperative Principle in High Stakes requires assuming that Keith believes that he and Rachel cannot reasonably go straight home. This similarity motivates the claim that, by uttering $\neg K$ in High Stakes, Keith implicates that Rachel and Keith cannot reasonably go straight home.

Now, here is the question: will a WAM built on this supposed implicature pass the cancelability test? Consider the following alternative ending to High Stakes.

... If their paychecks are not deposited before Monday morning, the important check they wrote will bounce, leaving them in a terrible situation. The bank isn't open on Sunday. Rachel reminds Keith of these facts. She then says, "What if our bank discontinued its Saturday hours since your last visit? Do you know that it will be open?" Remaining just as confident as he was before that the bank will be open, Keith replies, "Well, no, I don't know that the bank will be open, but we can still reasonably go straight home."

Is it admissible for Keith to utter 'I don't know that the bank will be open' to Rachel and then add that they can still reasonably go straight home? By uttering 'I don't know that the bank will be open, but we can still reasonably go straight home,' does Keith successfully convey to Rachel that he does not know that the bank will be open *without* conveying to her that they *cannot* reasonably go straight home? I'm not sure. That Keith does not know that the bank will be open does not entail that Rachel and Keith cannot rationally go straight home, as it *does* entail that Keith lacks actionable knowledge that the bank will be open. So perhaps Keith *can* utter $\neg K$ and then cancel the proposition that Rachel and Keith cannot rationally go straight home. Then again, given the details of High Stakes, it is hard to imagine what reason Keith could have for thinking *both* that he does not know that the bank will be open *and* that Rachel and Keith can reasonably go straight home. I have a hard time seeing what Keith could be telling Rachel here. So I am inclined to say that, by attempting to cancel

the proposition that Rachel and Keith cannot reasonably go straight home, Keith would make nonsense of his answer to Rachel's question. As a consequence, I am inclined to deny that, by uttering $\neg K$, Keith implicates that he and Rachel cannot reasonably go straight home. Given this, however, premise (3) of the propriety argument looks exactly right, and minimalists seem mistaken to rely on WAMs in response to High Stakes.¹⁸¹

6. The Propriety of Keith's Utterance in High Stakes

We are taking it for granted that Keith's utterance of $\neg K$ *is* appropriate, so, if we want to defend minimalism from the propriety argument, we must now reject premise (2). We must claim that the explanation of the propriety of Keith's utterance of $\neg K$ is neither that he says anything true by uttering $\neg K$ nor that he implicates anything true by uttering $\neg K$. But if this

¹⁸¹ Brown worries that the cancellability test causes problems for her WAM and says two things in response. First, she says that "the claim that '*S* knows that *p*, but *S* is not in a really strong epistemic position with respect to *p*' or '*S* knows that *p*, but her belief wouldn't match the facts in a really distant possible world' do not seem obviously inconsistent or uncomfortable." Second, Brown says that, even if these claims *are* uncomfortable, "Grice himself accepted that some implicatures cannot be 'comfortably' cancelled. Indeed, we should expect an uncomfortable cancellation in a case where speakers tend to confuse what's literally said by an utterance with what it pragmatically conveys." See Brown (2006), p. 428.

I doubt that either of these comments adequately addresses my claim that the proposition that Keith's epistemic position is not strong enough for him to rule out such possibilities as the possibility that the bank has recently changed its Saturday hours is not a cancellable implicature of Keith's utterance of $\neg K$. First, '*S* knows that *p*, but *S* is not in a really strong epistemic position with respect to *p*' and '*S* knows that *p*, but her belief wouldn't match the facts in a really distant possible world' are not actual sentences—they are sentence *schemas*. But the sentence that we have been considering—namely, 'I don't know that the bank will be open tomorrow, but my epistemic position *is* strong enough for me to rule out such possibilities as the possibility that the bank has recently changed its Saturday hours'—does not fit either schema. So it is unclear how Brown's first point could even apply. Second, while Keith's utterance of 'I don't know that the bank will be open tomorrow, but my epistemic position *is* strong enough for me to rule out such possibilities as the possibility that the bank has recently changed its Saturday hours' is tangibly uncomfortable (to use Brown's words), the proper response to this fact is surely not that "we should expect an uncomfortable cancellation in a case where speakers tend to confuse what's literally said by an utterance with what it pragmatically conveys." As we saw above, Keith's utterance is uncomfortable because in High Stakes it is clear that either Keith *does* know that the bank will be open tomorrow, or his epistemic position with respect to the proposition that the bank will be open *is not* strong enough for him to rule out (e.g.) the possibility that the bank has recently changed its Saturday hours. That some speakers might be confused about what Keith does and does not literally say by uttering $\neg K$ seems beside the point.

is right, then what *does* explain the propriety of Keith's utterance of $\neg K$? We get a clue by considering cases like the following.

Ester and Moonflower are freshmen roommates with different views about health and medicine. Moonflower thinks she can cure any ailment with a concoction she makes in her blender. Ester thinks the concoction is bogus but she hasn't told Moonflower this. One day, when Ester is wearing sunglasses indoors, Moonflower asks Ester if she has a headache. Ester does have a headache, but she doesn't want Moonflower to make one of her silly concoctions, and she knows that Moonflower will make one of her concoctions if she admits that she has a headache. Ester's options are (a) telling Moonflower that she doesn't have a headache, (b) telling Moonflower that she has a headache and drinking one of her concoctions, and (c) telling Moonflower that she has a headache and trying to persuade Moonflower that her concoctions don't work. Ester weighs the costs and benefits of the various options and picks (a). By Ester's lights it's better to have Moonflower think she doesn't have a headache than to either drink one of Moonflower's concoctions or get into a complicated discussion (maybe even an argument) with Moonflower about health and medicine.

We can fill in the details of this case so that, everything considered, it is appropriate for Ester to utter 'I don't have a headache,' even though Ester neither says nor implicates anything true with this utterance. But similar considerations explain the propriety of Keith's utterance of $\neg K$ in High Stakes.

In High Stakes it is clear that, if Keith does not know that the bank will be open, this is because he cannot rule out the possibility that their bank has recently changed its Saturday hours. This is why Rachel raises the possibility that their bank has recently changed its Saturday hours and then, instead of asking Keith if he can rule this possibility out, asks Keith whether he knows that the bank will be open. So now consider two possibilities: first, that Keith is a minimalist, and second, that he is not a minimalist.

Suppose first that Keith is a minimalist. In this case, in response to Rachel's question whether Keith knows that the bank will be open, Keith has three options: (a) he can tell Rachel that he does not know that the bank will be open, (b) he can tell Rachel that he *does* know that the bank will be open and then go inside to make sure that it has not recently changed its Saturday hours, or (c) he can tell Rachel that he knows the bank will be open and then *skip* making sure that it has not recently changed its Saturday hours.¹⁸² The expected costs and benefits of these options make (a) the obvious choice. Keith obviously should *not* skip going inside to make sure that the bank has not recently changed its Saturday hours, so option (c) is out. And (b) is just too complicated, given their practical circumstances. Keith would have to explain himself to Rachel if he were to pursue this option, and offering such an explanation would require getting into the details of minimalists epistemology. Since this is neither the time nor place to get into the details of minimalist epistemology, option (b) is out. But (a) is the only option left, and the only cost of this option is a little white lie. So (a) is the best of Keith's options, and *this* is why it would be appropriate for Keith to utter $\neg K$, if he were a minimalist.

¹⁸² There are other options, of course, but (a), (b) and (c) are sufficiently representative of them.

Assume that Keith *is not* a minimalist then. Now it is only *more* obvious that uttering $\neg K$ is the best of Keith's options. If Keith is not a minimalist, then he has all of the above reasons for uttering $\neg K$, plus an additional one: now, by his own lights, he might be saying something *true* by uttering $\neg K$. Thus, whether Keith is a minimalist or not, the expected costs and benefits of his options make uttering $\neg K$ his obvious best choice. Explaining the propriety of Keith's utterance of $\neg K$ in High Stakes requires neither the thesis that Keith says something true by uttering $\neg K$ nor the thesis that he implicates something true by uttering $\neg K$. This shows that (2) is false, and the propriety argument fails.

7. Conclusion

In this paper, I have offered advice to minimalists. I *have not* challenged contextualism or pragmatism. Instead, I have raised problems for minimalism from *within* a minimalist framework and then shown how minimalists should respond to these problems. Minimalists should ditch the WAMs they have adopted in response to High Stakes and then respond to the propriety of Keith's utterance by appealing to mundane features of practical rationality. Following my advice here will result in a minimalism that is less susceptible to a criticism that contextualists and pragmatists are bound to make.

Chapter 6: The Nature of Belief and the Value of Knowledge*

Belief and confidence are intimately related, and these, in turn, bear some tight connection to behavior—specifically *assertion*. But how exactly *are* belief, confidence and assertion related to each other? Suppose *A* and *B* own one ticket in a 100-ticket lottery, and suppose *A* and *B* have no evidence that their ticket lost other than the odds. We can add details to this case to elicit different intuitions about their respective beliefs. Does confidence fix belief? What if they are both 99% confident but only *B* will assert flat-out that the ticket lost? What if *A* and *B* will both assert flat-out that their ticket lost, but *A* is considerably less confident than *B*? Do these differences entail a difference in belief?

The relation between belief, confidence and assertion is interesting in its own right, but it gains special importance when we see how questions in epistemology and psychology turn on it. How should we regulate our beliefs—if indeed we *can* regulate them? If practical matters bear on the evaluation of belief *vis-à-vis* a tight connection between belief and assertion, do practical considerations encroach on epistemological properties like justification and knowledge? How can we measure the effects of some factor (culture, say) on belief? Well, what signals a difference in belief? These questions are the tip of the iceberg.

The relationship between belief, confidence and assertion also bears on the value of knowledge. Distinguish the kind of ignorance that results when your true belief falls short of knowledge from the kind of ignorance that results when you either hold no belief or hold a

* Thanks to Lisa Miracchi and Ernest Sosa for helpful comments on an earlier draft of this paper.

false belief. Call the first variety of ignorance ‘informed ignorance,’ and call the latter variety of ignorance ‘total ignorance.’ Knowledge entails belief and, on certain conceptions of the relationship between belief, confidence and assertion, the difference between holding a belief and holding *no* belief can be utterly trivial. So, on certain conceptions of the relationship between belief, confidence and assertion, the difference between knowledge and total ignorance can be utterly trivial. But in this case, it seems that, not only is knowledge not distinctively valuable, knowledge needn’t even be better than total ignorance.¹⁸³ In this paper, I say how belief, confidence and assertion relate to each other, and how these relations bear on the value of knowledge. Knowing that *p* will always be epistemically better than total ignorance *vis-à-vis p*, I will argue, but informed ignorance might be epistemically better than nearby ways of knowing—in which case knowledge is not distinctively valuable. To deliver these results, I defend an “affirmative” view of belief similar to the view that Ernest Sosa has recently defended.¹⁸⁴ Along the way, however, I reject Sosa’s argument for the affirmative view and show how the affirmative view actually *prevents* us from endorsing Sosa’s conclusion that knowledge is distinctively valuable. I also reject an assumption that is ubiquitous in contemporary epistemology: that the belief condition on knowledge satisfies whatever confidence condition there might be on knowledge. Knowledge, I argue, entails considerably higher confidence than is entailed by belief. I conclude by relating the affirmative view of belief to a straightforward argument for the threshold view of belief and

¹⁸³ Knowledge is distinctively valuable just in case, for any value of ‘*p*’ where *p* is true, it is epistemically better to know that *p* than to believe that *p* without knowing that *p*.

¹⁸⁴ Ernest Sosa, “Value Matters in Epistemology,” *The Journal of Philosophy*, Vol. 107, No. 4 (2010): 167-90. Reprinted in Ernest Sosa, *Knowing Full Well* (Princeton: Princeton University Press, 2011), ch. 3.

then showing how the considerations at work in that argument for the threshold view actually lend more support to the affirmative view.

1. The Threshold View

As we just said, belief and confidence are intimately related. According to the threshold view, belief *just is* confidence—confidence above some universal threshold. For some level of confidence c , confidence at or above c is both necessary and sufficient for belief. The threshold view is easy to motivate. Let $\text{Pr}(p)$ be S 's confidence that p , and suppose $\text{Pr}(p) = 1$ just in case S is absolutely certain that p , $\text{Pr}(p) = 0$ just in case S is absolutely certain that $\neg p$, $\text{Pr}(p) = 0.5$ just in case S 's confidence that p equals her confidence that $\neg p$, and so on for the rest of the values in the unit interval. Now consider (1) and (2).

(1) If $\text{Pr}(p) = 1$ at t , then S believes that p at t .

(2) If S believes that p at t , then $\text{Pr}(p) \geq 0.5$ at t .

Plausibly, you can't be absolutely certain that p without believing that p , so (1) seems true. And plausibly, you can't believe that p if your confidence that it is *not* the case that p exceeds your confidence that it *is* the case that p , so (2) seems true. But (1) is true only if absolute certainty suffices for belief, and (2) is true only if confidence at least as great as 0.5 is necessary for belief. So (1) and (2) are true only if some level of confidence entails belief and some level of confidence is necessary for belief. Thus, (1) and (2) give us (3) and (4), respectively.

$$(3) \quad \exists x \forall S p t \square (\text{Pr}(p) = x \text{ at } t \supset S \text{ believes that } p \text{ at } t)$$

$$(4) \quad \exists x \forall S p t \square (S \text{ believes that } p \text{ at } t \supset \text{Pr}(p) \geq x \text{ at } t)$$

But (3) entails (5), and, conjoining (4) and (5), it's easy to argue for (6).

$$(5) \quad \exists x \forall S p t \square (\text{Pr}(p) \geq x \text{ at } t \supset S \text{ believes that } p \text{ at } t)$$

$$(6) \quad \exists x \forall S p t \square (S \text{ believes that } p \text{ at } t \equiv \text{Pr}(p) \geq x \text{ at } t)$$

The argument for (6) goes as follows. Let m be the lowest value of 'x' for which (5) is true, and let n be the lowest value of 'x' for which (4) is true. Then one of three things is the case: either m minus n is less than zero, m minus n is greater than zero, or m minus n equals zero. If m minus n is less than zero, we can derive contradictions from (4) and (5). Suppose m is 0.8 and n is 0.9. In this case, (5) says S can't be 0.8 confident without believing, while (4) says S can't believe without being at least 0.9 confident. So if S is 0.85 confident, (5) says S believes while (4) says she doesn't. This is an unacceptable result. Suppose that m minus n is greater than zero, then. In this case we get a mystery. By hypothesis, confidence as high as m entails that S believes, and confidence below n entails that S doesn't believe. Since m minus n is greater than zero, m is greater than n . So there is a gap between m and n where S 's level of confidence leaves it an open question whether S believes. But isn't this a strange view? If confidence in the interval $n \times m$ leaves it open whether S believes, why should confidence *outside* this interval fully determine whether S believes? What factor, that can't make a difference outside this interval, can make a difference *inside* this interval? Why isn't

confidence the whole story about belief inside this interval, given that it's the whole story about belief outside this interval? What is so special about this interval? Without answers to these questions, we should be skeptical that m minus n is greater than zero.¹⁸⁵ But if we deny that m minus n is greater than zero, then we have to agree that m minus n equals zero, since this is the only option left. Yet if m minus n equals zero, then m equals n ; and if m equals n , then the lowest level of confidence that entails belief is also the lowest level of confidence that is entailed by belief. So, if m equals n , there is some level of confidence that is both the lowest level of confidence that entails belief and the lowest level of confidence that is entailed by belief. This is exactly what (6) says. So, we have some motivation for accepting (6). Thus, the threshold view of belief is easy to motivate. Or so we might argue.

