

Stakes, withholding, and pragmatic encroachment on knowledge

Mark Schroeder

Published online: 8 April 2011
© Springer Science+Business Media B.V. 2011

Abstract Several authors have recently endorsed the thesis that there is what has been called pragmatic encroachment on knowledge—in other words, that two people who are in the same situation with respect to truth-related factors may differ in whether they know something, due to a difference in their practical circumstances. This paper aims not to defend this thesis, but to explore how it could be true. What I aim to do, is to show how practical factors could play a role in defeating knowledge by defeating epistemic rationality—the very kind of rationality that is entailed by knowledge, and in which Pascalian considerations do not play any role—even though epistemic rationality consists in having adequate evidence.

Keywords Epistemology · Knowledge · Belief · Pragmatic encroachment · Rationality · Reasons · John Hawthorne · Jason Stanley · Jeremy Fantl · Matthew McGrath

1 Introduction

Several authors have recently defended the thesis that whether someone knows something can depend on actual or perceived *practical* factors of her situation, in addition to whether she believes it, the reasons for which she believes it, the available evidence, the truth-conduciveness of her environment, whether her faculties are reliable, and such other mundane factors on which knowledge is ordinarily allowed to depend. Call this the *pragmatic encroachment* thesis.¹

¹ The term ‘pragmatic encroachment’ is apparently originally due to Jonathan Kvanvig, in a blog post on June 12, 2004 entitled ‘Pragmatic aspects of knowledge?’. See http://el-prod.baylor.edu/certain_doubts/.

M. Schroeder (✉)
Mudd Hall of Philosophy, University of Southern California, 3709 Trousdale Pkwy, Los Angeles,
CA 90089, USA
e-mail: maschroe@usc.edu

According to John Hawthorne (2004) the pragmatic encroachment thesis can be used to solve certain puzzles deriving from thinking about lottery cases. According to Jason Stanley (2005) it is directly motivated by intuitions about cases. And according to Jeremy Fantl and Matthew McGrath (2010), it is the best way to make sense of the fallibility of knowledge.

This paper also explores the pragmatic encroachment thesis, and also as a sympathetic voice. But rather than arguing that the pragmatic encroachment thesis is true, my primary task will be to try to make sense of how it *could* be true. For the most obvious problem with the pragmatic encroachment thesis is that it *doesn't* seem like the kind of thing that could be true. We're all brought up in epistemology on the foil of Pascal's Wager, which shows how it might be advisable or beneficial to have some belief, independently of whether it is true. But Pascalian considerations, we observe, can't ground *knowledge*. Moreover, we observe, if there is any sense of 'rational' in which you don't know something unless it is rational for you to believe it, Pascalian considerations can't affect this kind of rationality—usually called *epistemic* rationality—either. If Pascalian considerations are our paradigm of practical factors, then it is puzzling in the extreme how practical factors could affect knowledge.

It is this challenge that I aim to address in this paper. What I aim to do, is to show how practical factors could play a role in defeating knowledge by defeating epistemic rationality—the very kind of rationality that is entailed by knowledge, and in which Pascalian considerations do not play any role. I won't quite be arguing that practical factors *do* defeat knowledge, because I won't quite be arguing that practical factors do defeat epistemic rationality. But I will be presenting a clear picture of *how* this could be the case, under a certain condition, and trying to show that this picture leads to detailed and attractive predictions, while departing in relatively minimal ways from orthodoxy.

2 Pragmatic encroachment

So what kind of thing do I have in mind, in saying that practical factors can affect whether someone knows? I am thinking, first and foremost, of cases like the following:

Low Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah remembers the bank being open on Saturday morning a few weeks ago, so she says, 'Fortunately, it will be open tomorrow, so we can just come back.' In fact, Hannah is right—the bank will be open on Saturday.

High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since their mortgage payment is due on Sunday, they have very

little in their account, and they are on the brink of foreclosure, it is very important that they deposit their paychecks by Saturday. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah remembers the bank being open on Saturday morning a few weeks ago, so she says, ‘Fortunately, it will be open tomorrow, so we can just come back.’ In fact, Hannah is right—the bank will be open on Saturday.²

These two cases are designed to differ on the basis of a practical feature of the situation—in this case, Hannah and Sarah’s impending mortgage payment. Intuitively, whether Hannah and Sarah have impending bills has nothing to do with whether the bank will be open on Saturday—it is not as if the bank manager is out to get them, or banks tend to close right before mortgage payments are due, so their impending mortgage payment doesn’t weigh as evidence either way, in terms of whether the bank will be open on Saturday. Nor does it affect whether any of their other evidence is reliable. If it makes a difference in whether they know, therefore, it would seem that it must make a difference simply because it makes the question of whether the bank will be open on Saturday a more pressing question.

The most direct way to motivate the pragmatic encroachment thesis on the basis of cases like Low Stakes and High Stakes, is on the basis of the intuitive judgment that in Low Stakes, Hannah knows that the bank will be open on Saturday, but she does not know this in High Stakes. But the pragmatic encroachment thesis can be motivated even if you don’t share these intuitive judgments. For example, even if you think that Hannah knows in both cases, or that she knows in neither, you will see the allure of the pragmatic encroachment thesis, if it seems to you that it would take *more evidence* to ground knowledge that the bank will be open on Saturday in the High Stakes case than in the Low Stakes case.

Let me head off from the beginning one possible misinterpretation of what is significant about pragmatic encroachment. This is the observation that Hannah and Sarah have *distinctive* practical circumstances that other people don’t share—because theirs is the mortgage payment that is coming due, which makes the question as to whether the bank will be open on Saturday more pressing for *them* than for others. This happens to be true about the cases as described, but I take it to be completely incidental. I could just as well have illustrated pragmatic encroachment with a moral example, in which the topic of the belief in question is equally pressing for everyone. For example, imagine a pair of cases in which Maria is the only doctor around, and Hannah has a belief about where Maria is, based on experience with Maria’s ordinary habits—and in fact Hannah is right, because today is a habitual day for Maria. In the low stakes version of this case we may imagine that there is no special need for a doctor, while in the high stakes version of the case we may imagine that Hannah has encountered a critically injured stranger (fill in the details). This pair of cases raises the same issues about pragmatic

² These two cases are adapted from Stanley (2005). In Stanley’s original cases, the subjects self-ascribe knowledge, which creates a complication for whether this is best interpreted in a contextualist framework. I have eliminated this complication from my version of the cases.

encroachment, but Hannah's practical circumstances would be shared by anyone else who encountered the injured stranger.³

3 Knowledge and epistemic rationality

Although the pragmatic encroachment thesis is first and foremost a thesis about *knowledge*, I am not going to provide a picture of how practical factors defeat knowledge *directly*. Rather, what I am going to focus on for most of this paper, is how Hannah's belief that the stakes are high in her situation could make it fail to be *rational* for her to believe that the bank will be open on Saturday. In doing so, I will be relying on what I take to be a highly plausible assumption about how knowledge and rationality are related to one another: in particular, I will be assuming that you know p only if it is rational for you to believe p .

It is important to be careful, however, how to interpret this thesis. Suppose, for example, that a villain offers a very credible threat to kill your family if you believe that $2 + 2 = 4$. Under such circumstances, if the threat is really credible, it would be irrational for you not to do everything in your power to not believe that $2 + 2 = 4$. It may even be 'irrational', in *some* sense, for you to believe that $2 + 2 = 4$. But unless you successfully get yourself to stop believing this, it is still something that you know. Hence, if there is a sense in which in this case it is irrational for you to believe that $2 + 2 = 4$, it is not in *this* sense that you know p only if it is rational for you to believe p . So the thesis that you know p only if it is rational for you to believe p must be qualified. You know p , I will say, only if it is *epistemically* rational for you to believe p . Epistemic rationality, I will say, is the kind of rationality that is entailed by knowledge—and the strongest such kind, if more than one kind of rationality is entailed by knowledge.

So what do we know about epistemic rationality? Well, one thing that we can say up front, is that threats like that of the villain do not affect what it is epistemically rational for you to believe. That, after all, was the point of introducing a distinctive *epistemic* kind of rationality, in the first place. But this only tells us about what does *not* affect epistemic rationality. So what *does* affect epistemic rationality? The most obvious answer is: the *evidence*. It is epistemically rational for someone to believe p , philosophers often say, just in case p is adequately supported by her *evidence*. To say this, given the characterization of epistemic rationality as the strongest kind of rationality that is entailed by knowledge, is to say (1) that you don't know p unless p is adequately supported by your evidence, (2) that being adequately supported by your evidence counts as making it *rational* to believe p , in some sense of 'rational', and (3) that there is no further or stronger rationality condition on knowledge. I will accept this characterization of epistemic rationality in what follows.

This characterization of epistemic rationality is, of course, exactly what makes it so hard to see what role practical factors like those arising in High Stakes cases

³ In conversation, Stewart Cohen has suggested to me that the important difference between epistemic and practical rationality is that epistemic rationality is 'categorical', whereas practical rationality is 'hypothetical'. The case of Dr. Maria illustrates that the issue of pragmatic encroachment cross-cuts the 'categorical'/'hypothetical' distinction.

could play, in defeating knowledge. Nevertheless, in the remainder of this paper I will be outlining a picture on which subjects like Hannah lack knowledge in cases like High Stakes because in such cases it is not epistemically rational for them to believe, and knowledge requires the epistemic rationality of belief. The characterization of epistemic rationality just above may make this look like a difficult needle to thread, but we will see that much turns on what it takes for a proposition to be *adequately supported* by the evidence.

Before we go on, let me head off one possible source of confusion. Some philosophers may hold that Hannah's lack of knowledge in High Stakes cases *can't* be explained by the fact that it is epistemically irrational for her to believe in that case, not because of general assumptions about the relationship between epistemic rationality and evidence or truth, but simply because they find it intuitively much more natural to describe High Stakes as a case in which it is rational for Hannah not to rely on her beliefs. Philosophers who take this view emphasize the importance of distinguishing between what it is rational to believe and what it is rational to do, and would not have us conflate those two questions. I agree—let's not conflate the question of what it is rational to believe and what it is rational to do, no matter how we are understanding rationality. Nothing I will say in this paper will be inconsistent with the thesis that *some* High Stakes-like cases are cases in which it *is* rational for Hannah to believe, but it is not rational for Hannah to rely on this belief. Indeed, in my view there clearly are cases like this. If there are such cases, I will say that they are cases in which Hannah knows, but she should not act on her knowledge.⁴ All that the pragmatic encroachment thesis as explained on the picture I will be offering requires, is that in *some* High Stakes cases, belief is not epistemically rational.

4 Ignorant high stakes

Having laid out my strategy in advance, of arguing that subjects lack knowledge in High Stakes cases because it is not epistemically rational for them to believe, I must confront the obvious problem that the intuitive force of High Stakes cases does not seem to depend on whether the subject is aware of the stakes or not. For example, similar problems seem to be raised by cases like the following:

Ignorant High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since their mortgage payment is due on Sunday, they have very little in their account, and they are on the brink of foreclosure, it is very important that they deposit their paychecks by Saturday, but neither Hannah nor Sarah is aware of this. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah remembers the bank being open on Saturday morning a few weeks ago, so she says, 'Fortunately, it will be open tomorrow, so we can just come back.' In fact, Hannah is right—the bank will be open on Saturday.

⁴ Note that both Fantl and McGrath (2010) and Hawthorne and Stanley (2008) appear to deny this possibility.

The problem for my proposed strategy raised by Ignorant High Stakes arises from the assumption that whether it is *rational* for Hannah to believe that the bank will be open tomorrow should not depend on features of her situation of which she is unaware—rather, it should supervene on her mental states. But if Ignorant High Stakes has just as much a claim to be a case in which Hannah lacks knowledge as the original High Stakes case was, then it would seem that in Ignorant High Stakes, the reason why Hannah lacks knowledge that the bank will be open on Saturday can't be that it is not epistemically rational for her to believe.

This is true. In Ignorant High Stakes, it *is* epistemically rational for Hannah to believe that the bank will be open on Saturday, but she lacks knowledge nonetheless. But rather than taking this to be a problem for my approach, I take it to be a virtue. For Ignorant High Stakes relates to Apparent High Stakes (below) in exactly the same way that a whole range of Gettier cases relate to defeating conditions on epistemic rationality:

Apparent High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. They believe, having followed their statements closely, that their mortgage payment is due on Sunday, that they have very little in their account, and that they are on the brink of foreclosure, and so believe that it is very important that they deposit their paychecks by Saturday. However, unbeknownst to them, an unexpected reimbursement check from the IRS which suffices to cover their mortgage payment has just today been directly deposited to their account, and so it is not actually important for them to deposit their checks before Monday. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah remembers the bank being open on Saturday morning a few weeks ago, so she says, 'Fortunately, it will be open tomorrow, so we can just come back.' In fact, Hannah is right—the bank will be open on Saturday.

In Apparent High Stakes, the stakes are not actually high—they just seem to be. In Ignorant High Stakes, the stakes are really high, even though they don't seem to be. Advocates of pragmatic encroachment say that both are defeating conditions on knowledge—but epistemic rationality is at best defeated in Apparent High Stakes. What I'll now seek to remind you of, is that the relationship between Ignorant High Stakes and Apparent High Stakes holds for every other kind of defeating condition of knowledge.