2. The Affirmative View of Belief

Ernest Sosa rejects the threshold view.¹⁸⁶ According to Sosa, the threshold view conflicts with the platitudes below, and this conflict suffices for rejecting the threshold view.

THE VALUE PLATTITUDES

- (KA) If one takes up a question, it is epistemically better to know the answer than not to know it. More specifically: One's conscious answer to the question is epistemically better than one's conscious suspension of judgment, provided one's answer constitutes knowledge.

¹⁸⁵ If you think you *have* satisfying answers to these questions, you are probably right. But hold your doubts about this argument for §8.

¹⁸⁶ See Sosa *ibid.*, pp. 167-73.

- (JA) If one takes up a question, it is epistemically better to have an answer than not to have an answer, provided the answer is justified (competent). More specifically: One's conscious answer to the question is epistemically better than one's conscious suspension of judgment, provided one's answer is justified (competent).
- (AB) If one takes up a question, it is epistemically better to answer that question aptly than not to answer it at all. More specifically: One's apt answer to the question is epistemically better than any attitude that falls short of that, amounting only to suspending judgment and not venturing an answer.¹⁸⁷

According to Sosa, (6) causes trouble for the value platitudes as follows. First, assume that (6) is true and let ' b ' name the value of ' x ' for which (6) is true. In this case, S believes that p just in case $\text{Pr}(p) \geq b$. Now suppose that Assertive knows that p , suppose that his confidence that p is b , and suppose that his evidence that p just barely justifies him in having confidence b that p . That is, suppose that Assertive's confidence that p would outstrip his justification for p if he were any more confident than b that p . Finally, suppose that Diffident has much better evidence that p than Assertive has, and suppose that, because Diffident has the intellectual virtue of being doxastically cautious, she is just slightly less confident than b that p . Her confidence that p is b minus 0.0001, let us say.¹⁸⁸ Now, since Diffident's confidence

¹⁸⁷ (KA), (JA) and (AB) are taken *verbatim* from *ibid.*, pp. 170 and 171.

¹⁸⁸ Let a be Assertive's ideal level of confidence that p , given his evidence for p . The idea here is that, given Assertive's evidence, Assertive would be justified in any level of confidence from x to b such that $x \geq a$, but no level of confidence outside the interval $x \geq a$. Because Assertive is slightly intellectually impetuous, his confidence is b , which is as high as it could possibly be within the limits of his evidence. In contrast, where d is

that p is *below* b , she doesn't believe that p . Thus, while Assertive knows that p , Diffident doesn't. But now (KA) tells us that Assertive is in an epistemically better state than Diffident is in, and this seems wrong. There are two relevant differences between the respective attitudes of Assertive and Diffident towards the proposition that p : (a) Diffident has much better evidence for the proposition that p than Assertive has—if Diffident were to proportion her confidence to her evidence, her belief would be a paradigm of knowledge—and (b) Assertive is only 0.0001 more confident that p than Diffident is. But surely, given that these are the only relevant differences between their attitudes toward the proposition that p , Assertive *is not* in an epistemically better state than Diffident is in. So, if we conjoin (6) and (KA), we get an incorrect result. Moreover, given the harmless assumption that knowing the answer to a question entails both having an answer to that question and answering that question aptly, we get an incorrect result if we conjoin (6) with (JA) or (AB). Given that (6) is true, (JA) and (AB) both tell us the same thing that (KA) tells us: Assertive is in an epistemically better state than Diffident is in. Since Assertive *is not* in an epistemically better state than Diffident is in, we should either reject (6) or reject all of the value platitudes.¹⁸⁹

Sosa rejects (6) and, in its place, he gives us an account of belief modelled on affirmation—one that sits well with the value platitudes. Sosa endorses the following view.

Diffident's ideal level of confidence that p , given her evidence for p , Diffident would be justified in any level of confidence from x to y such that $x < b < d < y$, but no level of confidence outside the interval $x < b < d < y$. Because Diffident has the virtue of being intellectually cautious, her confidence just barely falls below b , but it remains higher than x . The result is that Diffident and Assertive *both* depart from their ideal level of confidence, and both stay within the range of their justified levels of confidence, but Diffident departs in the direction of agnosticism while Assertive departs in the direction of certainty. Moreover, and crucially, Diffident's departure is explained by her possessing the intellectual virtue of doxastic caution while Assertive's departure is explained by his being slightly intellectually impetuous.

¹⁸⁹ Ibid., pp. 170-2.

- (7) S believes that p just in case she is disposed to affirm that p (where S affirms that p just in case she either publically asserts that p or privately assents to the proposition that p).¹⁹⁰

Call this view of belief, ‘the affirmative view.’ As Sosa points out, no matter how confident I am that p , I will not be able to use the proposition that p as a premise in reasoning or share the information that p with those who want to know unless I am disposed to affirm that p . If we retell the above story so that Diffident lacks belief, not for lack of confidence, but for lack of any disposition to affirm, we again get the result that Assertive is in an epistemically more valuable state than Diffident is in. But now this result seems exactly *right*. Surely, if Assertive can use the proposition that p as a premise in reasoning and share the information that p with those who want to know, and if Diffident can neither use the proposition that p as a premise in reasoning nor share the information that p with those who want to know, then Assertive *is* in an epistemically better state than Diffident is in *vis-à-vis* the proposition that p . So there seems to be no problem conjoining (7) with the value platitudes.¹⁹¹

3. The Threshold View Defended

Unfortunately, this argument against the threshold view runs into trouble. Reconsider (4).

- (4) $\exists x \forall S p t \Box (S \text{ believes that } p \text{ at } t \supset \text{Pr}(p) \geq x \text{ at } t)$

If (4) is false, then (8) is true, since (8) is equivalent to (4)’s negation.

- (8) $\forall x \exists S p t \Diamond (S \text{ believes that } p \text{ at } t \wedge \text{Pr}(p) < x \text{ at } t)$

¹⁹⁰ Ibid., p. 172.

¹⁹¹ Ibid. pp. 172-3.

But (8) is unambiguously false. If (8) is true, then, for some values of ‘ S ,’ ‘ p ’ and ‘ t ,’ S might believe that p while her confidence that p is less than 0. Since it is not even *possible* to be less confident than 0—as we said above, ‘0’ names the lowest level of confidence that’s possible—there is clearly no value of ‘ S ,’ ‘ p ’ or ‘ t ’ where S might believe that p at t while her confidence that p is less than 0 at t . So (8) is false and we are saddled with (4). Now consider the following claim.

$$(9) \quad \exists x \forall S p t \square (S \text{ knows that } p \text{ at } t \supset \text{Pr}(p) \geq x \text{ at } t)$$

Since knowledge entails belief, the truth of (4) guarantees the truth of (9). So let ‘ b ’ name the highest value of ‘ x ’ for which (4) is true, let ‘ k ’ name the highest value of ‘ x ’ for which (9) is true, and consider (10), which says that b is significantly lower than k .

$$(10) \quad b \ll k$$

If (10) is true, there is a lowest level of confidence that is consistent with belief, and there is also a lowest level of confidence that is consistent with knowledge, but the lowest level of confidence that is consistent with belief is significantly lower than the lowest level of confidence that is consistent with knowledge. So if (10) is true, S might fall below the lowest level of confidence that is consistent with knowledge while staying well above the lowest level of confidence that is consistent with belief.

Sosa’s argument against the threshold view takes it for granted that (10) is false. Sosa rejects the threshold view on the grounds that it conflicts with the value platitudes. In order to bring the threshold view into conflict with the value platitudes, Sosa stipulates that

Assertive knows that p (justifiedly believes that p , aptly believes that p) even while his confidence that p equals b , the lowest level of confidence that is consistent with belief. But if (10) is true, this stipulation is incoherent. Given (10), either Assertive's confidence that p is as low as b , in which case Assertive is not nearly confident enough for knowledge, or Assertive *is* confident enough for knowledge, in which case his confidence that p is not anywhere near as low as b . So in order for Sosa to run his argument, he must either reject (10) or find some way to bring the threshold view into conflict with (KA) *without* stipulating that Assertive knows that p even while his confidence that p equals b . We will consider the first of these options below. How about the second option, then? Is there any way to bring the threshold view into conflict with (KA) without stipulating that Assertive knows that p even while his confidence that p equals b ? It doesn't look like it. Unless we stipulate that Assertive knows that p , Assertive's attitude toward the proposition that p is not even *relevant* to (KA). But once we *do* stipulate that Assertive knows that p , (10) tells us that Assertive is much more confident that p than Diffident is, since Diffident lacks belief. Yet if Assertive is much more confident that p than Diffident is, it is no longer problematic to claim that Assertive is in an epistemically better state than Diffident is in. So, if (10) is correct, there is no conflict between the threshold view and (KA). And the same holds for (JA) and (AB). Principles analogous to (10) that swap out knowledge for justified and apt belief would render the threshold view compatible with (JA) and (AB) in exactly the way that (10) renders the threshold view compatible with (KA).¹⁹² So Sosa's argument against the threshold view works only if (10) and its analogues for justified and apt belief are all false.

¹⁹² At least, this is true so long as knowledge-level justification is the justification at issue here.

Without Sosa's arguments against the threshold view, however, it is unclear what motivates the affirmative view. So let us assume that (10) and its analogues *are* all false. Are we now in position to accept the affirmative view? No. Because (4) is true, Sosa's argument against the threshold view fails even if (10) is false.

If (10) is false, one of three things is the case: either (i) b is greater than k , (ii) b is equal to k , or (iii) b is less than k , but by so little that the difference between b and k is not significant. Which is it, then? We have stipulated that b is the lowest level of confidence (whatever it is) that's consistent with belief, and we have stipulated that k is the lowest level of confidence (whatever it is) that's consistent with knowledge. We already know that the lowest level of confidence that's consistent with belief is not *greater* than the lowest level of confidence that's consistent with knowledge—if it *were*, you could know without believing—so we know that b is not greater than k . Thus (i) is false and rejecting (10) commits us to either (ii) or (iii). In other words, if (10) is false, then b is either equal to k or insignificantly less than k . But there's no point in even distinguishing between the claim that b is equal to k and the claim that b is insignificantly less than k : b is *insignificantly* less than k iff, for all intents and purposes, b is *equal* to k . This is the force of the word 'insignificantly.' So we can just ignore the second of these options. In this case, (10) is false only if b equals k . By rejecting (10), we therefore saddle ourselves with (11).

$$(11) \quad b = k$$

But Sosa's argument against the threshold view works only if (10) is false. So preserving Sosa's argument against the threshold view saddles us with (11).

With (11), however, we run into new trouble. First, if (11) is true, then the affirmative view is no longer defensible. We have stipulated that k is the lowest level of confidence that is consistent with knowledge, so, whatever k is, it's pretty high. But this means that, if (11) is true, then b is pretty high too. Now by stipulation, b is the lowest level of confidence (whatever it is) that's consistent with belief. So, if belief is just the disposition to affirm (as the affirmative view says), then b is also the lowest level of confidence that's consistent with *the disposition to affirm*. It follows that, if (11) is true, then the lowest level of confidence that's consistent with the disposition to affirm is quite high. The problem is, the lowest level of confidence that's consistent with the disposition to affirm seems quite *low*. It seems that S can be disposed to affirm that p so long as she is more confident than not that p . Here's why.

The affirmative view tells us that S believes that p only if she is disposed to affirm that p . Presumably, S is disposed to affirm that p only if she *does* affirm that p (or at least *would* affirm that p) in certain relevant circumstances. So suppose that S is absolutely certain that it's not the case that p , and consider the circumstances in which, *being absolutely certain that it's not the case that p* , she would affirm that p . Are any of these circumstances among the relevant ones? Clearly not. If S is absolutely certain that it's not the case that p , then she would only affirm that p in circumstances where she doesn't mean to affirm a *truth* by affirming that p (e.g., in circumstances where she means to *lie* by affirming that p). And surely, if S would only affirm that p in circumstances where she doesn't mean to affirm a truth by affirming that p , then S would not affirm that p in any of the relevant circumstances. But now, if S would not affirm that p in any of the relevant circumstances, then, according to the affirmative view, S doesn't believe that p . So the affirmative view tells us that S can't believe that p unless her

confidence that p is greater than 0. Since the affirmative view says that S believes that that p just in case she is disposed to affirm that p , the affirmative view says that S can't be disposed to affirm that p unless her confidence that p is greater than 0.

Now, plausibly, arguments analogous to this one can show that, just as S can't be disposed to affirm that p unless her confidence that p is greater than 0, she can't be disposed to affirm that p unless her confidence that p is at least as high as 0.5. Perhaps arguments analogous to this one might even show that S can't be disposed to affirm that p unless her confidence that p is *higher* than 0.5. But what about (for example) 0.51? Can we show that S can't be disposed to affirm that p unless her confidence that p is greater than 0.51? I don't see how. If S is 0.51 confident that p , then she's more confident than not that p , and I see no reason why someone who's more confident than not that p could not be disposed to affirm that p . The style of argument that we just used to show that S can't be disposed to affirm that p unless her confidence that p is greater than 0 gives us no reason to think that S can't be disposed to affirm that p unless her confidence that p is greater than 0.51, and I see no other available arguments for this conclusion. In fact, to my mind, S intuitively *can* be disposed to affirm that p while her confidence that p is as low as 0.51. (I will provide a case that shows as much later.) So it seems highly unlikely that the disposition to affirm even requires confidence greater than 0.51. But now, if belief is just the disposition to affirm, then it's possible to believe that p while being 0.51 confident that p . Thus, if the affirmative view is correct, then b isn't higher than 0.51. But we're assuming that (11) is true, and, according to (11), b equals k . So if the affirmative view is correct, it's possible to *know* that p while being 0.51 confident that p . But surely, knowledge requires considerably greater confidence

than 0.51. The suggestion that S might know that p even while her confidence that p is 0.51 seems crazy.¹⁹³ So, given that (11) is true, the affirmative view seems indefensible.

But now, without the affirmative view, we once again lose the value platitudes. Again, suppose that Assertive's confidence that p is b , suppose that Assertive's confidence is proportionate to his evidence, and assume that Assertive satisfies all the necessary conditions for knowing that p .¹⁹⁴ And again, suppose that Diffident has much better evidence that p than Assertive has, suppose Diffident lacks knowledge that p , and suppose that Diffident lacks knowledge that p for no other reason than that her confidence that p just barely falls below b . Given these suppositions, (KA) once again entails that Assertive is in an epistemically better state than Diffident is in. And once again, this seems wrong. There are two relevant differences between the respective attitudes of Diffident and Assertive towards the proposition that p : (a) Diffident has much better evidence for the proposition that p than Assertive has—if Diffident were to proportion her confidence to her evidence, her belief would be a paradigm of knowledge—and (b) because Assertive is intellectually much less cautious than Diffident is, Assertive is just barely more confident that p than Diffident is. Since we are assuming that (11) is true, and since the affirmative view is indefensible given that (11) is true, we can't add to (a) and (b) that Assertive is disposed to affirm while Diffident is not. Without the affirmative view, it is entirely possible that Assertive and Diffident are *both* disposed to affirm. So let us stipulate that they are. Now surely, given that

¹⁹³ Though see §2 of the Appendix.

¹⁹⁴ When we first considered Assertive and Diffident, we were assuming that the threshold view is true, and so we were assuming that, since Assertive's confidence that p is b , Assertive *had* to believe that p . Now we are no longer assuming that the threshold view is true, and so we do not get the result that Assertive believes that p simply because his confidence that p is b . Instead, we get this result by stipulation. Nothing prevents us from stipulating that Assertive believes that p while his confidence is exactly b , and so we *do* stipulate that Assertive believes that p while his confidence is exactly b .

(a) and (b) are the only relevant differences between their attitudes toward the proposition that p , Assertive *is not* in an epistemically better state than Diffident is in. So (KA) comes out false. And of course, (JA) and (AB) also incorrectly say that Assertive is in an epistemically better state than Diffident is in. So we once again find ourselves poorly positioned to accept the value platitudes.

What led us to this uncomfortable position? Not the threshold view of belief. The threshold view is equivalent to (6) and, for everything we have said, (6) might be false.

$$(6) \quad \exists x \forall S p t \Box (S \text{ believes that } p \text{ at } t \equiv \text{Pr}(p) \geq x \text{ at } t)$$

After all, for everything we have said, the lowest level of confidence that's consistent with belief is *lower* than the lowest level of confidence that entails belief. Rather, we got here by accepting (4) and rejecting (10).

$$(4) \quad \exists x \forall S p t \Box (S \text{ believes that } p \text{ at } t \supset \text{Pr}(p) \geq x \text{ at } t)$$

$$(10) \quad b \ll k$$

We accepted (4) because, without it, we get the absurd result that, for some values of ' S ,' ' p ' and ' t ,' S might believe that p at t while her confidence that p is below 0 at t . Rejecting (4) is clearly not the way to go. If (4) conflicts with the value platitudes, so much the worse for the value platitudes. What about (10), then? Did we go wrong by rejecting it? Well, what would we gain by *accepting* it? We would get the value platitudes back. If (10) is true, and if Assertive

knows that p while Diffident's confidence that p is too low for belief, then Diffident is *much* less confident than Assertive is that p —in which case it is not at all implausible that Assertive is in an epistemically better state than Diffident is in. If (10) is true, (KA) gets the right result. Moreover, it would be arbitrary to accept (10) while rejecting its analogues for justified and apt belief,¹⁹⁵ and these analogues would get the same result—that Assertive is in an epistemically better state than Diffident is in. So, if we accept (10) and its analogues for justified and apt belief, we can continue to embrace the value platitudes.

Where does this leave us? Sosa argued for the affirmative view by arguing that it's consistent with the platitudes while the threshold view isn't. As we have seen, however, the platitudes and the affirmative view are both *false* unless (10) is true, and yet, if (10) *is* true, then the platitudes are *consistent* with the threshold view. Since the platitudes can't cause trouble for the threshold view if they are either false or consistent with the threshold view, the platitudes can't cause trouble for the threshold view. But now the affirmative view seems unmotivated. Moreover, nothing we have said so far undermines the motivation we found for the threshold view in §1, so the threshold view *is* motivated. The argument we considered in §1 doesn't provide a conclusive case for the threshold view, but it does provide some reason to accept it, and, right now, we seem to have *no* reason to accept the affirmative view. At this juncture, then, the threshold view looks like the best motivated of our available options. So Sosa's argument for the affirmative view falls short.

4. Is Knowledge Distinctively Valuable?

¹⁹⁵ Remember, we are worried about knowledge-level justification here.

What consequences follow from this conclusion? Sosa relies on the affirmative view to establish that knowledge is distinctively valuable (by which we will mean that, for any value of ' p ' where p is true, it is epistemically better to know that p than to believe that p without knowing that p). But if the threshold view is better motivated than the affirmative view, we might lose this result. How does Sosa argue that knowledge is distinctively valuable, then? He arrives at this conclusion *via* a novel argument from the knowledge norm of assertion, as follows.

If knowledge is the norm of assertion, it is plausibly also the norm of affirmation, whether the affirming be private or public. Affirmation that p moreover seems epistemically proper and worthy if, and only if, the disposition to so affirm is then epistemically proper and worthy. We can now argue as follows:

- i. Knowledge is the norm of affirmation: *i.e.*, to affirm that p with full epistemic propriety or worth requires knowing that p .
- ii. Knowledge is the norm of belief: *i.e.*, to believe that p —to be disposed to affirm that p —with full epistemic propriety or worth, requires knowing that p .
- iii. It is epistemically better to believe with full epistemic propriety or worth than to believe without such propriety or worth.
- iv. Therefore, knowledge is epistemically better than merely true belief, which is true belief that falls short.¹⁹⁶

¹⁹⁶ *Ibid.*, pp. 177.

Sosa doesn't just assume that the knowledge norm holds and then argue that knowledge is distinctively valuable. Before arguing that knowledge is distinctively valuable, he offers a novel argument for the knowledge norm. Let us take it for granted that this argument succeeds and that knowledge *is* the norm of assertion. Now, while Sosa's argument for the knowledge norm doesn't require that the threshold view is false, the above argument *does* require that the threshold view is false. The proposition that you believe iff you are disposed to affirm is inconsistent with the threshold view; yet this proposition appears explicitly in Sosa's second premise. Because the threshold view is thus far better motivated than the affirmative view, Sosa's argument for the distinctive value of knowledge also falls short.

5. Back to the Affirmative View (or Something Near Enough)

Perhaps there's a fix, however. If we can establish the affirmative view on other grounds than that it is consistent with the value platitudes while the threshold view is not, then perhaps we can endorse Sosa's argument for the distinctive value of knowledge. So consider the following case. Mark and Teresa are watching a PBS documentary on gambling. The host of the show puts 100 marbles in a bag—99 red ones and one green one—and jostles the bag until the green marble could be anywhere in the bag. The host explains the contents of the bag to a passerby, tells the passerby that he is going to draw a random marble from the bag, and offers the passerby a bet that pays 1¢ if the draw produces a red marble and costs 99¢ if the draw produces a green one. Now Mark and Teresa know that there is a 99% chance that the draw *will* produce a red marble, so suppose that they both think that the bet is perfectly fair. That is, suppose that they both think that the passerby should be indifferent between bets 1 and 2, below.

Bet 1: The host puts 1¢ in the pot and the passerby puts 99¢ in the pot. The passerby gets the whole pot if the draw produces a red marble. The host gets the whole pot if the draw produces the green one.

Bet 2: The passerby puts 1¢ in the pot and the host puts 99¢ in the pot. The host gets the whole pot if the draw produces a red marble. The passerby gets the whole pot if the draw produces the green one.

Then it seems plausible that Mark and Teresa are equally confident that the draw will produce a red marble.¹⁹⁷ Specifically, it seems plausible that they are both exactly 0.99 confident that the draw will produce a red marble. Compatibly with all this, however, it seems that one of them might be disposed to affirm that the draw will produce a red marble while the other one is not.

With respect to the question whether the draw will produce a red marble, Mark and Teresa might differ in at least this respect: one of them might think it is extremely important to affirm the truth while the other thinks it is good to affirm the truth but more important to avoid affirming something false.¹⁹⁸ So let us assume that Mark and Teresa's values differ in exactly this way. Mark thinks it is extremely important to affirm the truth while Teresa thinks it is good to affirm the truth but more important to avoid affirming falsehoods. We can say that, with respect to affirming, Teresa has slightly skeptical dispositions while Mark has

¹⁹⁷ Of course there are any number of reasons why one or both of them might think the passerby should refuse the bet even while they think that the bet is fair. For a nice discussion of the relationship between confidence, actual betting, and estimates of bets as fair, see David Christensen, "Dutch-Book Arguments Depragmatized: Epistemic Consistency for Partial Believers," *The Journal of Philosophy*, Vol. 93, No. 9 (1996): 450-79.

¹⁹⁸ Cf., Sosa, "Value Matters in Epistemology," *The Journal of Philosophy*, Vol. 107, No. 4 (2010): 168-71.

slightly dogmatic dispositions. Now we can imagine that Mark and Teresa overhear Kim ask whether the draw will produce a red marble, we can imagine that Mark affirms that the draw will produce a red marble while Teresa doesn't, and we can imagine that Mark and Teresa differ in this respect for no other reason than that Teresa is a bit more skeptical than Mark.

As we are imagining the case, Mark and Teresa do not care about anything but affirming a truth and avoiding a falsehood. In order to affirm a truth, Mark affirms that the draw will produce a red marble, and, in order to safeguard against affirming a falsehood, Teresa refrains from affirming that the draw will produce a red marble. Mark is not affirming that the draw will produce a red marble in an attempt to lie, or in an attempt to affirm the first proposition that happens to come to mind, or because he really likes propositions about red things, or anything like that. He is affirming that the draw will produce a red marble in order to affirm a truth. And Teresa is not refraining from affirming that the draw will produce a red marble because she is distracted, or because she will be shot if she doesn't remain completely silent, or because she is meditating and trying to clear her mind, or anything like that. She is refraining from affirming in order to safeguard against affirming a falsehood. This is the whole story about Mark's affirming and Teresa's withholding. But then, is it not plausible that Mark *believes* that the draw will produce a red marble while Teresa doesn't? I want to say, 'yes.' In fact, we can add to this story that Mark is certain that he believes that the draw will produce a red marble, and that Teresa is certain that she doesn't believe that the draw will produce a red marble. Is Mark mistaken about what he believes? Is Teresa? It seems highly implausible that, just because Mark and Teresa are equally confident, one of them *must* be mistaken. Instead, it seems pretty obvious that Mark believes that the draw will

produce a red marble while Teresa doesn't. But if Mark believes that the draw will produce a red marble while Teresa doesn't, then the threshold view is false. To all appearances, then, Mark and Teresa give us a case where people differ with respect to belief while they are identical with respect to confidence—in which case the threshold view is false.

Now, I think, we have grounds for rejecting the threshold view. Whether the threshold view is uniquely inconsistent with the value platitudes or not, the threshold view seems to have counterexamples. Can we accept the *affirmative view*, then, and along with it Sosa's argument for the distinctive value of knowledge? Perhaps. But as we will see below, even if we can, we can only accept the affirmative view and Sosa's argument for the distinctive value of knowledge at a price.

Our conclusion that Mark and Teresa differ with respect to belief is motivated entirely by our stipulation that only Mark affirms that the marble will be red, in combination with our stipulation that this difference results from nothing more than Teresa's being more skeptical than Mark. Since (4) is true, and since we are calling the lowest value of 'x' for which (4) is true 'b,' the fact that Mark and Teresa differ with respect to belief suggests the following thesis.

- (13) *S* believes that *p* just in case (i) she affirms that *p* and (ii) her confidence that *p* is greater than or equal to *b*.

But (13) is ambiguous between (14) and (15).

- (14) *S* believes that *p* at *t* just in case, at *t*, (i) she *does* affirm that *p* and (ii) her confidence that *p* is greater than or equal to *b*.

- (15) *S* believes that *p* at *t* just in case, at *t*, (i) she is *disposed* to affirm that *p* and (ii) her confidence that *p* is greater than or equal to *b*.

Since (14) is clearly false—believing *p* at *t* doesn't require actually affirming *p* at *t*—reflection on Mark and Teresa puts us in good position to accept (15). So now, (15) is the best of our available options.

How does (15) relate to the affirmative view? Being disposed to affirm that *p* either entails having confidence greater than or equal to *b* that *p*, or it doesn't.¹⁹⁹ If it does, then (15) is *equivalent* to the affirmative view, and the fact that the value platitudes are either false or consistent with the threshold view does nothing to undermine Sosa's argument for the distinctive value of knowledge. On the other hand, if being disposed to affirm that *p* *doesn't* entail having confidence greater than or equal to *b* that *p*, then (15) gives us grounds for rejecting the affirmative view and, with it, the second premise of Sosa's argument. So, at first blush, the conclusion that (15) is not equivalent to the threshold view appears to undermine the distinctive value of knowledge. Actually, however, this appearance is misleading. Even if being disposed to affirm that *p* *doesn't* entail having confidence greater than or equal to *b* that *p*, (15) still tells us that *S* can't believe unless she is disposed to affirm, and this is all we need to run an argument very much like Sosa's. Even if the affirmative view turns out false, (15) puts us in good position to argue like this:

¹⁹⁹ Cf. *ibid.*, p. 173.

Knowledge is the norm of assertion and of affirmation more generally, so the knowledge norm governs the disposition to affirm. But the disposition to affirm is a necessary condition on belief, so the knowledge norm governs belief. Yet the knowledge norm governs belief only if beliefs that fall short of knowledge are not fully epistemically proper. So, beliefs that fall short of knowledge are not fully epistemically proper. In this case it is epistemically better to know than to have a belief that falls short of knowledge—even a true belief that falls short of knowledge. Thus, knowledge is distinctively valuable.

This argument establishes the distinctive value of knowledge no less than Sosa's original argument. The only difference between this argument and Sosa's is that Sosa's argument says believing *is* being disposed to affirm whereas this argument only says that believing *requires* being disposed to affirm. So, for purposes of evaluating Sosa's argument for the distinctive value of knowledge, the question whether (15) is equivalent to the affirmative view turns out to be trivial.

To summarize, even if we can't employ the value platitudes to reject the threshold view and thereby motivate the affirmative view, we have other means of motivating the affirmative view. To motivate the affirmative view, we simply argue for (15) as we did above and then argue that being disposed to affirm that p entails having confidence greater than or equal to b that p . Moreover, even if the affirmative view turns out false, we are still in excellent position to accept (15), in which case we are in excellent position to accept an argument very much

like Sosa's original argument for the distinctive value of knowledge. So it appears that, whether (15) is equivalent to the affirmative view or not, we have found an intuitive, well-motivated theory of belief that preserves both the value platitudes and the distinctive value of knowledge.

6. The Platitudes *versus* the Distinctive Value of Knowledge

Unfortunately, there is a hang-up. Take a third look at (10).

$$(10) \quad b \ll k$$

If (10) is true, then the lowest level of confidence that is consistent with belief is significantly lower than the lowest level of confidence that is consistent with knowledge. So if (10) is true, S might fall below the lowest level of confidence that is consistent with knowledge while staying well above the lowest level of confidence that is consistent with belief. So assume that (10) *is* true, and consider the case of Smith and Jones. Smith and Jones are like Diffident and Assertive, with three small differences. First, Smith's confidence that p is exactly k (rather than b), her confidence is proportionate to her evidence, and she satisfies all of the necessary conditions for knowing that p . Second, Jones has much better evidence that p than Smith has, and Jones lacks knowledge that p for no other reason than that his confidence that p just barely falls below k (rather than b). Third, Jones *believes* that p . (Jones's confidence that p just barely falls below the requirement for *knowledge*, but not the requirement for *belief*, so we can *stipulate* that Jones believes that p .) Now, like Assertive and Diffident, there are only two relevant differences between the respective attitudes of Jones and Smith towards the proposition that p : (a) Jones has much better evidence for the proposition that p than

Smith has—if Jones were to proportion his confidence to his evidence, his belief would be a paradigm of knowledge—and (b) because Smith is much less intellectually cautious than Jones is, Smith is just barely more confident that p than Jones is. But given that these are the only relevant differences between their attitudes toward the proposition that p , Smith is not in an epistemically better state than Jones is in. And yet, if (10) is true, Smith knows that p while Jones believes that p without knowing that p . So if (10) is true, it is false that knowledge is distinctively valuable. In this case, if (10) is true and knowledge is the norm of assertion (as we are assuming), then the knowledge norm doesn't entail that knowledge is distinctively valuable. And in this case Sosa's original argument for the distinctive value of knowledge and our revision of it both fail. The inference from Sosa's second premise to his third premise is invalid, and so is the analogous inference in our revision of Sosa's argument.