First take the case of *counterevidence*. Counterevidence can defeat knowledge either in *apparent* cases or in *ignorant* cases:

Apparent Counterevidence. Hannah and Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks, but as they drive past the bank, Sarah observes that the lines inside are very long. Hannah remembers the bank being open on a Saturday morning a few weeks ago, so she says, keeping her eyes on the road, 'Fortunately, it will be open tomorrow'. Exercising her mischievous streak, Sarah says, 'but there's a sign out front that says CLOSED TOMORROW!' Hannah believes Sarah, but in fact, there is no sign out front, and the bank will be open on Saturday.

Ignorant Counterevidence. Hannah and Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks, but as they drive past the bank, Hannah observes that the lines inside are very long. Hannah remembers the bank being open on a Saturday morning a few weeks ago, so she says, watching all of the people standing in line, ‘Fortunately, it will be open tomorrow’. While Hannah is watching the people, she misses seeing the enormous yellow sign in front that says CLOSED TOMORROW. But in fact, the sign was put out by mistake, and the bank will actually be open on Saturday, as Hannah believes.

In neither Apparent Counterevidence nor Ignorant Counterevidence does Hannah know that the bank will be open on Saturday—and they bear the same relationship as Apparent High Stakes and Ignorant High Stakes do. In Apparent Counterevidence, not only does Hannah not know that the bank will be open on Saturday, it is epistemically irrational for her to believe that it will be; whereas Ignorant Counterevidence is its Gettier Counterpart—in Ignorant Counterevidence it is rational for Hannah to believe that the bank will be open on Saturday, but she lacks knowledge, due to the existence of counterevidence of which she is unaware. The same goes for other kinds of defeat. For example:

Apparent Undercutting. Hannah and Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks, but as they drive past the bank, Sarah observes that the lines inside are very long. Hannah remembers reading in the newspaper this morning that the bank has begun keeping Saturday hours, so she says, keeping her eyes on the road, ‘Fortunately, it will be open tomorrow’. Exercising her mischievous streak, Sarah says, ‘Did you read that in the Press-Gazette? You can’t count on them, because they’ve started making up stories in order to deal with budget cuts.’ Hannah believes Sarah, but in fact, the Press-Gazette is reliable, and the bank will be open on Saturday.

Ignorant Undercutting. Hannah and Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks, but as they drive past the bank, Hannah observes that the lines inside are very long. Hannah remembers reading in the newspaper this morning that the bank has begun keeping Saturday hours, so she says, ‘Fortunately, it will be open tomorrow’. Unbeknownst to Hannah, however, due to budget cuts, the Press-Gazette has started manufacturing local news stories, and consequently isn’t a very reliable source of news. Nevertheless, Hannah is right, and the bank will be open on Saturday.

Again, knowledge can be defeated either because of apparent undercutting defeaters for the agent’s evidence, or because of undercutting defeaters of which the agent is unaware. In the former case it would not even be epistemically rational for Hannah to believe that the bank will be open tomorrow; in the latter case, its Gettier counterpart, she lacks knowledge even though it would be rational for her to believe.

I infer from this set of cases—which can be extended to other kinds of defeaters for knowledge—that whatever the relationship is between knowledge and epistemic

rationality, it is one that sustains this general relationship: whenever belief in q is the kind of thing to make it epistemically irrational to believe p , a subject may fail to know p either because she believes q , or because q is true. In the 1970's a number of philosophers tried to capture this insight by proposing conditional-fallacy-prone analyses of knowledge, according to which to know is just to have a belief that is not only rational, but *would remain* rational given the addition of more true beliefs.⁵ Though there are significant obstacles to such theories, all I need for my explanatory strategy is the underlying insight that defeaters for knowledge come paired in this way. I will be offering a picture to explain how apparent high stakes can make it epistemically irrational to believe, and I will delegate the explanation of Ignorant High Stakes cases to this general relationship between epistemic rationality and knowledge.⁶

Now that we've finished with preliminaries, let's move on to develop a positive picture of how it could be that practical factors could affect epistemic rationality, and in the right way to explain the difference between High Stakes and Low Stakes.

5 So how could stakes affect epistemic rationality?

So this leave us with our puzzle: how could it be that epistemic rationality—the strongest kind of rationality condition on knowledge, which I've allowed above can be characterized by saying that it is epistemically rational for someone to believe p just in case her evidence adequately supports p —itself suffers from pragmatic encroachment? The answer lies in two small but important words: *adequately supports*.

What does it take for evidence to adequately support a conclusion? It will be helpful here to compare epistemic rationality to the ordinary practical rationality of action, where it is very natural to connect the rational course of action with the one that is best supported by the reasons available to the agent at the time⁷:

Sufficiency It is rational for S to do A just in case S has at least as much reason to do A as to not do A.

According to Harman (2002), one of the things that makes epistemic rationality so different from practical rationality is that though Sufficiency is correct for practical rationality, it is not correct for epistemic rationality. Is this so?

Harman starts with the correct observation that sometimes, even though S has at least as much evidence for p as for $\sim p$, it is not rational for S to believe p —because given that the evidence is fairly balanced, or that there is not very much of it, or that further, conclusive evidence is expected to soon be forthcoming, the rational thing to do is to *withhold* belief—to not make up one's mind about the matter in question.

⁵ See, in particular, the extensive discussion of such theories in Shope (1983).

⁶ See Schroeder (2010b) for further discussion.

⁷ Insofar as rationality can be connected to reasons, it is connected to *subjective* reasons, sometimes referred to as reasons that agents *have*, as opposed to reasons that there are, but which no one has. So I intend, here. See Schroeder (2008) for discussion.

This observation is correct—and important. But since it is an observation about evidence, and Sufficiency is a principle about reasons, in order to connect Harman's observation with the principle of Sufficiency, we need an auxiliary principle connecting evidence to reasons. The closest thing to what we need is the principle of *Special Evidentialism*:

Special Evidentialism S has at least as much epistemic reason to believe p as to believe $\sim p$ just in case S has at least as much evidence for p as for $\sim p$.

Note that in my formulation of Special Evidentialism I appeal to *epistemic* reasons. This is to finesse the problem that just as there may be some sense of 'rational' in which it is irrational for you to believe that $2 + 2 = 4$ when your family member's lives are credibly threatened by the villain, there may be some sense of 'reason' in which the fact that they have been so threatened is a 'reason' for you not to believe that $2 + 2 = 4$. Fortunately, since we have already introduced the term 'epistemic rationality', we can appeal to it here; I will stipulatively use 'epistemic reasons' for those reasons, whatever they are, which bear on epistemic rationality. Since we know that the villain's threat does not bear on the epistemic rationality of your believing that $2 + 2 = 4$, we may therefore safely assume that it is not an epistemic reason.

Special Evidentialism has the virtues of being almost exactly what Harman needs for his argument, of being highly plausible in its own right (especially once we have granted the characterization of epistemic rationality in terms of evidence), and of being a trivial corollary of *General Evidentialism*, which is plausible in its own right:

General Evidentialism S has at least as much epistemic reason to believe p as to believe q just in case S has at least as much evidence for p as for q .