Can we show that (10) is false, then? And should we even want to, since, if (10) is false, then so are the value platitudes? We will address this second question later. In the meantime, call the view we get if (15) is *not* equivalent to the affirmative view, 'the confident affirmative view' (or 'CAV' for short). We are committed to (15), so we are committed to either CAV or the affirmative view. The question is, can we show that (10) is false, given that either CAV or the affirmative view is true?

As we have already seen, we can't show that (10) is false, given that the affirmative view is true. If (10) is false, then b equals k ; but the affirmative view seems indefensible given that b equals k . Thus, given that the affirmative view is true, we are in terrible position to show that (10) is false. Suppose that CAV is true, then. In this case, belief requires both the disposition to affirm and confidence at least as high as b , but the disposition to affirm doesn't itself

require confidence at least as high as b . Plausibly, CAV does put us in good position to deny (10). At least, here is an argument to this effect: Consider just the thesis that b equals k . This thesis is *prima facie* plausible—so plausible that Sosa takes it for granted in his argument that the threshold view conflicts with the value platitudes.²⁰⁰ But given this plausibility, we should assume that b equals k unless we have some positive reason to think otherwise. And yet, if b equals k , then (10) is false. So we should assume that (10) is false unless we have some positive reason to think that b doesn't equal k . Do we have any positive reason to think that b doesn't equal k , then? If the affirmative view is true, we do. We just saw as much. But right now we are assuming that CAV (rather than the affirmative view) is true, and, on CAV, b can be much higher than the lowest level of confidence that is consistent with the disposition to affirm. So it seems that, unlike the affirmative view, CAV gives us no reason to deny that b equals k . Thus, downstream from CAV, we should assume that b equals k . But b equals k only if (10) is false. So, if we accept CAV, we should deny (10). CAV therefore puts us in great position to deny (10).

Whether this argument works or not, rejecting the affirmative view in favor of CAV is our only chance at preserving the distinctive value of knowledge.

Let us take stock. We are committed to (15), and, courtesy of our commitment to (15), we are committed to either CAV or the affirmative view. While CAV might put us in position to deny (10), the affirmative view saddles us with (10). If (10) is true, then knowledge is not

²⁰⁰ Sosa stipulates that Assertive knows that p , that Diffident is just barely less confident than Assertive that p , and that, in virtue of nothing but Diffident's being less confident than Assertive that p , Diffident doesn't believe that p . But as we saw in §3, these stipulations are jointly coherent only if b is at least as great as k . Since b can't be *greater* than k , Sosa's stipulations assume that b equals k .

distinctively valuable. So CAV has at least this advantage over the affirmative view: unless we accept CAV, we can't preserve the distinctive value of knowledge. But CAV allows us to preserve the distinctive value of knowledge only if we deny (10), and, as we saw in §3, once we deny (10), we must also deny the value platitudes. So CAV only saves the distinctive value of knowledge at the cost of forcing us to deny the value platitudes. Of course, we can *accept* the value platitudes if we accept the affirmative view instead of CAV. But this way, we saddle ourselves with (10) and thereby lose the distinctive value of knowledge. Because knowledge is not distinctively valuable if (10) is true, and because the value platitudes are false if (10) is false, neither the affirmative view nor CAV gets us everything we want. Whether we accept CAV or the affirmative view, we must say something *prima facie* undesirable about the value of knowledge—either knowledge *is not* distinctively valuable or the value platitudes are all *false*.²⁰¹

7. The Nature of Belief and the Valuable of Knowledge

Which is it, then? Are the platitudes true, or is knowledge distinctively valuable? Is CAV correct, or is the affirmative view the right theory of belief? Does *b* equal *k*, or not? Which package should we accept: the affirmative view plus (10) and the platitudes, or CAV plus (10)'s negation and the distinctive value of knowledge—or perhaps some third package, like CAV plus (10) and the platitudes? And how should we go about deciding between these options?

²⁰¹ And note, giving up (15) will not help here. So long as (4) is true, knowledge isn't distinctively valuable if (10) is true, and the platitudes are false if (10) is false. So long as (4) is true, these conditionals hold no matter which theory of belief turns out to be true.

I find the platitudes harder to give up than the distinctive value of knowledge, and this is a happy result, since, even considered apart from our decision between the platitudes and the distinctive value of knowledge, I find (10) considerably more plausible than its negation. Consider the following case.

Assertive Kim: A marble has been randomly drawn from a bag containing 49 green marbles and 51 red ones. Jerome and Kim know the contents of the bag and that the draw was random, but the outcome of the draw has not been announced. Jerome wonders aloud whether the draw produced a red marble or a green one, and Kim responds, “it produced a red one.” Jerome finds Kim’s assertion jarring, so he assumes that Kim is just guessing—he assumes that, contrary to the appearance given by her flat-out asserting that it produced a red marble, she doesn’t really *believe* that the draw produced a red marble. But Kim insists that she *is not* guessing. “It is not a guess,” she tells Jerome. “I *know* that the draw produced a red marble.” Giving Kim the benefit of the doubt, Jerome assumes that Kim must have some inside knowledge of the draw. But Kim insists that she doesn’t have any information about the draw that Jerome doesn’t have. Now Jerome is confused. He is convinced that Kim really does believe that the draw produced a red marble, but he doesn’t see how she could believe this, given that she knows there is only a 51% chance that the marble is red. Jerome concludes that, while Kim doesn’t have anything to go on beyond the 51% chance of a red marble, her confidence must be well above 0.51. He concludes that Kim must be flouting the Principal Principle. To verify his theory, Jerome asks Kim how confident she is that the marble is red. Kim responds with a

question: “Are you asking me to place a numerical value on my confidence?” Jerome says he is, and Kim replies that, where 1 equals certainty that the marble is red and 0 equals certainty that the marble is not red, her confidence that the marble is red is exactly 0.51. At this point, Jerome suspects that Kim is mistaken about her confidence, so he devises a test. He describes bets 3 and 4 below, tells Kim that he is willing to enter into either bet, and asks Kim which bet she would prefer.

Bet 3: Jerome puts \$49 in the pot and Kim puts \$51 in the pot. Kim gets the whole pot if the draw produced a red marble. Jerome gets the whole pot if the draw produced a green one.

Bet 4: Kim puts \$49 in the pot and Jerome puts \$51 in the pot. Jerome gets the whole pot if the draw produced a red marble. Kim gets the whole pot if the draw produced a green one.

Jerome thinks that Kim is really much more confident than 0.51 that the draw produced a red marble, and, thus, that Kim will prefer bet 3 to bet 4. But to Jerome’s surprise, Kim says it doesn’t matter. Both bets are perfectly fair, she says. As far as she is concerned, *Jerome* can pick. Now Jerome *does* pick. He picks bet 4, and Kim says that that is fine by her. Jerome devises many similar tests and each time Kim acts exactly the way we would expect someone who was only 0.51 confident to act. At this point Jerome becomes convinced that Kim really is 0.51 confident that the

marble is red. But now he is also convinced that Kim no longer *believes* that the marble is red. “No, I still believe that it is red,” Kim replies. “In fact, it *is* red. I *know* that it is,” she says. And she adds that her knowledge of this outcome is based solely on the 51% chance that the draw produced a red marble.

What should we say about this case? It is definitely weird. Kim seems kind of crazy. Given that she has no reason to think the draw produced a red marble other than the 51% chance that it did, she clearly should not believe that the draw produced a red marble, and she definitely doesn't *know* that it produced a red marble. But all this is beside the point. The question is not whether Kim *should* believe that the draw produced a red marble; it is whether she *does* believe that the draw produced a red marble. And if she does, is her confidence really 0.51? I think the answer to both questions is ‘yes’ (or at least, that it could be ‘yes’). Despite her quirks, I think Kim shows us how belief could be coupled with confidence as low as 0.51.²⁰² So, to my mind, 0.51 confidence is consistent with belief. But 0.51 confidence is clearly inconsistent with *knowledge*, so Kim shows us how belief could be coupled with insufficient confidence for knowledge. But now it follows that *b* is significantly lower than *k*, and (10) is true. Thus, reflecting on cases like Assertive Kim, (10) looks exactly right. Even

²⁰² But is there not a sense in which Kim is actually *much* more confident than 0.51 that the draw produced a red marble? I do not think so. The way I am imagining the case, Kim manifests all the behavior we would expect from someone who is only 0.51 confident that the draw produced a red marble, with one exception: when Kim considers the question whether the draw produced a red marble, in order to affirm a truth, she asserts flat-out that the draw produced a red marble and gives unqualified mental assent to the proposition that the draw produced a red marble. The right thing to say about Kim is not that, while there is a sense in which Kim's confidence is 0.51, there is also a sense in which her confidence is much higher than 0.51. Rather, the right thing to say is just that, with respect to affirming the proposition that the draw produced a red marble, Kim doesn't act the way most people who were only 0.51 confident would act. Typically, people are only as assertive as Assertive Kim when they are much more confident in a proposition than 0.51.

apart from our choices between the affirmative view, CAV, the value platitudes and the distinctive value of knowledge, we have good reason to accept (10).

Moreover, when we consider (10) in relation to the affirmative view, CAV, the value platitudes and the distinctive value of knowledge, we gain even more reason to embrace (10). As we have seen, (10) is true only if knowledge is not distinctively valuable, and (10) is false only if the value platitudes are all false. So setting cases like Assertive Kim aside, we can ask which is more plausible: that knowledge *is* distinctively valuable, or that the value platitudes are all *true*?

Let us distinguish the kind of ignorance that results when your true belief falls short of knowledge from the kind of ignorance that results when you hold either no belief or a false belief. Let us call the first variety of ignorance ‘informed ignorance,’ and let us call the latter variety of ignorance ‘total ignorance.’ The value platitudes are true just in case, for every value of ‘ p ’ where p is true, it is epistemically better to know that p , justifiedly believe that p or aptly believe that p than to be totally ignorant that p . And knowledge is distinctively valuable just in case, for every value of ‘ p ’ where p is true, it is epistemically better to know that p than to have informed ignorance that p . Now, the platitudes are platitudes because they seem so obvious. How could knowing that p , justifiedly believing that p or aptly believing that p fail to be epistemically better than total ignorance that p ? If knowledge is at the top of the value spectrum, total ignorance is not just below it, it is all the way at the bottom.²⁰³ There seems to be no contest between total ignorance that p and knowing that p , justifiedly believing that p or aptly believing that p . This is why, upon concluding that the

²⁰³ Plausibly, the ignorance that accompanies false belief is epistemically even worse than the ignorance that accompanies no belief.

threshold view allows total ignorance that p to be epistemically better than knowing that p , Sosa did not hesitate to reject the threshold view. Now compare the platitudes to the distinctive value of knowledge. It is not so clear that knowledge is distinctively valuable. Philosophers have not just struggled to say why it is better to know that p than to hold the justified true but Gettiered belief that p , they have struggled to say why it is even better to know that p than to hold the *unjustified* true belief that p . So while the platitudes seem obviously true, it is not implausible that knowledge *is not* distinctively valuable. In our choice between the platitudes and the distinctive value of knowledge, then, we should prefer the platitudes. But the platitudes saddle us with (10). So (10) is plausible in its own right, and it follows from the platitudes, which are considerably harder to give up than the distinctive value of knowledge.²⁰⁴

Everything considered, then, we should accept (10) and the value platitudes at the cost of rejecting the distinctive value of knowledge. This leaves us with a final question: which should we add to our package, CAV or the affirmative view? The answer has to be ‘the affirmative view.’ First, the affirmative view is simpler and more elegant than CAV—especially given that (10) is true. If the affirmative view is right, then belief is just the disposition to affirm. If CAV is right, then belief is the disposition to affirm plus confidence above some level that is higher than the level of confidence required for the disposition to

²⁰⁴ Does (10) saddle us with an unwieldy theory of knowledge—something like *knowledge is justified true un-Gettiered belief plus sufficient confidence*, or *knowledge is apt belief plus sufficient confidence*? No. Take the claim that justified true un-Gettiered belief suffices for knowledge. When epistemologists make this claim, they do not mean that, once you have a true un-Gettiered belief, then, so long as you have *some* justification for that belief, you have knowledge. It is implicit in all discussion of the JTB theory that the justification at issue is *knowledge-level* justification. But since knowledge-level justification is the justification at issue here, it is not implausible that justified belief entails high confidence. In this case, however, people working in the JTB tradition do not need to add a high-confidence condition to their theories of knowledge any more than they need to add, as a necessary condition for knowing, that somebody somewhere believes something or that $2 + 2 = 4$. (These, of course, are necessary conditions for knowing.)

affirm, but lower than the level required for knowledge. CAV is rather cumbersome. Second, the affirmative view fits very nicely with (10) while CAV doesn't. As we said, if belief is just the disposition to affirm, then b is not even as high as 0.51, in which case b is significantly lower than k . Because the affirmative view says belief *is* just the disposition to affirm, the affirmative view can explain why (10) is true. It is true because the disposition to affirm doesn't require very much confidence. But on the other hand, if CAV is true, it is rather mysterious why (10) is true. If the disposition to affirm is not responsible for the fact that b is significantly lower than k , then what is? Why is b so low? If CAV were true, we would expect b to equal k . Third, it is not clear what motivates CAV. We left CAV on the table because we thought that it might allow us to reject (10) and thereby accept the distinctive value of knowledge. But as we have seen, (first) we do not *want* to accept the distinctive value of knowledge, since doing so would require us to reject the value platitudes, and (second) CAV *doesn't* allow us to reject (10). Cases like Assertive Kim motivate (10) whether we accept CAV or not. Everything considered, then, the affirmative view is surely a better addition to (10) and the value platitudes than CAV. This, then, is the view that we ought to accept: belief is the disposition to affirm, the lowest level of confidence that is consistent with belief is significantly lower than the lowest level of confidence that is consistent with knowledge, and, while knowledge is better than total ignorance, knowledge is not distinctively valuable.

8. The (Previously) Mysterious Interval $m \times n$

In §1, we let m name the lowest level of confidence that is consistent with belief, and we let n name the lowest level of confidence that entails belief. We then considered the possibility

that m is less than n , and we eliminated this possibility on the grounds that it would saddle us with a mystery. If m is less than n , then confidence in the interval $m \times n$ leaves it open whether S believes. *But if confidence in the interval $m \times n$ leaves it open whether S believes, then why (we asked without knowing how to answer) should confidence **outside** this interval fully determine whether S believes? What factor, that can't make a difference outside this interval, can make a difference **inside** this interval? What is so special about this interval?* Now we can answer these questions. The disposition to affirm is not consistent with just any level of confidence, and some levels of confidence entail the disposition to affirm. To a first approximation,

- (16) S is disposed to affirm that p just in case, if S were to either affirm that p for the sole purpose of affirming a truth or refrain from affirming that p for the sole purpose of safeguarding against affirming a falsehood, then S would affirm that p .²⁰⁵

As we saw in §3, it falls out of (16) that S can't be disposed to affirm that p while she is absolutely certain that it is not the case that p . So now suppose that S is absolutely certain that it *is* the case that p . In the nearest worlds where S either affirms that p for the sole purpose of affirming a truth or refrains from affirming that p for the sole purpose of safeguarding against affirming a falsehood, she doesn't refrain from affirming that p for the sole purpose of safeguarding against affirming a falsehood. (She is absolutely certain that p , so, by her own lights, p is true in every possible world. But in this case, if she refrains from affirming that p , she doesn't do so to safeguard against affirming a *falsehood*.) It follows that,

²⁰⁵ Read “for the sole purpose of affirming a truth” as short for “for the sole purpose of affirming a truth *in the act of affirming that p* ,” and make the analogous modification for “for the sole purpose of safeguarding against affirming a falsehood.”

in the nearest worlds where S either affirms that p for the sole purpose of affirming a truth or refrains from affirming that p for the sole purpose of safeguarding against affirming a falsehood, she affirms that p . But then, according to (16), S is disposed to affirm that p . Thus, while confidence 0 is inconsistent with the disposition to affirm, confidence 1 entails the disposition to affirm.

Generalizing, we get this result: some level of confidence entails the disposition to affirm and some level of confidence is necessary for the disposition to affirm. Now let ' t ' name the lowest level of confidence that is consistent with the disposition to affirm, and let ' u ' name the lowest level of confidence that entails the disposition to affirm. In this case, while confidence outside the interval $t \times u$ fully determines whether S is disposed to affirm, confidence inside this interval leaves it an open question whether S is disposed to affirm. But S is disposed to affirm just in case she believes. So t equals m , and u equals n . The special thing about the interval $m \times n$, then, is that, inside this interval but not outside of it, it is indeterminate whether one is disposed to affirm. Inside this interval but not outside of it, one's values *vis-à-vis* affirming truths and avoiding falsehoods can make a difference to whether one *believes*.

Appendix: Certainty Instead of Belief, and Uncertain Knowledge

Here I consider and reject two strategies for repairing Sosa's arguments for the affirmative view and the distinctive value of knowledge, each proposed by Sosa himself.²⁰⁶

1. Certainty Instead of Belief

²⁰⁶ I owe these strategies to personal correspondence with Sosa.

As we have said several times, knowledge requires some high level of confidence. Let us call this level of confidence (whatever exactly it is) ‘certainty.’ In this case, you can’t know that p unless you are certain that p . (Knowledge doesn’t require *absolute* certainty, of course, so certainty doesn’t require confidence as high as 1. It just requires something in the neighbourhood—confidence as high as 0.9, say.) Now, plausibly, when epistemologists say things like “knowledge is justified true un-Gettiered *belief*,” they really have in mind justified true un-Gettiered *certainty* (in the sense of ‘certainty’ that we just defined). But if this is right, then the above considerations might be off the mark, for Perhaps Sosa’s arguments should be understood as targeting *certainty* (so defined) rather than *belief*. If so, then, in outline, Sosa’s arguments really go as follows.

Take the threshold view of certainty, which is equivalent to (6’), and consider revisions of the value platitudes that discuss certainty rather than belief.

$$(6') \quad \exists x \forall S p t \Box (S \text{ is certain that } p \text{ at } t \equiv \text{Pr}(p) \geq x \text{ at } t)$$

If we let ‘ c ’ name the value of ‘ x ’ for which (6’) is true, we can bring (6’) into conflict with the revised value platitudes exactly the way that we brought (6) into conflict with the original value platitudes. We just stipulate (first) that Assertive knows (is justifiedly certain, has apt certainty) that p , that his confidence that p is c , and that his evidence that p just barely justifies him in having confidence c that p ; and (second) that Diffident has much better evidence that p than Assertive has, and that, because Diffident is intellectually very cautious, she is just slightly less confident than c that p . Now, (4’) is no less a fact of life than (4) is, since certainty entails belief.

$$(4') \quad \exists x \forall S p t \Box (S \text{ is certain that } p \text{ at } t \supset \Pr(p) \geq x \text{ at } t)$$

So we will have to deal with (4'). This is actually not difficult, however. The appropriate modification of (7) gives us (7').

$$(7') \quad S \text{ is certain that } p \text{ just in case she is disposed to affirm that } p.$$

Now, if ' ϵ ' names the lowest level of confidence that is consistent with certainty (*i.e.* the lowest level of ' x ' for which (4') is true), then (7') entails that S is disposed to affirm that p only if her confidence that p is at least as great as ϵ . In this case, we can employ (7') to save the revised platitudes in exactly the way that Sosa proposed that we employ (7) to save the original platitudes. Because Diffident's confidence is below ϵ , she is not disposed to affirm that p . But now, she can't use the proposition that p as a premise in reasoning or share the information that p with those who want to know. Meanwhile, we can just stipulate that Assertive *is* certain that p , from which it follows from (7') that Assertive *is* disposed to affirm, and so Assertive *can* use the proposition that p as a premise reasoning and share the information that p with those who want to know. In this case, the revised platitudes get the right result. Assertive *is* in an epistemically better state than Diffident is in. Now consider (10').

$$(10') \quad \epsilon \ll k$$

Where (10) created problems for Sosa's original argument, (10') doesn't create problems for this argument. First, we stipulated that ϵ equals k , so we can reject (10') and espouse (11') right off.

$$(11') \quad c = k$$

Second, with the proper characterization of the disposition at issue in (7')—and note, the disposition at issue in (7') is not the same disposition as the disposition at issue in (7)—it is intuitively obvious that *S* can't be disposed to affirm unless her confidence is at least as great as *k*. So (10') doesn't create any of the problems for (7') that (10) created for (7). Moreover, we can stipulate that knowledge is distinctively valuable just in case, for any value of '*p*' where *p* is true, it is epistemically better to know that *p* than to be *certain* that *p* without knowing that *p*; and with this stipulation, we can endorse an argument almost identical to Sosa's original argument for the distinctive value of knowledge. So, by simply switching from belief to certainty (as we defined 'certainty' above), we get a sound argument against the threshold view, we motivate the affirmative view, and we get to keep both the value platitudes and the distinctive value of knowledge. In other words, we get everything Sosa was after. So the arguments in §§1-7 really miss the mark. They are relevant to the belief-conception of knowledge—the conception of knowledge according to which knowledge is something like justified true un-Gettiered *belief*—but they are irrelevant to the conception of knowledge that epistemologists are really interested in: the certainty-conception.

Unfortunately, this reply contains a simple, fatal flaw. Call the disposition at issue in (7) the belief disposition, and call the disposition at issue in (7') the certainty disposition. Given (7), anyone who is less confident than *b* will lack the belief disposition; and given (7'), anyone who is less confident than *c* will lack the certainty disposition. Now, while it is plausible that anyone who is less confident than *b* will not be able to use the proposition that *p* as a premise in reasoning or share the information that *p* with those who want to know, it is *not*

plausible that anyone who is less confident than c will not be able to use the proposition that p as a premise in reasoning or share the information that p with those who want to know. Remember, while b is lower than 0.51, c is *much higher* than 0.51. But the above argument against (6') and for (7') depends crucially on the assumption that, because Diffident's confidence is below c , diffident can't use the proposition that p as a premise in reasoning or share the information that p with those who want to know. In other words, it depends crucially on the following highly doubtful conditional.

- (C) If S 's confidence that p is below c (the minimum level of confidence that is consistent with *certainty*), then S can't use the proposition that p as a premise in reasoning or share the information that p with those who want to know.

Since (C) is doubtful, why can't we just stipulate that, even though Diffident's confidence is below c , Diffident can (and indeed *does*) use the proposition that p as a premise in reasoning and share the information that p with those who want to know? Remember, Diffident's confidence that p is well above b , so we *can* just stipulate that Diffident *believes* that p . And since Diffident has outstanding evidence for his belief that p —such good evidence that Diffident could be justified in being almost *absolutely* certain that p —there is no *normative* sense in which it seems plausible that Diffident can't use the proposition that p as a premise in reasoning or share the information that p with those who want to know. It seems that Diffident *should* use the proposition that p as a premise in reasoning and share the information that p with those who want to know.²⁰⁷ So, to all appearances, nothing prevents

²⁰⁷ In which case the knowledge norm of assertion runs into trouble.

us from just stipulating that Assertive can use the proposition that p as a premise in reasoning and share the information that p with those who want to know, in which case we have a counterexample to the revised value platitudes—one that holds even while (7') is true.

If we have a counterexample to the revised platitudes even while (7') is true, however, then the conflict between (6') and the revised platitudes gives us no reason to prefer (7') to (6'). And in this case, the revised argument for the distinctive value of knowledge fails. So, swapping out the belief-conception of knowledge for the certainty-conception of knowledge will not get us *both* the value platitudes *and* the distinctive value of knowledge. With the certainty-conception of knowledge, we face basically the same problems as we did under the belief-conception of knowledge.

2. *Unconfident Knowledge*

According to a second strategy suggested by Sosa, knowledge doesn't require more confidence than belief, but this is not because belief requires *high* confidence; it is because knowledge suffers *low* confidence. Knowledge *does* require the disposition to affirm, however—not hesitantly, but outright. If S is neither disposed to flat-out assert that p nor disposed to give full mental assent to p , then S is not in a mental state that is consistent with knowing that p . Now, apt affirmation is the right way to think about knowledge. If S is disposed to affirm that p , if her disposition to affirm that p manifests a competence, and if her getting it right about p also manifests that competence, then S knows that p . Thus, even if her confidence that p is as low as (say) 0.51, S still knows. But now we can embrace both (7) and (11)—in which case the second horn of the dilemma in §3 fails and Sosa's original argument goes through.

To my mind, it is very implausible that S could know that p even while she is not at all confident that p . We found ourselves unable establish that b is higher than 0.51, so suppose S 's confidence that p is 0.51. In this case, by S 's own lights, it is barely more probable than not that p . But how can S know that p while, by her own lights, it is barely more probable than not that p ? Let e be S 's total evidence for p and take the objective probability of p on e . That is, take the epistemic probability that p .²⁰⁸ Either S 's actual confidence matches her ideal confidence, or it doesn't. If her actual confidence matches her ideal confidence, then the epistemic probability of p is only 0.51, in which case S has far too little evidence for knowledge. So suppose on the other hand that the epistemic probability of p is quite high (high enough for knowledge). In this case, S is failing miserably to take proper stock of her evidence that p . Her evidence conclusively supports p and yet she is almost as confident that $\neg p$ as she is that p . In this case S is exhibiting massive internal irrationality. But do we really want to say that S can know that p even while she is massively internally irrational in her confidence that p ? I have a hard time saying 'yes.' So I want to resist the suggestion that S could know that p even while she is not at all confident that p .

Moreover, if we *do* accept the suggestion that knowledge suffers confidence as low as b , what do we buy ourselves? Just an argument for the distinctive value of knowledge, which (as we noted in §7) is doubtful even apart from any of the considerations in this paper. Whether we accept the suggestion that knowledge suffers confidence as low as b or not, we still have the value platitudes and a convincing argument for the affirmative view of belief. So, even if it were just a choice between the distinctive value of knowledge and the claim

²⁰⁸ Evidence should be given a very loose characterization here, as it standardly is in discussion of epistemic probability.

that S could know that p without being at all confident that p , it seems that it would be perfectly sensible to give up the distinctive value of knowledge in order to affirm that, in order for S to know that p , she must be pretty confident that p . The claim that S could know that p without being at all confident that p is both insufficiently motivated and problematic in its own right.

Chapter 7: Knowledge and Intellectual Interests

Knowledge entails true belief, but beyond this there is controversy. Ask yourself when David Hume was born. Whether you know the answer to this question depends on whether you have a belief about Hume's birth date, and whether your belief is true (if you do have one). But even supposing that your belief is true, whether you know when Hume was born plausibly also depends on the evidence for your belief, whether your belief is the product of reliable cognitive processes or faculties, whether your belief is safe or sensitive, and so on. The arguments for and against each of these conditions on knowledge are sophisticated and difficult to adjudicate. Fortunately, some properties are clearly unnecessary for knowledge. To know when Hume was born you needn't wear a size 9 shoe or know that penguins eat lantern fish, for example.

The clear irrelevance of these factors raises a question, though: why is it plausible that knowledge entails reliable cognitive faculties (for example) but implausible that knowledge entails a particular shoe size? The answer seems obvious enough, but turns out to be controversial. "Only *truth-relevant* factors are necessary for knowledge," it is tempting to say. By "truth-relevant," I mean factors that affect your grip on the truth. You don't have *any* grip on the truth *vis-à-vis* Hume's date-of-birth if you have no belief about Hume's date-of-birth, or if you have latched onto the *wrong* date. And if you have a *true* belief about Hume's birth date, the *strength* of your grip on the truth depends on your evidence, the reliability of the cognitive faculties responsible for your belief, whether your belief is safe, and so on. If

your belief is insufficiently supported by the evidence, or the product of insufficiently reliable faculties, or not sufficiently safe, then your grip on the truth will be too weak for your belief to count as *knowledge*. Of course, you do not literally *grip* the truth. This is just a metaphor. But it is not a new one. It traces back at least as far as Plato's *Meno*. In terms familiar from the *Meno*, whether you are adequately *tethered* to the truth may depend on the strength of your evidence, the reliability of your cognitive faculties, or the safety of your belief. In stark contrast, whether you are tethered to the truth *vis-à-vis* Hume's birth date clearly *need not* (and almost certainly *does not*) depend on your shoe size or your knowledge of penguin food.

Plausibly, truth-relevant factors like evidence, reliability, and safety *exhaust* the factors relevant to knowledge. If so, then two people who are identical with respect to truth-relevant factors are also identical with respect to their knowledge. This assumption has recently fallen under criticism, however, and it is the subject of dispute between what I will call 'interest-relativism' and 'interest-objectivism.' Interest-relativism is a component of interest-relative *invariantism*—the view popularized by Jason Stanley.²⁰⁹ Interest-relative invariantism conjoins the interest-relativist thesis that knowledge depends in part on our interests with the invariantist thesis that 'knows' is not a context-sensitive word.²¹⁰ Neither thesis entails the other, and I will focus solely on the debate between interest-relativism and interest-objectivism in this paper. Interest-objectivism is just the denial of interest-relativism. Since your interests will rarely affect your grip on the truth, interest-relativism entails that truth-*irrelevant* factors can prevent you from knowing. If interest-relativism is true, you can go

²⁰⁹ See Stanley (2005).

²¹⁰ The word 'knows' is context sensitive if it semantically expresses different properties or relations in different contexts of utterance.

from having a sufficiently strong grip on the truth, to having an insufficiently strong grip on the truth, without any change in your grip on the truth. Interest-objectivists resist interest-relativism precisely because interest-relativism makes knowledge depend on truth-irrelevant factors.²¹¹ I am strongly inclined toward interest-objectivism, but here I will defend interest-relativism.

Interest-relativism comes in two varieties: what I will call ‘pragmatism’ and ‘intellectualism.’ Pragmatism is the view that *practical* interests can make a difference to knowledge, while intellectualism is the view that *intellectual* interests can make a difference to knowledge. I will say exactly what these views amount to below. As we will see, pragmatism and intellectualism might both be true, but neither view entails the other. Jeremy Fantl, Matthew McGrath, John Hawthorne, Jason Stanley, and many others defend pragmatism.²¹² While pragmatism has received considerable attention in the literature, intellectualism has scarcely been identified as a position in logical space, and it has no defenders. In this paper, I defend interest-relativism by defending intellectualism. There will be little point in defending intellectualism if the arguments against pragmatism already undermine intellectualism, however, so I take it for granted in this paper that pragmatism is false.²¹³ Apart from this paper, the arguments for pragmatism exhaust the arguments for interest-relativism. As a result, many philosophers think they can resist interest-relativism by simply defending *purism* (the denial of pragmatism). This paper shows otherwise. Interest-relativism is true even if

²¹¹ See the introductory section of Grimm (2011) for a nice discussion of this point.

²¹² See, for example, Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005), Hawthorne and Stanley (2008), Fantl and McGrath (2009), and Ross and Schroeder (forthcoming).