We don't actually need General Evidentialism to see the allure of Harman's argument, but it will be useful to refer back to, later. Now, if epistemic rationality obeys the principle of Sufficiency, then (substituting):

Ep.Sufficiency It is epistemically rational for S to believe p just in case S has at least as much epistemic reason to believe p as to not believe p .

And now it is straightforward to see that we can use this together with Special Evidentialism to derive the negation of Harman's observation, if only we could appeal as well to the following assumption:

Problematic S has at least as much epistemic reason to believe p as to not believe p just in case S has at least as much epistemic reason to believe p as to believe $\sim p$.

If Ep.Sufficiency, Problematic, and Special Evidentialism are all true, then Harman's observation is false. So since Harman's observation is true, one of Ep.Sufficiency, Problematic, and Special Evidentialism has got to go. It should be clear which assumption is the one I think is Problematic.

The problem with Problematic is that there is more than one way to not believe p . You can not believe p by believing $\sim p$ —or, you can not believe p by withholding belief with respect to p . The Problematic assumption simply assumes that there cannot be any epistemic reasons to withhold belief. For if there were such reasons, then it could be that even though S has at least as much reason to believe p as to believe $\sim p$, S has even more reason to withhold with respect to p —and hence more reason to not believe p than to believe p . In fact, this is an incredibly natural diagnosis of the cases which satisfy Harman’s observation—for by Harman’s own lights, they are cases in which the only rational option is to withhold belief. Ep.Sufficiency can describe these same cases by saying that in these cases, S has more reason to withhold than to have either belief.

6 The importance of reasons for withholding

In light of the considerations in the last section, I am unpersuaded by Harman’s attempt to argue that epistemic rationality does not obey the principle of Sufficiency. We can allow that it does obey this principle, so long as we recognize the existence of epistemic reasons to withhold. Actually, it will turn out to be more convenient in what follows to work with the following principle of *Generalized Sufficiency*:

Gen.Sufficiency It is rational for S to do A just in case S has at least as much reason to do A as in favor of any of the alternatives to doing A .⁸

On the assumption that the alternatives to believing p are believing $\sim p$ and withholding with respect to p , this yields what we are looking for:

Bf.Sufficiency It is epistemically rational for S to believe p just in case S has at least as much epistemic reason to believe p as to believe $\sim p$ and S has at least as much epistemic reason to believe p as to withhold with respect to p .

This principle tells us what it takes for evidence to be *adequate*. To be adequate, S ’s evidence for p must not only be better than S ’s evidence for $\sim p$; it must also amount to better reason to believe p than all of S ’s epistemic reasons to withhold with respect to p .

So far, I am hypothesizing that the reason *why* it can be epistemically irrational for S to believe p even though S has at least as much—or even more—evidence for p as for $\sim p$, is that there is an important class of epistemic reasons to withhold—that is, reasons to withhold that are of the right kind to play a role in affecting epistemic rationality. Importantly, though to many this will sound like a surprising

⁸ The principle of *Generalized Sufficiency* can be thought of as a natural generalization of *Sufficiency* from a two-option case to a multiple-option case, and so it is convenient to use in comparing the three options of believing p , believing $\sim p$, and withholding with respect to p . In fact, however, I am now inclined to think that there are reasons to believe that there may be substantive tradeoffs in deciding between *Sufficiency* and *Generalized Sufficiency*. But I don’t believe that any of the points I will be making in this paper turn on these tradeoffs.

claim, it is a hypothesis that everyone can accept—friends and foes of pragmatic encroachment alike—for this hypothesis makes no claim about what the reasons to withhold are like.

Indeed, it is a hypothesis that quite arguably—although I won't defend this claim here in full—everyone *should* accept. For epistemologists have repeatedly found it notoriously difficult, over a very long period of time, to say what it is for evidence to be sufficient to make belief epistemically rational. In some cases—for example, with regard to beliefs about the distant past or about difficult theoretical questions about which there is little evidence to be had either way—belief seems to be epistemically rational even though there is relatively little evidence. Yet in other cases belief seems to be irrational even though a significant preponderance of evidence makes the conclusion extremely likely. For example, many philosophers believe that it is epistemically irrational to believe that a single lottery ticket will lose, no matter the size of the lottery, on the basis simply of the odds. But in arbitrarily large lotteries, this conclusion may be arbitrarily well supported by the evidence—a hard standard for most beliefs about science, history, philosophy, or pop culture trivia to have to pass.

One piece of circumstantial evidence for just how difficult it is to give an adequate account of what it is for evidence to be sufficient for epistemic rationality, is the number of accounts in the literature which have resorted to fallacious conditional formulations—saying for example that some evidence is sufficient for rationality just in case *were* the agent to believe on the basis of that evidence, she *would* be rational.⁹ It is tempting to resort to conditional formulations precisely because it is so difficult to say in categorical terms what makes evidence sufficient. Nevertheless, the principle of Generalized Sufficiency does allow us to say so: S's epistemic reasons to believe *p* are sufficient just in case they outweigh the epistemic reasons for the alternatives—i.e., for withholding and for believing $\sim p$. This account isn't a panacea, because it leaves us with the problem of diagnosing just what the epistemic reasons to withhold are.¹⁰ But as we will see in what follows,

⁹ Compare Shope's (1983) survey of attempts to appeal to the sufficiency or conclusiveness of reasons in attempts to analyze knowledge, which is one of the main places I know of where the concept of sufficiency has been put to work in epistemology. *Conclusive reasons* accounts of knowledge can be thought of as motivated by the correlation between 'ignorant' and 'apparent' cases in which knowledge is defeated, and propose (roughly) that for a true belief to be knowledge, the reasons for which it is held have to be sufficiently good to make it rational, even were the subject to learn more true things. A number of such accounts were offered in the literature, and they typically shared the feature that the appeal to *sufficiency* is effectively discharged by being understood in terms of whether it *would* still be rational (or 'justified') for the subject to believe, were she to be better-informed. In contrast, if we can appeal to a categorical concept of sufficiency, then it is possible to defend a view in this family that does not suffer from conditional-fallacy problems. See Schroeder (2010b) for elaboration.

¹⁰ In this section and to follow I will be referring to 'reasons to withhold' that are not evidence. If the principle of *Generalized Sufficiency* is correct, then this is the right way of talking. If, however, the principle of *Sufficiency* is correct, it would be more accurate to speak of reasons not to believe *p* that are not evidence that $\sim p$. Intuitively, however, the explanation of why there have to be epistemic reasons not to believe *p* that are not evidence that $\sim p$ has something important to do with the possibility of withholding. For this reason, and because I think the issues discussed here are orthogonal to which of *Sufficiency* or *Generalized Sufficiency* is correct (or perhaps they are compatible), I will stick to talk about reasons to withhold.

it *does* productively constrain this investigation. Consequently, I believe that Generalized Sufficiency and in particular, the principle of Belief Sufficiency, can and should be accepted by all parties.