²¹³ For a sample of the best arguments against pragmatism, see Brown (2008), (2012), and (forthcoming), and Reed (2010) and (2012).

pragmatism is false. The upshot is that knowledge depends on truth-irrelevant factors even if it does not depend on practical interests.²¹⁴

In section 1, I explain how pragmatism and intellectualism make knowledge depend on truth-irrelevant factors. In section 2, I lay out some principles relevant to my argument for intellectualism. In section 3, I present my argument for intellectualism and, in section 4, I defend my argument against possible objections. Section 4 does not merely tie up loose ends. As I show in section 4, there is no way to resist my argument for intellectualism without espousing some view that is less plausible than intellectualism itself. Finally, in section 5, I draw out the implications of intellectualism for the pragmatism/purism debate.

1. Interests, Practical *versus* Intellectual

1.1. Practical Interests

What do interest-relativists mean when they say that knowledge depends on our *interests*? Let's start with the pragmatist's notion of *practical* interests. Paradigm versions of pragmatism endorse principles like the following *knowledge-action* principle: *S* knows that *p* at *t* only if she can rationally act as if *p* at *t*.²¹⁵ Whether *S* can rationally act as if *p* will depend on the probability of *p* and the relative values of the consequences that might follow from acting as if *p*. But how should we think of these probabilities and values? Subjectively, in terms of *S*'s personal level of confidence that *p* and the values that she personally places on the consequences that might follow from acting as if *p*? Or more objectively, in terms of the

²¹⁴ Readers familiar with the pragmatic encroachment literature will note that I am hijacking the term 'intellectualism' from Jason Stanley. As Stanley uses it, 'intellectualism' names the view that I am calling 'purism.'

²¹⁵ Again, see Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005), Hawthorne and Stanley (2008), Fantl and McGrath (2009), and Ross and Schroeder (forthcoming).

epistemic probability of p and the values that S *ought* to place on the consequences that might follow from acting as if p ? Since the knowledge-action principle gets implausible results on the subjective interpretation, and since no pragmatist explicitly endorses the subjective interpretation, I will take the objective interpretation for granted. Holding this interpretation fixed, the following familiar scenarios help clarify the pragmatist's notion of practical interests.

Case A: Hannah is driving past the bank on Friday afternoon. She has her paycheck in hand and she plans to deposit it, but there are long lines, so she considers returning to deposit her paycheck on Saturday morning. She knows that it does not matter much when she deposits the paycheck.

Case B: Hannah is driving past the bank on Friday afternoon. She has her paycheck in hand and she plans to deposit it, but there are long lines, so she considers returning to deposit her paycheck on Saturday morning. But in this case she knows that she will incur an enormous fine if she does not deposit her check before Saturday afternoon.

Assume that Hannah knows that the bank will be open on Saturday in Case A, and concentrate on her choice between stopping at the bank and going straight home in Case B. By hypothesis, the following decision table accurately represents Hannah's choice between stopping at the bank and going straight home in Case B.

Table 1

	The bank will be open	The bank will not be open
Stop at the bank	<i>Waste a little time</i>	<i>Avoid enormous fine</i>
Go straight home	<i>Save a little time</i>	<i>Incur enormous fine</i>

Conditional on the proposition that the bank will be open on Saturday, Hannah wastes an insignificant amount of time by stopping at the bank on Friday, and saves that much time by going straight home on Friday. In contrast, conditional on the proposition that the bank will *not* be open on Saturday, Hannah loses a large sum of money by going straight home on Friday, but saves that much money by stopping at the bank on Friday. Since it is better to save time than to waste it, Hannah acts as if the bank will be open only if she goes straight home. But Hannah cannot *rationally* go straight home. Given her evidence that the bank will be open on Saturday, and the severe consequences that she knows will follow if she fails to deposit her check before noon on Saturday, she can see that going straight home would be too risky. This means that Hannah cannot rationally act as if the bank will be open. So, according to the knowledge-action principle, Hannah cannot know that the bank will be open in Case B. It follows that, if the knowledge-action principle is true, Hannah knows that the bank will be open in Case A but not Case B.

Now the probabilities at issue in Hannah's choice between stopping at the bank and going straight home can make a difference to whether Hannah knows that the bank will be open. Hannah cannot know that the bank will be open if the bank will probably be *closed*. There is no controversy here. When pragmatists say that Hannah's *practical interests* can make

a difference to whether Hannah knows that the bank will be open, they mean that the relative values of the consequences listed in Table 1 can make a difference to whether she knows. This is the controversial feature of pragmatism, for values like these are standardly truth-irrelevant.

1.2. Intellectual Interests

Intellectual interests mirror practical interests. Imagine that Hannah is considering the evidence for and against some proposition p and, as a result, she faces a choice between believing that p and withholding with respect to p . We can represent Hannah's choice with a decision table like Table 1, so long as we know what consequences will follow from these options. What consequences *will* follow from these options, then? The answer depends on a huge number of factors. What proposition are we talking about? What intellectual and practical environment does Hannah occupy? How does belief relate to action? Does believing that p entail acting as if p ? Different answers to these and many other questions will make significant differences to the relevant decision table.

At least this much is clear, however: just as the consequences of going straight home depend on whether the bank will be open tomorrow (and many other things), the consequences of forming the belief that p depend on whether p is true, whether Hannah has sufficient evidence that p , and so on. If Hannah forms the belief that p and p is true, or if she forms the belief that p on the basis of good evidence, or if she employs reliable cognitive faculties in the formation of her belief that p , she will get a belief with positive epistemic properties. On the other hand, if Hannah forms the belief that p and p is false, or if she

forms the belief that p in the absence of sufficient evidence that p , or if she employs *unreliable* cognitive faculties in the formation of her belief that p , she will get a belief with negative epistemic properties. And crucially, just as Hannah could stop at the bank in order to safeguard against incurring the large fine, she could withhold in order to safeguard against forming a belief with negative epistemic properties. So now, if we understand intellectual interests in terms of the *values* of epistemic properties, we can understand intellectualism as the view that knowledge depends not just on the presence or absence of epistemic properties, but also (in certain circumstances) on the *values* of these properties—on the fact that some epistemic properties are *good*, some epistemic properties are *bad*, and some epistemic properties are better or worse than others.

Understood this way, intellectualism is structurally analogous to pragmatism and controversial for the same reason. As we will see below, both views make knowledge depend on truth-irrelevant factors. Intellectualism *differs* importantly from pragmatism for at least this reason, however: where pragmatism attempts to trace the interest-relativity of knowledge that p back to the rationality of *acting as if p* , intellectualism will attempt to trace the interest-relativity of knowledge that p back to the rationality of *believing that p* . Since knowing that p might entail rationally believing that p even if it does not entail that one can rationally act as if p , intellectualism is not just a species of pragmatism, and purists cannot simply transfer their arguments against pragmatism over to intellectualism.²¹⁶

²¹⁶ Consider Hannah's belief that the bank will be open in Case B. If purism is right, then Hannah knows that the bank will be open in Case B even while she cannot *rationally act* as if the bank will be open in Case B. But purists would not deny that, if Hannah knows that the bank will be open in Case B, then she *rationally believes* that the bank will be open in Case B. Because purists think the connection between knowledge and rational belief is considerably more intimate than the connection between knowledge and *rational action*, purists must regard intellectualism as importantly different from pragmatism.

2. Three Claims about Rationality

2.1. Mere Evidence of Negative Consequences

To see why knowledge might depend on intellectual interests, we must note three things about rationality. First, consider Case B again. In Case B, Hannah will incur a large fine if she does not deposit her paycheck before Saturday afternoon, and she *knows* that she will incur this large fine if she does not deposit her paycheck before Saturday afternoon. These facts help explain why Hannah cannot rationally go straight home, as she would prefer to conditional on the proposition that the bank will be open. Hannah's inability to rationally go straight home does not depend on these facts, however. This is apparent when we compare Case B to Case B', below.

Case B': Hannah is driving past the bank on Friday afternoon. She has her paycheck in hand and she plans to deposit it, but there are long lines so she considers returning to deposit her paycheck on Saturday morning. On the basis of outstanding evidence, Hannah is convinced that she will incur an enormous fine if she does not deposit her check before Saturday afternoon. But as it turns out, Hannah is wrong about this. The relevant policy changed just moments ago and Hannah's belief that she will incur the large fine is now *false*. Contrary to how things appear to Hannah, the consequences of failing to deposit the check before Saturday afternoon are actually negligible.

Hannah cannot rationally go straight home in Case B' any more than she can rationally go straight home in Case B, since, in both cases, the chances of incurring the fine give her a

reason to stop at the bank that outweighs her reasons for going straight home.²¹⁷ There needn't actually *be* severe consequences for failing to deposit the check before noon; mere *evidence* of severe consequences for failing to deposit the check before noon can make it irrational for her to go straight home. In general: mere evidence of negative consequences for φ -ing can make it irrational to φ , even if those consequences will not actually result from φ -ing.

2.2. Irrational Action without Full Belief

Second, consider another variation on Case B, Case B'' below.

Case B'': Hannah is driving past the bank on Friday afternoon. She has her paycheck in hand and she plans to deposit it, but there are long lines so she considers returning to deposit her paycheck on Saturday morning. At this point Sarah tells Hannah that, according to the bank's policies, Hannah will incur an enormous fine if she does not deposit her check before Saturday afternoon. Hannah does not remember any policy like this, but Sarah is generally reliable, so Hannah starts to worry that she will in fact incur an enormous fine if she does not deposit her check before Saturday afternoon.

Where b is the proposition that the bank will be open on Saturday morning, and f is the proposition that Hannah will incur an enormous fine if she does not deposit her check

²¹⁷ I will occasionally speak of chances instead of epistemic probabilities. As Hawthorne and Stanley point out (2008, p. 582), chance-talk in ordinary language is often epistemic in character. Throughout, by 'chances,' I will mean the relevant epistemic probabilities.

before noon on Saturday, we can represent Hannah's choice between stopping at the bank and going straight home with the following table.

Table 2

	$b \wedge f$	$b \wedge \neg f$	$\neg b \wedge f$	$\neg b \wedge \neg f$
Stop at the bank	<i>Waste a little time</i>	<i>Waste a little time</i>	<i>Avoid enormous fine</i>	<i>Waste a little time</i>
Go straight home	<i>Save a little time</i>	<i>Save a little time</i>	<i>Incur enormous fine</i>	<i>Save a little time</i>

So long as Hannah cannot be highly confident that f is false—so long as she cannot be highly confident that she will *not* incur a large fine by failing to deposit her check before Saturday afternoon—the chances that she *will* incur such a fine give her a reason to stop at the bank that outweighs her reasons for going straight home. Since we are supposing that Hannah *cannot* be highly confident that f is false, we have another case where Hannah cannot rationally go straight home. And as we just saw, whether Hannah will *actually* incur the large fine does not matter. Even if we stipulate that, in Case B'', Hannah will *not* incur the large fine, Hannah cannot rationally go straight home, since she should be worried that she *will* incur the large fine. Generalizing, mere evidence of negative consequences for φ -ing can make it irrational for me to φ , even if that evidence does not justify full belief that there will be negative consequences for φ -ing, and even if those consequences would not actually follow from my φ -ing.

2.3. A Puzzle about Withholding

Third, imagine an urn containing 75 red marbles, 25 green marbles, and nothing else. I know that a random marble was drawn from the urn, and I know that the odds that the draw produced a red marble were 3:1. I also know that the outcome of the draw will never be announced. In fact, I know that nobody will ever know anything about the outcome of the draw other than that the odds of a red marble were 3:1. Now I have three options here: I can believe that the draw produced a red marble, I can believe that the draw *did not* produce a red marble, or I can withhold. Call the proposition that the draw produced a red marble ‘*r*,’ and call the proposition that the draw did not produce a red marble ‘*g*.’ (In the scenario, the draw did not produce a red marble just in case it *did* produce a *green* marble.) Which option should I choose? Intuitively, I should withhold.²¹⁸ But *why*? I cannot be rationally obliged to withhold without sufficient *reason* to withhold, so I must have sufficient reason to withhold. Since I withhold only if I resist both *r* and *g*, I must have sufficient reason to resist *r* and sufficient reason to resist *g*.²¹⁹ My reason for resisting *g* is clear enough: if I were to believe that *g*, my belief would probably be *false*. But this cannot be my reason for resisting *r*. After all, if I were to believe that *r*, my belief would probably be *true*. So why should I resist *r*? Intuitively, even though my belief would probably be true, it is not probable *enough* that it would be true. The probability must be higher than 0.75, and this is why I should resist *r*.

²¹⁸ If you don’t share this intuition, then change the details so that the urn contains 60 red marbles and 40 green ones, or even so that it contains 51 red marbles and 49 green ones. Any percentage will work, so long as there are more red marbles than green marbles, and so long as the percentage of red marbles is too low for me to believe that a red marble was drawn.

²¹⁹ By ‘resist,’ I mean that I either withhold with respect to the relevant proposition or believe the negation of that proposition. Since withholding is inconsistent with believing the negation of the withheld proposition, resisting a proposition is not the same thing as withholding with respect to that proposition. For similar distinctions, see Schroeder (2012), pp. 272-4.

But why should the probability be higher than 0.75? This probability guarantees that it is three times more likely than not that I would form a *true* belief. These are pretty good odds. So again, why should I resist *r*?

The answer to this question lies in an asymmetry between truth and falsehood *vis-à-vis* norms of belief. By ‘norms of belief,’ I mean true propositions of the form ‘it is epistemically appropriate for *S* to believe that *p* only if *S*’s belief that *p* has property φ .’ A belief is epistemically *flawless* (I will say) just in case it satisfies every norm of belief, and epistemically *flawed* just in case it fails at least one norm of belief.²²⁰ A belief is epistemically *perfect* just in case there is no epistemic dimension on which it could be better. Since beliefs can be epistemically appropriate without being epistemically *perfect*, beliefs can be epistemically *flawless* (in my parlance) without being epistemically perfect.

Now, while falsehood guarantees that a belief is epistemically flawed, truth does not guarantee that a belief is epistemically flawless. Beliefs can be epistemically flawed even while true. Paradigmatically, beliefs produced by unreliable processes, or based on insufficient evidence, are epistemically flawed even if true. And plausibly, a belief formed on the basis of nothing better than 3:1 odds would be the product of insufficiently reliable processes, or based on too little evidence. This means that, if I were to believe *r*, the chances that I would thereby form an epistemically flawed belief are quite high. This is the case even though the chances that I would form a *false* belief are quite low. And just as the high chance that I would form a *false* belief gives me sufficient reason to resist *g*, the high chance that I would form an *epistemically flawed* belief gives me sufficient reason to resist *r*. *This* is why I cannot

²²⁰ Cf. Sosa (2010), p. 175.

rationality believe that r . More generally, a high chance that I would form an epistemically flawed belief by believing some proposition can make it irrational for me to believe that proposition.

When we combine this generalization with the points about rationality that I made in §2.1 and §2.2, we arrive at the following insight: evidence that my belief that p would have some epistemic flaw could make it irrational for me to believe that p , even if this evidence does not justify full belief that my belief that p would have this epistemic flaw, and even if my belief that p would actually *not* have this epistemic flaw. And note carefully, the rationality now at issue seems clearly *epistemic*, since a high chance of forming a *false* belief is a paradigm epistemic reason to withhold, and since a high chance of forming an *epistemically flawed* belief seems no less epistemic than a high chance of forming a false belief. This insight will play a crucial role in the argument to follow.²²¹

3. An Argument for Intellectualism

What does epistemically flawless belief require? According to many philosophers, it requires *knowledge*. According to Ernest Sosa, for example, inapt performances are flawed, belief is a species of performance, and knowledge is just apt belief.²²² So now consider the following cases, and suppose that they differ only insofar as the stipulated differences in Hannah's intellectual interests require that they differ.

Case C: After considering the evidence for and against p , Hannah forms the belief that p and thereby acquires an item of knowledge. But while she knows that p , her

²²¹ Following Schroeder (2012), p. 273, I stipulatively define 'epistemic reasons' as "those reasons, whatever they are, which bear on epistemic rationality."

²²² Sosa (2010), p. 175.

epistemic position *vis-à-vis* p is not strong enough for her to know that she knows that p .

Case D: Hannah is considering the question whether p , just as in Case C. But in this case, Hannah is an epistemologist interested in norms of belief. In Case C, Hannah was oblivious to the chances of forming an epistemically flawed belief. But in this case these chances are salient, since Hannah sees that Sosa and others might be right that beliefs that fall short of knowledge are epistemically flawed. Hannah knows there is a high chance that, if she were to form the belief that p , she would thereby form an epistemically flawed belief. Since she knows that epistemically flawed beliefs are epistemically inappropriate, Hannah sees that it would be epistemically better to have no belief than an epistemically flawed belief.

What should we make of these cases? Case C stipulates that Hannah knows that p , and I do not have any strong intuitive response to Case D. Still, there are excellent reasons to think that, while Hannah knows that p in Case C, she cannot know that p in Case D.

In Case C, Hannah is not in position to know whether her belief is an item of knowledge. Case D differs from Case C only insofar as the stipulated difference in Hannah's intellectual interests requires that they differ. So, in Case D, Hannah is not in position to know whether her belief would be an item of knowledge. But in Case D, Hannah thinks that beliefs that fall short of knowledge might be epistemically flawed. She does not just *happen* to think this. Her belief is perfectly rational, since her evidence strongly supports this conclusion.²²³ As a result,

²²³ See Williamson (2000), Adler (2002), Engel (2005), Bird (2007), Huemer (2007), Sutton (2007), Bach (2008), Sosa (2010), McHugh (2011), and Smithies (2012) for considerations suggesting that beliefs that fall short of knowledge are epistemically flawed. Aidan McGlynn (forthcoming) has a nice a nice argument that beliefs can

Hannah is reasonably worried that, if she were to form the belief that p , her belief would be epistemically flawed. This means that Hannah's choice between believing that p and withholding in Case D is analogous to her choice between stopping at the bank and going straight home in Case B''. Conditional on Hannah's evidence in Case B'', there is a decent chance that Hannah will incur a large fine if she fails to deposit her check before noon on Saturday. Similarly, conditional on Hannah's evidence in Case D, there is a decent chance that Hannah would form an epistemically flawed belief if she were to believe that p . We can see the relevant parallels between Case B'' and Case D by drawing up the decision table for Hannah's choice in Case D. Where ' f ' names the proposition that Hannah's belief would be epistemically flawed for falling short of knowledge, we can represent Hannah's choice between believing that p and withholding with Table 3, below.

Table 3

	$p \wedge f$	$p \wedge \neg f$	$\neg p \wedge f$	$\neg p \wedge \neg f$
Believe that p	<i>Form belief that is epistemically flawed for falling short of knowledge</i>	<i>Form epistemically flawless belief</i>	<i>Form belief that is epistemically flawed for being false and for failing some other requirement</i>	<i>Form belief that is epistemically flawed for being false</i>

fall short of knowledge without being epistemically flawed, but we may suppose that Hannah has only read the literature already in print and is therefore not yet aware of arguments like McGynn's.

		<i>for knowledge</i>	
Withhold	<i>Form no belief</i>	<i>Form no belief</i>	<i>Form no belief</i>

Just as Table 1 represents the only relevant difference between Case A and Case B, and just as Table 2 represents the only relevant difference between Case A and Case B', Table 3 represents the only relevant difference between Case C and Case D. But we could fill in the details of Case D so that, just as the chances of incurring the large fine give Hannah a reason to stop at the bank that outweighs her reasons for going straight home, the chances of forming an epistemically flawed belief give Hannah a reason to withhold that outweighs her reasons to believe that *p*. So let's assume that the chances of forming an epistemically flawed belief *do* give Hannah a reason to withhold that outweighs her reasons to believe. In this case, Hannah cannot *rationaly* believe that *p*. And note, the rationality at issue here is *epistemic*, not merely practical.²²⁴ Since knowledge entails epistemically rational belief, it follows that Hannah does not *know* that *p*, even if she believes that *p*. Since Hannah knows that *p* in Case C, and since Case D differs from Case C only insofar as the stipulated differences in Hannah's intellectual interests require that they differ, it follows that intellectual interests can make a difference to knowledge. This is just what intellectualism claims, so intellectualism is true.

Putting this reasoning formally, and making it explicit that epistemic reasons and rationality are at issue here, my argument for intellectualism goes as follows.²²⁵

²²⁴ Or so I argue below.

²²⁵ Again, following Schroeder (2012), p. 273, I stipulatively define 'epistemic reasons' as "those reasons, whatever they are, which bear on epistemic rationality."

- (1) If S 's epistemic reasons for withholding with respect to p outweigh her epistemic reasons for believing that p , then it is not epistemically rational for S to believe that p .
- (2) In Case D, Hannah's epistemic reasons for withholding with respect to p outweigh her epistemic reasons for believing that p .
- ∴ (3) In Case D, it is not epistemically rational for Hannah to believe that p . [from (1) and (2)]
- (4) If it is not epistemically rational for S to believe that p , then S does not know that p .
- ∴ (5) In Case D, Hannah does not know that p . [from (3) and (4)]
- (6) In Case C, Hannah *does* know that p .
- (7) Case C differs from Case D only insofar as the stipulated differences in Hannah's intellectual interests require that they differ.
- ∴ (8) A difference in knowledge might follow from as little as a difference in intellectual interests. [from (5), (6) and (7)]
- (9) If a difference in knowledge might follow from as little as a difference in intellectual interests, then intellectualism is true.
- ∴ (10) Intellectualism is true. [from (8) and (9)]

But of course, if intellectualism is true, then so is interest-relativism, since intellectualism is a species of interest-relativism. Since intellectualism is true, knowledge depends on truth-irrelevant factors, even if pragmatism is false.

4. Objections and Replies

What should we make of this argument? It is clearly valid, so we must resist one or more of its premises if we want to resist its conclusion. Premise (7) just tells us how to understand Cases C and D, and premise (9) is true by our stipulative definition of ‘intellectualism,’ so we must take (7) and (9) for granted. Moreover, intellectualism seems considerably more plausible than the claim that Hannah might know that p even while it would be epistemically irrational for her to believe that p , so I will also take (4) for granted. Since the argument is valid, rejecting (3), (5) or (8) requires rejecting one of the premises that jointly entail them. This leaves us with (1), (2), and (6). By my lights, the following objections constitute the most substantial reasons to deny (1), (2) or (6).

Objection 1: Suppose a Cartesian demon gives you the following credible threat: unless you withhold with respect to the proposition that you have hands, you will only form epistemically flawed beliefs from here forward.²²⁶ In this case, your epistemic reasons for withholding with respect to the proposition that you have hands would easily outweigh your epistemic reasons for believing that you have hands, so (1) says that it is not epistemically rational for you to continue believing that you have hands. Since it clearly *is* epistemically rational for you to continue believing that you have hands, (1) is therefore false.

Reply to Objection 1: Perhaps there is some sense in which the demon gives me reasons for withholding with respect to the proposition that I have hands that outweigh my reasons for believing that I have hands, but it is doubtful that the demon gives me genuinely *epistemic* reasons for withholding. In fact, there is good reason to insist that these reasons could not

²²⁶ Richard Fumerton considers a case somewhat like this at Fumerton (2001), p. 55. (Thanks to Kurt Sylvan for pointing out the similarity between this case and Fumerton’s.)

be genuinely epistemic. Following Schroeder, I have stipulatively defined epistemic reasons as those reasons (whatever they are) that bear on epistemic rationality. So the fact that it would be epistemically rational for me to continue believing that I have hands strongly suggests that my reasons for withholding with respect to this proposition are not genuinely epistemic.

Objection 2: Okay, the demon in Objection 1 does not give you any genuinely *epistemic* reason to withhold. But the demon does give you *some* reason to withhold, and this reason for withholding is the same *kind* of reason as Hannah's reason for withholding in Case D. So Hannah's reason for withholding in Case D is not genuinely epistemic either. This means that (2) is false.

Reply to Objection 2: In the urn case from §2.3, my reason for withholding is that there is a high chance that I would form an epistemically flawed belief by believing the proposition in question. But this is exactly Hannah's reason for withholding in Case D. Since my reason for withholding in the urn case is genuinely epistemic, Hannah's reason for withholding in Case D is also genuinely epistemic. With respect to epistemic reasons and rationality, there is a clear difference between withholding with respect to p because *your belief that p* might have some negative epistemic status, and withholding with respect to p because *some unrelated future beliefs* might have some negative epistemic status. My withholding in the urn case, and Hannah's withholding in Case D, would both fall on the former side of this divide, while my withholding in the demon case would fall on the latter side of this divide. So the demon Case from Objection 1 causes no problems for premises (1) or (2).

This is one way we might respond to Objection 2, at least. More cautiously, we might respond by posing a challenge for the proponent of Objection 2. *If* my reasons for withholding in the urn case are genuinely epistemic, and *if* there is no relevant difference between my reasons for withholding in the urn case and Hannah's reasons for withholding in Case D, *then* Hannah's reasons for withholding in Case D are also genuinely epistemic. This much is clear. So the proponent of Objection 2 must either show that my reasons for withholding in the urn case are *not* genuinely epistemic, or show that my reasons for withholding in the urn case differ relevantly from Hannah's reasons for withholding in Case D. The latter option seems hopeless. The only plausible difference *at all* between my reasons for withholding in the urn case and Hannah's reasons for withholding in Case D is a difference *in strength*; the chances that I would form an epistemically flawed belief by believing the relevant proposition are plausibly higher than the chances that Hannah would form an epistemically flawed belief by believing the relevant proposition. (Note that, for all we have said, the relevant proposition is the *same* proposition in both cases.) But there is clearly no difference *in kind* between my reasons for withholding in the urn case and Hannah's reasons for withholding in Case D. Again, Hannah's reason for withholding in Case D is *identical* to my reason for withholding in the urn case. So it seems clear that, if my reasons for withholding in the urn case are genuinely epistemic, then Hannah's reasons for withholding in Case D are also genuinely epistemic. The proponent of Objection 2 must therefore show that my reasons for withholding in the urn case are *not* genuinely epistemic.

But it is unclear how the proponent of Objection 2 could show this. If my reasons for withholding in the urn case are not genuinely epistemic, then it seems doubtful that I have

any *epistemic* obligation to withhold in the urn case. Yet if I have no epistemic obligation to withhold in the urn case, then either I have no obligation to withhold in the urn case *at all*, or epistemic *rationality* has nothing to do with my obligation to withhold in the urn case. Since I clearly *am* obliged to withhold in the urn case, the proponent of Objection 2 must say that epistemic rationality has nothing to do with my obligation to withhold in the urn case. But the claim that epistemic rationality has nothing to do with my obligation to withhold in the urn case seems patently false, for it seems that I am obliged to withhold in the urn case precisely because it would be epistemically irrational for me to do otherwise. So Objection 2 looks like a dead end.²²⁷

Objection 3: Okay, let's take it for granted that Hannah's reasons for withholding in Case D *are* genuinely epistemic, since we run into a serious difficulties by denying this claim. Given this assumption, Objection 2 fails. But premise (2) still comes out false, since Hannah's reasons for withholding with respect to *p* actually *do not* outweigh her reasons for believing that *p*. As Table 3 makes clear, whether Hannah's reasons for withholding outweigh her reasons for believing depends on the value of epistemically flawed belief, the value of epistemically flawless belief, and the value of withholding. If the positive value of epistemically flawless belief significantly outweighs the negative value of epistemically flawed belief, or if withholding *itself* has negative value, then Hannah's reasons for withholding with respect to *p* *will not* outweigh her reasons for believing that *p*, and premise (2) will come out false. But the positive value of epistemically flawless belief *does* significantly outweigh the

²²⁷ See Schroeder (2012) for an account of epistemic reasons and rationality that fits very nicely with premise (2) and thereby causes substantial obstacles for Objection 2.

negative value of epistemically flawed belief, and withholding *does* have negative value. So even if Hannah's reasons for withholding in Case D *are* genuinely epistemic, (2) is still false.

Reply to Objection 3: The first thing to note in response to this objection is that, even if withholding *does* have negative value, it does not follow that Hannah's reasons for withholding do not outweigh her reasons for believing. After all, if the negative value of epistemically flawed belief significantly outweighs the positive value of epistemically flawless belief, Hannah's reasons for withholding could outweigh her reasons for believing *even if* withholding has negative value. So it is simply false that, if withholding itself has negative value, then Hannah's reasons for withholding will not outweigh her reasons for believing.

The second thing to note in response to this objection is that, while it might be *practically* very good for Hannah to believe some proposition and *practically* very bad for her to withhold with respect to that proposition (perhaps because believing that proposition will give Hannah confidence to face the day), the value we are interested in here is distinctively *epistemic* value. This is important because it seems doubtful that the positive *epistemic* value of epistemically flawless belief significantly outweighs the negative epistemic value of epistemically flawed belief. Remember, epistemically flawless beliefs are just beliefs that satisfy every norm of belief. Plausibly, it's rather easy to satisfy every norm of belief, even if *knowledge* is a norm of belief. But if epistemically flawless belief is easy to acquire, then why shouldn't the negative value of epistemically flawed belief outweigh the positive value of epistemically flawless belief, as I am tempted to say? At the very least, it seems much more plausible that the positive value of epistemically flawless belief is *exactly proportionate* to the

negative value of epistemically flawed belief than it is that the former significantly outweighs the latter. But this means that Objection 3 relies on an implausible assumption.

Objection 4: The problem with (2) is not that Hannah's epistemic reasons for withholding with respect to p do not outweigh her epistemic reasons for believing that p , as Objections 2 and 3 suggest. Rather, the problem is that epistemic reasons to withhold do not exist. Withholding is a kind of default state. There are no epistemic reasons to be in this default state. There are only epistemic reasons to get out of it, either by believing the relevant proposition or by believing its negation. Epistemic rationality can oblige you to withhold, of course. This happens whenever your epistemic reasons are too weak to support full belief. But this is the whole story about the epistemic rationality of withholding, and epistemic reasons to withhold are no part of it. Since epistemic reasons for *belief* exhaust the epistemic reasons there are, (2) is false.

Reply to Objection 4: Conditional on Hannah's evidence in Case D, there is a high chance that she would form an epistemically flawed belief by believing that p . Hannah knows this, so Hannah has an epistemic reason to resist p . But this epistemic reason to resist p is not an epistemic reason to believe that $\neg p$. If it were, then Hannah's credence that p in Case D should be lower than her credence that p in Case C, since Hannah has this epistemic reason to resist p in Case D but lacks this reason to resist p in Case C. So there are epistemic reasons to resist p that are not epistemic reasons to believe that $\neg p$.²²⁸ Are these reasons epistemic reasons to *withhold* with respect to p , then? I think they must be, but for present purposes it does not matter. If the answer is 'no,' we can simply replace premise (1) with the principle

²²⁸ Cf. Ibid, p. 273.

that it is epistemically irrational for S to believe that p if her epistemic reasons for resisting p outweigh her epistemic reasons for believing that p , and replace premise (2) with the proposition that, in Case D, Hannah's epistemic reasons for resisting p outweigh her epistemic reasons for believing that p . Instead of modifying the argument, however, I will stipulatively define 'reasons to withhold' as follows: R is a reason for S to withhold with respect to p just in case (a) R is a reason for S to resist p , and (b) R is not a reason for S to believe that $\neg p$. With this stipulative definition, we can leave (1) and (2) untouched.