7 Reasons to withhold can't be evidence

Nevertheless, though Belief Sufficiency can be accepted by all parties, once we allow that there can be epistemic reasons for withholding, it becomes much less obvious that the reasons to withhold cannot be practical in nature. This is because as long as we cling to the idea that only evidence can be an epistemic reason, it is easy to dismiss pragmatic factors out of hand as being the wrong kind of thing entirely to bear on epistemic rationality. But reasons to withhold *can't* be evidence. Consequently, once we allow for epistemic reasons to withhold, we must allow that epistemic reasons are *not* exhausted by the evidence. And this is how reasons to withhold that are practical in nature can get their foot in the door.

Why is it that reasons to withhold cannot be evidence? It is because the evidence is exhausted by evidence which supports p and evidence which supports $\sim p$. But the evidence which supports p is reason to believe p , and the evidence which supports $\sim p$ is reason to believe $\sim p$. Consequently the reasons to withhold must come from somewhere else. So they cannot be evidence.

You might think that this reasoning goes too quickly.¹¹ You might think, for example, that though it is true that any *particular* piece of evidence must either support p or support $\sim p$, and hence be a (defeasible) reason to believe p or a (defeasible) reason to believe $\sim p$, the *total evidence* can support withholding belief (all-things-considered). But I think this thought is making a mistake. The all-things-considered way in which the totality of reasons supports belief in p , belief in $\sim p$, or withholding is importantly different from the defeasible way in which particular reasons provide their support. The all-things-considered support provided by the totality of reasons isn't just one of the things which goes into the balance in determining which of belief in p , belief in $\sim p$, or withholding is rational; it is the result of balancing the considerations which bear on which of these states is rational. So to say that the totality of the evidence can weigh decisively in favor of withholding is just to say that the question of which state it is most rational to be in is determined solely by the evidence. Naturally, this is the orthodox view, but it does not tell us anything about *how* the evidence determines which state is rational—that is, the puzzle about what makes evidence sufficient.

On the picture I've been suggesting, however, the question of which state it is rational to be in is a matter of the relative weight of the epistemic reasons in favor of believing p , believing $\sim p$, and withholding with respect to p . *Given that picture*, the individual considerations which are epistemic reasons to withhold don't look like they can be evidence. But if not, then what sort of thing could they be?

In part 3 I will develop what I think is one natural picture about the nature of reasons to withhold, according to which at least some of them *are* due to practical

¹¹ Thanks to Stewart Cohen for discussion of this point.

factors—at least, practical factors of a certain kind. What I will show, is that this picture results from intuitive claims about the nature of reasons to withhold, that it entails the pragmatic encroachment thesis, and that it allows for more fine-grained predictions than other formulations of the thesis of pragmatic encroachment. What I won't argue, is that this is the *right* picture of the nature of epistemic reasons to withhold; my aim is simply to show how natural and conservative this view would be, if it were right. The reason why I don't defend a stronger conclusion in this paper is that fully evaluating which sorts of things are and are not epistemic reasons is a large further project that would take us too far afield. But I will begin to sketch the contours of this larger project in part 4—in order to show what work would need to be done in order to defend the picture outlined in part 3. But first, the picture.

8 A first pass at reasons to withhold

If you want to know what sort of thing could be a reason for or against withholding, it pays to pay attention to what makes withholding different from belief. To withhold is to *not make up your mind*, to have formed no belief. Consequently, any disadvantage of forming beliefs—of making up your mind—is potentially a reason to withhold.

So 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*. But what, exactly, *are* the costs of type-1 and type-2 error? In various circumstances different sorts of things could turn out to be downsides of having a false belief, or downsides of not having formed a belief, but I will focus on general costs—costs that we can expect to accrue to false beliefs or to the lack of belief no matter the situation, just because of the kind of state that belief is.

Conceived in this way, the 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. For example, if you want to find the Lady and avoid the Tiger, then if you believe that the Lady is behind the left-hand door and the Tiger is behind the right-hand door, you will open the left-hand door. In this case, the cost of type-1 error is that you will get the Tiger instead. 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. But they're a *special* sort of practical factor—threats from villains about what will happen to your family if you don't withhold don't count. To be the right sort of cost of type-1 error, on this view, you have to be a cost that a belief gives rise to when it is false, due to its playing its normal role as a belief – the sort of cost that is intrinsic to the nature of belief. And to be the right sort of cost of type-2 error, on this view, you have to be

a benefit of having made up your mind that having made up your mind gives rise to, due to its playing its normal role—the sort of benefit that is intrinsic to the nature of belief.

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. Indeed, the main difference between the Low Stakes case and the High Stakes case is that the cost of falsely believing that the bank will be open on Saturday that will accrue to that false belief due to Hannah acting on it, is much higher in High Stakes. So in addition to being an intuitive thing to say about reasons to withhold, the hypothesis that reasons to withhold derive (at least in part) from the relative costs of type-1 and type-2 error, as elaborated above, looks like the right kind of thing to predict the right results in the contrast between High Stakes and Low Stakes.

In fact, understanding reasons to withhold in terms of the balance of the costs of type-1 and type-2 error is already sufficient to help us make more refined predictions than merely saying that ‘stakes can undermine knowledge’. For as we’ve just seen, what the stakes do directly, in the original bank cases, is to affect the cost of type-1 error. But it is also possible to construct cases in which raising the stakes raises the cost of type-2 error along with the cost of type-1 error. Our picture predicts that these cases will not affect knowledge in as clear-cut a way¹²:

Forced Choice, Low Stakes. Hannah and her wife Sarah are out driving on Saturday morning, at twenty minutes to noon. Sarah remembers that they still haven’t deposited their paychecks from Friday, but points out that just one of their bank’s two branches is open until noon on Saturdays, but she can’t remember which, and there is only time to try one. Hannah says, ‘Oh, I remember being at the branch on Chapala Street two weeks ago on Saturday. It’s the one that is open today.’ Hannah is right—the branch on Chapala Street is the one that is open on Saturday.

Forced Choice, High Stakes: Hannah and her wife Sarah are out driving on Saturday morning, at twenty minutes to noon. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks that day, but they have so far forgotten to do so. Sarah remembers that they still haven’t deposited their paychecks from Friday, but points out that just one of their bank’s two branches is open until noon on Saturdays, but she can’t remember which, and there is only time to try one. Hannah says, ‘Oh, I remember being at the branch on Chapala Street two weeks ago on Saturday. It’s the one that is open today.’ Hannah is right—the branch on Chapala Street is the one that is open on Saturday.

So as we can see, the hypothesis that knowledge is affected by reasons to withhold, and reasons to withhold derive from the preponderance of the costs of type-1 error

¹² Shaffer (2006) offers cases which exploit the same features diagnosed here as an objection to Hawthorne (2004) and Stanley (2005), although he acknowledges in a note that they may not raise a problem for the view of Fantl and McGrath (2002), which is more closely related to the one offered here.

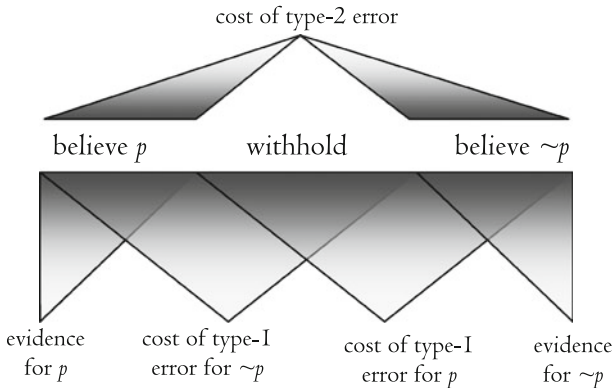
over the costs of type-2 error, allows for relatively fine-grained predictions, which make intuitive sense.¹³

9 In real life things are more complicated

In real life, however, things are more complicated. The main complication arises from the fact that there is, in fact, no such thing as the cost of type-1 error regarding some proposition. There is only the cost of type-1 error in believing it, and the cost of type-1 error in believing its negation. These two costs can easily come apart. For example, in High Stakes, the cost of wrongly believing that the bank is *not* open on Saturday is that Hannah and Sarah will have to stand in line, whereas the cost of wrongly believing that it *is* open on Saturday is that Hannah and Sarah will miss their impending bill. The latter cost can easily outweigh the former—for example, it could lead the bank to foreclose on their house, if we set the case up correctly.

Since there is strictly speaking no such thing as the cost of type-1 error with respect to a proposition *tout court*, strictly speaking the costs of type-1 error are not associated with reasons to withhold. Rather, the costs of type-1 error of believing p are associated with a reason to *not believe p* —i.e., to either withhold or believe $\sim p$.¹⁴ And similarly, the costs of type-1 error of believing $\sim p$ are associated with a reason to not believe $\sim p$ —i.e., to either withhold or believe p . Meanwhile, the costs of type-2 error are associated with reasons to *not withhold*—i.e., to either believe p or believe $\sim p$.

These observations are what motivate the model illustrated by the following picture, in which shaded triangles represent which of believing p , withholding, and believing $\sim p$ is supported by each kind of reason:



¹³ I've been writing as if *both* the costs of type-1 error and the costs of type-2 error can be epistemic reasons for and against withholding, respectively. But nothing I've said rules out the possibility that some costs of type-1 error are epistemic reasons to withhold, but no costs of type-2 error are ever reasons against withholding. Whether the principle of *General Evidentialism* is correct, and not just the principle of *Special Evidentialism*, turns on whether there are epistemic reasons *against* withholding, and hence on whether the costs of type-2 error are really epistemic reasons not to withhold.

¹⁴ On the picture provided by *Generalized Sufficiency*, reasons against an alternative need to be interpreted as reasons for the other alternatives.

This picture incorporates our assumptions that evidence for p is reason to believe p , evidence for $\sim p$ is reason to believe $\sim p$, the cost of type-2 error for p is reason to either believe p or believe $\sim p$, the cost of type-1 error for p is reason to not believe p , and the cost of type-1 error for $\sim p$ is reason to not believe $\sim p$. There are no direct reasons to withhold in this picture; only the net interaction effect of reasons to not believe p and reasons to not believe $\sim p$.

Keeping in mind that this is just an illustrative model, and that doing arithmetic with the terms involved is at very best an extreme idealization on how reasons ‘add up’ and weigh against one another,¹⁵ it will be useful in what follows to illustrate some of the structural features of this model by following through on this idealization. So I’ll introduce abbreviations as follows: believing p is supported by the evidence for p (Ev_p), the cost of type-2 error ($Err2$), and the cost of type-1 error in believing $\sim p$ ($Err1_{\sim p}$). Similarly, believing $\sim p$ is supported by the evidence for $\sim p$ ($Ev_{\sim p}$), the cost of type-2 error ($Err2$), and the cost of type-1 error in believing p ($Err1_p$). Finally, withholding is supported by the costs of each type of error ($Err1_p$ and $Err1_{\sim p}$).

The principle of Belief Sufficiency tells us that it is epistemically rational for S to believe p just in case S has at least as much epistemic reason to believe p as *both* S ’s epistemic reason to withhold *and* S ’s epistemic reason to believe $\sim p$. This model lets us say more about what those conditions are. In particular, S ’s epistemic reasons to believe p are at least as good as S ’s epistemic reasons to believe $\sim p$ just in case

$$Ev_p + Err2 + Err1_{\sim p} \geq Ev_{\sim p} + Err2 + Err1_p \quad \text{or, cancelling,} \\ Ev_p + Err1_{\sim p} \geq Ev_{\sim p} + Err1_p.$$

And similarly, S ’s epistemic reasons to believe p are at least as good as S ’s epistemic reasons to withhold just in case

$$Ev_p + Err2 + Err1_{\sim p} \geq Err1_p + Err1_{\sim p} \quad \text{or, cancelling,} \quad Ev_p + Err2 \geq Err1_p.$$

This means that there are correspondingly, at least in principle, *two* ways in which having better evidence that p than that $\sim p$ can still fail to make it rational to believe p . This can happen if the costs of type-1 error of believing p exceed the costs of type-2 error by a sufficient amount—enough to outweigh the evidence that p and make it more rational to withhold than to believe p ¹⁶:

$$\text{Withholding is more rational : } Err1_p - Err2 > Ev_p$$

Or it can happen if the costs of type-1 error of believing p exceed the costs of type-1 error of believing $\sim p$ by more than the evidence for p outweighs the evidence for $\sim p$.

$$\text{Believing } \sim p \text{ is more rational : } Err1_p - Err1_{\sim p} > Ev_p - Ev_{\sim p}$$

¹⁵ Compare chapter 7 of Schroeder (2007).

¹⁶ Note particularly that on the plausible assumption that it is always true that $Err2 < Err1_p$, for all p , it follows that it is never rational to believe p without evidence—even though strictly speaking there may be epistemic reasons to believe p that aren’t evidence (namely, $Err1_{\sim p}$). This is an attractive result, because it confirms the characterization from Sect. 1.2 that believing p is epistemically rational just in case the *evidence* for p is adequate.