Objection 5: This attention to premises (1) and (2) is actually misguided, since the *real* problem with the argument for intellectualism resides at premise (6). In Case D, Hannah does not know that p , and two facts are responsible for her ignorance here: the fact that she is not in position to know whether she knows that p , and the fact that, plausibly, beliefs that fall short of knowledge are epistemically flawed. In Case D, Hannah is *aware* of these facts, and this is why she cannot *rationally believe* that p . But the deeper reason why Hannah fails to *know* that p is simply that these facts *obtain*. By simply obtaining, these facts defeat Hannah's belief, whether she is aware of them or not. This is important, since these facts *also obtain in Case C*. So the very facts that prevent Hannah from rationally believing that p in Case D prevent her from knowing that p in Case C. Since (6) says that Hannah *does* know that p in Case C, (6) is false.²²⁹

Reply to Objection 5: According to this objection, the fact that Hannah is not in position to know whether she knows that p , and the fact that beliefs that fall short of knowledge are plausibly epistemically flawed, together entail that Hannah does not know that p . Since there

²²⁹ Cf. Schroeder's take on Gettier cases at *ibid*, pp. 269-72.

is nothing special about *Hannab* (she could be anybody), Objection 5 commits us to (C), below.

- (C) If *S* is not in position to know whether she knows that *p*, and if it is plausible that beliefs that fall short of knowledge are epistemically flawed, then *S* does not know that *p*.

But as the literature on norms of belief makes clear, it *is* plausible that beliefs that fall short of knowledge are epistemically flawed.²³⁰ So Objection 5 commits us to (C'), below:

- (C') If *S* is not in position to know whether she knows that *p*, then *S* does not know that *p*.

But (C') entails the dubious KK principle, according to which *S* knows that *p* only if she is in position to know that she knows that *p*. This means that Objection 5 commits us to the KK principle. There are powerful arguments against the KK principle, and many of the most plausible theories of knowledge predict that it fails.²³¹ Instead of rehearsing these arguments or getting into the details of any theory of knowledge, I will simply voice my conviction that the KK principle is false and invite those who share my conviction to consider adopting intellectualism.

²³⁰ Again, See Williamson (2000), Adler (2002), Engel (2005), Bird (2007), Huemer (2007), Sutton (2007), Bach (2008), Sosa (2010), McHugh (2011), and Smithies (2012) for considerations suggesting that beliefs that fall short of knowledge are epistemically flawed. These considerations may not *establish* that beliefs that fall short of knowledge are epistemically flawed. But at the very least, these considerations show that it is *plausible* that beliefs that fall short of knowledge are epistemically flawed.

²³¹ Regarding the first point, see Williamson's anti-luminosity argument and his application of this argument to the KK principle at Williamson (2000), chs. 4 and 5. Regarding the second, note that virtually *every* theory of knowledge predicts the possibility of a case where *S*'s belief that *p* satisfies all the necessary conditions for knowledge while her belief that she *knows* that *p* would fail at least one of them.

Objection 6: The argument for intellectualism assumes that a choice between believing and withholding could be relevantly analogous to a choice between stopping at the bank and going straight home. But this assumption is false, since we have a lot of control over choices like stopping at the bank, but very little control over our beliefs. The argument for intellectualism therefore rests on a false assumption.

Reply to Objection 6: The argument for intellectualism is valid and thus sound unless one of its premises is false, so which premise rests on the false assumption that we have as much control over our beliefs as we have over (*e.g.*) whether we stop at the bank? The only candidates are premises (1), (2), and (6), and it looks like they could all be true even if we have considerably *less* control over our beliefs than we have over choices like stopping at the bank. So it seems doubtful that the argument for intellectualism rests on the false assumption that we have as much control over our beliefs as we have over choices like stopping at the bank.

Objection 7: Okay, so intellectualism is true. And since intellectualism is a species of interest-relativism, interest-relativism is true too. But so what? Haven't epistemologists been intellectualists all along? Consider Case C and Case D. By hypothesis, they differ only insofar as the stipulated differences in Hannah's intellectual interests require that they differ. But as a result of these stipulated differences, Hannah has a genuinely epistemic reason to withhold in Case D that she lacks in Case C. Since epistemologists have *never* denied that a difference in genuinely epistemic reasons to withhold can make a difference to knowledge, intellectualism is just the traditional view.

Reply to Objection 7: Even though Hannah has a genuinely epistemic reason to withhold in Case D that she lacks in Case C, Case D is identical to Case C on every *truth-relevant* dimension. As we noted in response to Objection 4, Hannah's reason to withhold with respect to p is no reason (epistemic or otherwise) to believe that $\neg p$. More generally, Hannah's reason to withhold with respect to p makes no difference to the truth of p or the probability that p , to Hannah's first-order or higher-order evidence that p , to the reliability of the faculties that she might rely on in believing that p , to the counterfactual relations that would hold between Hannah and the truth of her belief that p , and so on. On all of these dimensions, Case D is identical to Case C. But the traditional view says that a reason can make a difference to knowledge only if it makes a difference to one of these truth-relevant dimensions. To put the same point in terms of epistemic position, by acquiring this reason to withhold, Hannah goes from an epistemic position that is strong enough for knowledge, to an epistemic position that is too weak for knowledge, without any change in the strength of her epistemic position. But the traditional view says that one *cannot* go from a sufficiently strong epistemic position to an insufficiently strong epistemic position without some change in the strength of one's epistemic position. So intellectualism is not the traditional view. If it is true, then the traditional view is mistaken. Since I see no promising way of resisting my argument for intellectualism, I conclude that the traditional view *is* mistaken.

5. Concluding Remarks

Intellectualism is interesting in its own right, since it is an interesting question how knowledge relates to truth. Intellectualism is especially interesting against the backdrop of the pragmatism/purism debate, however. The motivation for resisting pragmatism is

supposed to be *precisely* that it makes knowledge depend on truth-irrelevant factors—that, if it is true, one can go from an epistemic position that is strong enough for knowledge to an epistemic position that is too weak for knowledge without any change in the strength of one’s epistemic position. But if intellectualism is true, then one *can* go from an epistemic position that is strong enough for knowledge to an epistemic position that is too weak for knowledge without any change in the strength of one’s epistemic position. So if intellectualism is true, there seems little reason for resisting pragmatism. In addition to providing an interesting account of the relationship between knowledge and truth, intellectualism therefore has important implications for the debate between purists and pragmatists. The upshot is that purists must find some motivation for resisting pragmatism other than that it makes knowledge depend on truth-irrelevant factors.²³²

²³² Thanks to Tomas Bogardus, Daniel Fogal, Alvin Goldman, Jennifer Lackey, Lisa Miracchi, Kate Nolfi, Carlotta Pavese, Ernest Sosa, Jason Stanley, Kurt Sylvan, and Jenn Wang for helpful comments and conversation.

References

- Adler, Jonathan (2002), *Beliefs Own Ethics* (Cambridge, MA: MIT Press).
- Bird, Alexander (2007), "Justified Judging," *Philosophy and Phenomenological Research*, Vol. 74, No. 1: pp. 81-110.
- Bach, Kent (2006). "The Top 10 Misconceptions about Implicature," in *Drawing the Boundaries of Meaning*, Betty Birner and Gregory Ward, Eds. (Amsterdam: John Benjamins Publishing), pp. 21-30.
- (2008), "Applying Pragmatics to Epistemology," *Philosophical Issues*, Vol. 18: pp. 69-88.
- Brown, Jessica (2006). "Contextualism and Warranted Assertability Manoeuvres," *Philosophical Studies*, Vol. 130: pp. 407-35.
- (2008), "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning," *Noûs*, Vol. 42, No. 2: pp. 167-89.
- (2012), "Practical Reasoning, Decision Theory, and Anti-Intellectualism," *Episteme*, Vol. 9, No. 1: pp. 43-62.
- (Forthcoming), "Experimental Philosophy, Contextualism, and SSI," *Philosophy and Phenomenological Research*.
- (forthcoming), "Impurism, Practical Reasoning, and the Threshold Problem," *Noûs*.

- Buckwalter, Wesley (2010), "Knowledge is not Closed on Saturday: A Study in Ordinary Language," *Review of Philosophy and Psychology*, Vol. 1, No. 3: pp. 395–406.
- Buckwalter, Wesley and Jonathan Schaffer (forthcoming), "Knowledge, Stakes and Mistakes," *Nous*. DOI: 10.1111/nous.12017.
- Christensen, David. "Dutch-Book Arguments Depragmatized: Epistemic Consistency for Partial Believers," *The Journal of Philosophy*, Vol. 93, No. 9 (1996): 450-79.
- Cohen, Stewart (1988). "How to be a Fallibilist," *Philosophical Perspectives*, Vol. 2, Epistemology: pp. 91-123.
- (1999), "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives*, Vol. 13, Epistemology: pp. 57-89.
- (2012), "Does Practical Rationality Constrain Epistemic Rationality?," *Philosophy and Phenomenological Research*, Vol. 85, No. 2.
- Davis, Wayne (2007). "Knowledge Claims and Context: Loose Use," *Philosophical Studies*, Vol. 132: pp. 395-438.
- DeRose, Keith (1992). "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research*, Vol. 52: pp. 913-29.
- (1995). "Solving the Skeptical Problem," *Philosophical Review*, Vol. 104: pp. 1-52.
- (2002). "Assertion, Knowledge and Context," *Philosophical Review*, Vol. 111: pp. 167-203.
- (2009). *The Case for Contextualism* (Oxford: Oxford University Press).

- Engel, Pascal (2005), "Truth and the Aim of Belief," in *Laws and Models in Science*, D. Gillies ed. (London: King's College Publications).
- Fantl, Jeremy and Matthew McGrath (2002). "Evidence, Pragmatics and Justification," *The Philosophical Review*, Vol. 111: pp. 67-94.
- (2009), *Knowledge in an Uncertain World* (Oxford: Oxford University Press).
- Feltz, Adam, and Chris Zarpentine (2010), "Do you Know More When it Matters Less?," *Philosophical Psychology*, Vol. 23, No. 5: pp. 683-706.
- Fumerton, Richard (2001), "Epistemic Justification and Normativity," in *Knowledge, Truth, and Duty*, M. Steup ed. (Oxford: Oxford University Press).
- (2010), "Fencing Out Pragmatic Encroachment," in *Philosophical Perspectives*, Volume 24, Epistemology: pp. 243-53.
- Ganson, Dorit (2008). "Evidentialism and Pragmatic Constraints on Outright Belief," *Philosophical Studies*, Vol. 139: pp. 441-58.
- Grice, Paul (1989). *Studies in the Ways of Words* (Cambridge, MA: Harvard University Press).
- Grimm, Stephen (2011), "On Intellectualism in Epistemology," *Mind*, Vol. 120, No. 479: pp. 705-733.
- Hawthorn, John (2004). *Knowledge and Lotteries* (Oxford: Oxford University Press).
- Hawthorne, John and Jason Stanley (2008), "Knowledge and Action," *The Journal of Philosophy*, Vol. 105, No. 10.

- Huemer, Michael (2007), "Moore's Paradox and the Norm of Belief," in *Themes from G.E. Moore: New Essays in Epistemology and Ethics*, S. Nuccetelli and G. Seay eds. (Oxford: Oxford University Press).
- Kruglanski, Arie (2004), *The Psychology of Closed Mindedness* (New York: Taylor and Francis).
- Lackey, Jennifer (2008), "What Luck is Not," *Australasian Journal of Philosophy*, Vol. 86, No. 2: 255-67.
- (2011), "Assertion and Isolated Secondhand Knowledge," in *Assertion*, eds. J. Brown and H. Cappelen (Oxford: Oxford University Press), pp. 251-76.
- Locke, Dustin (forthcoming), "Practical Certainty," *Philosophy and Phenomenological Research*.
- May, Joshua, Walter Sinnott-Armstrong, Jay Hull and Aaron Zimmerman (2010), "Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study," *Review of Philosophy and Psychology*, Vol. 1, No. 2: pp. 265–73.
- McGlynn, Aidan (forthcoming), "Believing Things Unknown," *Noûs*. DOI: 10.1111/j.1468-0068.2011.00843.x.
- McHugh, Conor (2011), "What Do We Aim At When We Believe?," *Dialectica*, Vol. 65: pp. 327-82.
- Nagel, Jennifer (2008), "Knowledge Ascriptions and the Psychological Consequences of Changing Stakes," *Australasian Journal of Philosophy*, Vol. 86, No. 2: 279-93.
- Neta, Ram (2007), "Anti-intellectualism and the Knowledge-Action Principle," *Philosophy and Phenomenological Research*, Vol. 75, No. 1: pp. 180-7.

- (2012), “The Case Against Purity,” *Philosophy and Phenomenological Research*, Vol. 85, No. 2.
- Pinillos, Angel (2012), “Knowledge, Experiments and Practical Interests,” in *New Essays on Knowledge Ascriptions*, eds. J. Brown and M. Gerken (Oxford: Oxford University Press).
- Pritchard, Duncan (2005), *Epistemic Luck* (Oxford: Oxford University Press).
- Reed, Baron (2010), “A Defense of Stable Invariantism,” *Noûs*, Vol. 44, No. 2: pp. 224-44.
- (2012), “Resisting Encroachment,” *Philosophy and Phenomenological Research*, Vol. 85, No. 2.
- Ross, Jacob and Mark Schroeder (forthcoming), “Belief, Credence and Pragmatic Encroachment,” *Philosophy and Phenomenological Research*.
- Rysiew, Patrick (2001). “The Context Sensitivity of Knowledge Attributions,” *Noûs*, Vol. 35: pp. 477-514.
- Schaffer, Jonathon and Joshua Knobe (2013), “Contrastive Knowledge Surveyed,” *Noûs*, Vol. 46, No. 4: pp 675-708.
- Schroeder, Mark (forthcoming), “Stakes, Withholding, and Pragmatic Encroachment on Knowledge,” *Philosophical Studies*. DOI: 10.1007/s11098-011-9718-1.
- Smithies, Declan (2012), “The Normative Role of Knowledge,” *Noûs*, Vol. 46, No. 2: pp. 265-88.
- Sosa, Ernest (2007), *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1* (Oxford: Oxford University Press).

- (2010), “Value Matters in Epistemology,” *The Journal of Philosophy*, Vol. 107, No. 4, pp. 167-90. Reprinted in Ernest Sosa (2011), *Knowing Full Well* (Princeton: Princeton University Press), ch. 3.
- (2011), *Knowing Full Well* (Princeton: Princeton University Press, 2011).
- Sripada, Chandra Sekhar and Jason Stanley (forthcoming), “Empirical Tests of Interest-Relative Invariantism,” *Episteme*.
- Stanley, Jason (2005). *Knowledge and Practical Interests* (Oxford: Oxford University Press).
- Stanley, Jason and Jeffrey C. King (2007). “Semantics, Pragmatics, and the Role of Semantic Content,” in *Language in Context* (Oxford: Oxford University Press), pp. 133-81.
- Weatherson, Brian (2005), “Can We Do without Pragmatic Encroachment?,” *Philosophical Perspectives*, Vol. 19, Epistemology, pp. 417-43.
- (forthcoming), “Knowledge, Bets and Interests,” in *New Essays on Knowledge Ascriptions*, eds. J. Brown and M. Gerken (Oxford: Oxford University Press).
- Weiner, Matt (2007), “Norms of Assertion,” *Philosophy Compass*, Vol. 2, No. 2: 187-94.
- Williamson, Timothy (2000), *Knowledge and its Limits* (Oxford: Oxford University Press).