Note that our original High Stakes and Low Stakes cases, borrowed from Stanley (2005), trigger *both* conditions, by raising the cost of type-1 error in believing that the bank will be open tomorrow, without raising the cost of type-1 error in believing that the bank will not be open tomorrow, and without raising the cost of type-2 error. The forced choice cases trigger neither condition, because since they are symmetric, they raise the cost of type-1 error about both propositions at the same time, and the nature of the forced choice also raises the cost of type-2 error to match (as closely as possible) the cost of type-1 error. So the more detailed model makes good on the intuitive claims I made in the last two sections.

10 Further predictions

If this model is roughly correct, however, then we should expect there to be further cases which trigger one condition but not the other. It is easy to construct a case to trigger the first condition without triggering the second, making withholding more rational than believing p without making believing $\sim p$ more rational:

Nasa Engineering. Hannah and Sarah are engineers working on the design of NASA's next-generation shuttle, a multi-billion dollar project planned to operate over several decades and ultimately carry hundreds of astronauts into space, where error means death. Currently they are trying to decide which materials to use for an important component, and are investigating two new alloys, to see which will be more appropriate for the component. Citing preliminary research, Sarah notes that the first alloy holds up better under temperatures under 300°, and that most alloys which hold up well under 300° also perform well at shuttle temperatures. Hannah says, 'okay, so the first alloy will hold up better at shuttle temperatures.' In fact Hannah is correct; the first alloy does hold up better at higher temperatures.

Nasa Engineering is a case in which the costs of each direction of type-1 error are equally high, and both are much higher than the costs of type-2 error, since there is no forced decision—the engineers can simply wait on more research before deciding which alloy will hold up better. It is trickier to construct a case in which believing $\sim p$ is more rational than believing p but both are more rational than withholding. Indeed, there are general reasons to suspect that this sort of case is impossible, and to hope that a more accurate picture of what is going on would improve on the model I've been sketching here by building in constraints that rule this possibility out.

Nevertheless, here is my best go at what a case would have to be like, in order to have this feature:

Game Show. Hannah and Sarah are playing Go Big or Go Home, a successful game show on daytime television with a B-celebrity host. They have reached the final question, which is: 'will the bank be open tomorrow, on Saturday?'. The possible answers are 'yes' and 'no', and they must answer within the time limit, or they will lose all of their money (they have accumulated a very large

sum so far). If they answer and get it right, they double their money, but if they answer ‘yes’ and get it wrong, they lose all of their money and if they answer ‘no’ and get it wrong, they keep what they already have. Hannah tells Sarah, ‘The answer is ‘yes’—I was there three weeks ago on a Saturday morning, and it was open.’ In fact, Hannah is correct; the bank will be open tomorrow.

Game Show is a situation of forced choice, which raises the cost of type-2 error, but it is constructed to keep the relative costs of the two directions of type-1 error different. It is clearly a more controversial case, but if it works, it makes it more rational to believe $\sim p$ than to believe p , but not more rational to withhold, than to believe p .

The picture painted in this and the preceding sections is just a picture, and the most fine grained predictions that I drew out from it were based on highly idealizing assumptions. Nowhere have I defended in full the assumptions of the picture, to the effect that the costs of type-1 and type-2 error that I have been discussing are really epistemic reasons. But I hope that my picture shows three things: (1) that there is a highly natural way of understanding reasons to withhold in terms on which it is both easy to see their practical significance and nevertheless easy to see why Pascalian considerations will not count. (2) That this picture of how reasons to withhold work fits smoothly with the cases that motivate pragmatic encroachment, and in fact promise to make more fine-grained intuitive predictions than some other views. And (3) that all of this is consistent with the thesis that it is epistemically rational for S to believe p just in case S has adequate evidence for p . Reasons to withhold simply raise the bar on how good the evidence needs to be, in order to be adequate.

11 Are reasons to withhold really ‘epistemic’?

The main proposal of this paper can be put by saying that evidence is only one kind of epistemic reason.¹⁷ Of course, if we had simply defined ‘epistemic reason’ to mean ‘evidence’, then the idea thus expressed would be incoherent—but I didn’t define ‘epistemic reason’ that way. Rather, I allowed that any reason may count as ‘epistemic’, if it contributes to *epistemic rationality*, and I introduced epistemic rationality as the strongest kind of rationality that is entailed by knowledge. This way of framing things successfully ruled out Pascalian considerations from counting as epistemic reasons, without *ipso facto* deciding whether or not evidence is the only kind of epistemic reason.

I then offered some evidence that there really are epistemic reasons to withhold belief, in addition to epistemic reasons to believe—in particular, such reasons can explain why having more evidence for p than for $\sim p$ does not suffice to make it rational to believe p . Moreover, differences in the reasons to withhold can explain why in some cases a greater amount or preponderance of evidence is required in order to make belief rational, whereas in other cases only less or a lesser preponderance of evidence is required. These arguments didn’t tell us what the

¹⁷ Compare Owens (2000).

epistemic reasons to withhold *are*; they only suggested that there *must* be *some* epistemic reasons to withhold. But one thing *is* clear: whatever the reasons to withhold with respect to *p* are, if the picture I've offered about the sufficiency of reasons is right, then reasons to withhold must be something other than evidence for or against *p*.

So how can we find a principled way of triangulating on what the epistemic reasons to withhold really are, and in turn, of evaluating whether the kinds of cost of type-1 and type-2 errors that I've discussed in part 3 should really be counted as *epistemic* reasons, or lumped together with Pascalian considerations? Must it come down to intuition-mongering with respect to the kinds of high-stakes cases we've been discussing all along?

I think the answer is 'no'. For it turns out both that there are plausible *independent* markers of the distinction between epistemic and non-epistemic reasons for belief, and that the distinction between epistemic and non-epistemic reasons for belief is matched by parallel distinctions among reasons for intention and other kinds of attitude. This means, first, that the independent markers of the distinction between epistemic and non-epistemic reasons for belief give us *clues* as to which reasons to withhold should count as epistemic, and second, that the parallel distinctions for other attitudes give us an independent place to look for leverage, in order to sort out exactly what *does* distinguish epistemic from non-epistemic reasons.

The chief plausible marker of the distinction between epistemic and non-epistemic reasons, and the one that will be most helpful for our purposes, here, is that, as a wide range of philosophers have noticed, it is much less straightforward to believe *for* or *on the basis of* non-epistemic reasons, such as the considerations provided by Pascal.¹⁸ Being convinced by Pascal's Wager is the kind of thing to motivate you to spend less time with your atheist friends and more time taking the Eucharist, but not the kind of thing to make it any easier to believe in God. Spending less time with your atheist friends and more time taking the Eucharist are what we might call *indirect strategies* to induce belief in God; paradigmatic non-epistemic reasons require such indirect strategies, but believing for paradigmatic epistemic reasons (evidence) doesn't require any indirect strategies.

Note, therefore, that just as there is a contrast between reasons to believe on the basis of which it is easy to believe without resorting to indirect strategies, and reasons to believe on the basis of which it is not easy to believe without resorting to indirect strategies, there is a similar contrast among reasons to withhold. For example, if you are waiting on the results of a biopsy to confirm whether you have cancer, the fact that the results will soon be announced is a reason to withhold belief—to wait to make up your mind—even if you already have a fair bit of evidence one way or the other. Note that no indirect strategies are required, in order to get yourself to withhold belief, on the basis of the fact that a conclusive test is about to be announced. In contrast, if the villain offers to kill your family unless you withhold as to whether $2 + 2 = 4$, you will need to resort to indirect strategies—just as if the threat had been about *belief*.

¹⁸ Compare Williams (1973), Bennett (1990), and Hieronymi (2006).

Observations like this one support the thesis that the fact that further, potentially conclusive, evidence is shortly forthcoming is the right kind of thing to be an epistemic reason to withhold—even though it is not itself evidence, one way or the other. And indeed, this is a highly plausible claim in its own right—it *does* seem epistemically irrational to make up your mind in advance of getting evidence that you know will be conclusive and could tell the other way. Moreover, it does seem like the kind of thing that can undermine knowledge. If these claims are right, then the marker based on the need for indirect strategies converges on independent plausible judgments about which reasons affect epistemic rationality—for reasons to *withhold*, as well as for reasons to believe. So if that's right, then we can use it to test our hypothesis from part 3 about stakes-involving reasons to withhold, based on the costs of type-1 and type-2 error.

So: does withholding belief in high-stakes cases require indirect strategies? Or is it a perfectly normal thing, to hold out for more evidence in cases in which we are aware that the cost of being wrong are high? I venture that we do *not* require indirect strategies in such cases, and that it *is* a perfectly normal thing to hold out for further evidence, in such cases. I also venture that it is highly plausible that we don't require indirect strategies to make up our minds, in cases in which we see that the costs of indecision that derive from our inability to act are high. *If* I am right about these things, then that is evidence that the kinds of reasons I've appealed to in part 3 plausibly *are* epistemic.

12 The right-kind/wrong-kind distinction is a *general* problem

As mentioned above, a different place to look to, in order to gain independent leverage on the question of what makes reasons epistemic rather than non-epistemic, is the case of other kinds of attitudes, with respect to which there are parallel distinctions. For example, as many have taken Gregory Kavka's (1983) famous toxin puzzle to illustrate, offering someone a financial reward for having an intention, such that the award can be gained without ever having to do the thing intended, is very different from giving them an ordinary reason to intend—it's very hard to intend *for* the prospect of such a reward, for example, without taking indirect strategies such as placing independent side-bets. And although someone who manages to have the rewarded intention may be 'rational' in *some* sense, we don't ordinarily think of this as a distinctively rational state to be in, qua intention. Similar distinctions arise for all kinds of other attitudes—for example, the fact that Jones has offered you one thousand dollars if you will admire her daughter makes it *advantageous*, but not rational qua admiration, to admire Jones's daughter, and requires indirect strategies, whereas the fact that Jones's daughter is hard-working and generous does make it rational to admire her, and no indirect strategy is required to admire her on that basis. These parallel distinctions are often referred to as the distinction between the 'right' and 'wrong' kinds of reason for each kind of attitude.¹⁹

¹⁹ Compare D'Arms and Jacobson (2000a, b); Rabinowicz and Rønnow-Rasmussen (2004); Schroeder (2010a).

Given the parallel distinctions between ‘right’ and ‘wrong’ kinds of reason, for different kinds of attitude, it is highly plausible that any account of what the difference is, between epistemic and non-epistemic reasons for belief, should fall out from a more general account of the difference between the ‘right’ and ‘wrong’ kinds of reason—which should in turn be informed by the cases of these other attitudes, in addition to belief. Solving this problem is far too large a task for the closing section of a paper about pragmatic encroachment, but ultimately it is to a satisfactory solution of this *general* problem that we should expect our answer to whether the reasons to withhold discussed in part 3 are of the ‘right kind’ to be *epistemic*, to be beholden.

That is why I haven’t sought, in this paper, to defend the thesis that these really are epistemic reasons, outright. It is enough, for now, to have shown how it *could* be that practical factors could undermine knowledge—and to have shown what broader questions would need to be resolved, in order to make good on this.

Acknowledgments Special thanks to Jake Ross, Stew Cohen, Jason Stanley, Jeremy Fantl, Matthew McGrath, Josh Dever, Michael Bratman, Kritika Yegnashankaran, Eric Wiland, and John Brunero. Thanks also to audiences at Stanford University, the University of Texas at Austin, and to an audience for a related paper at the first annual Saint Louis Area Conference on Reasons and Rationality, as well as to an extraordinarily helpful blind referee for that paper for the *European Journal of Philosophy*. This paper was written under the support of a grant from the Advancing Scholarship in the Humanities and Social Sciences initiative at USC.

References

- Bennett, J. (1990). Why is belief involuntary? *Analysis*, 50, 93.
- D’Arms, J., & Jacobson, D. (2000a). Sentiment and value. *Ethics*, 110(4), 722–748.
- D’Arms, J., & Jacobson, D. (2000b). The moralistic fallacy. *Philosophy and Phenomenological Research*, 61(1), 65–90.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, 111(1), 67–94.
- Fantl, J., & McGrath, M. (2010). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Harman, G. (2002). Practical aspects of theoretical reasoning. In A. Mele & P. Rawling (Eds.), *The oxford handbook to rationality*. Oxford: Oxford University Press.
- Hawthorne, John. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105(10), 571–590.
- Hieronimi, P. (2006). Controlling attitudes. *Pacific Philosophical Quarterly*, 87(1), 45–74.
- Kavka, G. (1983). The toxin puzzle. *Analysis*, 43(1), 33–36.
- Owens, D. (2000). *Reason without freedom*. New York: Routledge.
- Rabinowicz, W., & Rønnow-Rasmussen, T. (2004). The strike of the demon: on fitting pro-attitudes and value. *Ethics*, 114, 391–423.
- Schroeder, M. (2007). *Slaves of the passions*. Oxford: Oxford University Press.
- Schroeder, M. (2008). Having reasons. *Philosophical Studies*, 139(1), 57–71.
- Schroeder, M. (2010a). Value and the right kind of reason. *Oxford Studies in Metaethics*, 5, 25–55.
- Schroeder, M. (2010b). Knowledge is belief for sufficient (objective and subjective) reason.
- Shaffer, J. (2006). The irrelevance of the subject: Against subject-sensitive invariantism. *Philosophical Studies*, 127, 87–107.
- Shope, R. (1983). *The analysis of knowing*. Princeton: Princeton University Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Williams, B. (1973). Deciding to believe. In: *Problems of the self* (p. 148). Cambridge: Cambridge University Press.