

Can Pragmatists Be Moderate?

Alex Worsnip

Forthcoming in *Philosophy & Phenomenological Research*

(Winner of the 2019 Young Epistemologist Prize)

Penultimate version; please cite the published version where possible

0. Introduction: hard vs. moderate pragmatism

In the debate between evidentialists and pragmatists about reasons for belief, the following sort of case is frequently discussed:

Bribe. You have strong (but not utterly infallible) evidence that the game starts at 3pm. But I, an eccentric millionaire, offer you a bribe: if you can avoid believing that the game starts at 3pm, I will give you \$1,000,000.¹

Meanwhile, in the debate about “pragmatic encroachment” in epistemology, the following sort of case is frequently discussed:

Risks. You have strong (but not utterly infallible) evidence that your interview starts at 3pm. But if you take it for granted that the interview starts at 3pm, and you’re wrong, you will miss the interview, and lose your only shot at your dream job. (Let’s suppose you could, instead of taking the start time of the interview for granted, call the interviewer to check on the time of the interview.)²

These two cases have important differences, but share two core features. First, they are both cases where you have strong (but not utterly infallible) evidence for some proposition p . Second, they are both cases where a *pragmatic fact* holds: namely, that believing or relying on p leads to, or risks leading to, a costly outcome. This pragmatic fact is evidentially irrelevant to whether p . Now, in both cases, we can ask the following question: is this pragmatic fact nevertheless capable of making a difference to whether you ought to believe p ?

For helpful discussion of the ideas in this paper, I’m grateful to Jeremy Fantl, Daniel Fogal, Jamie Fritz, Georgi Gardiner, Zoë Johnson King, Alex King, Stephanie Leary, Alexandra Lloyd, Cameron McCulloch, Matt McGrath, Lisa Miracchi, Sarah Moss, Ram Neta, Keshav Singh, Daniel Wodak, and audiences at the University of Michigan and the Rutgers Epistemology Conference.

¹ Many (e.g. Foley 1987: ch. 5; Reisner 2009; Rinard 2018; Berker 2018) discuss versions of this case where the bribe is *for* belief. For discussions of bribes *not* to believe specifically, see Fantl & McGrath (2002: 82); Fritz (forthcoming).

² Cases like this one are discussed by, e.g., Fantl & McGrath (2002); Stanley (2005).

Let us call those who answer this question in the affirmative even in cases like Bribe *hard pragmatists*.³ Conversely, let us call those who answer it in the affirmative in cases like Risks, but not in cases like Bribe, *moderate pragmatists*. Many philosophers, including most self-identified pragmatic encroachers, seem to be moderate pragmatists.⁴

Moderate pragmatists incur an explanatory burden. At the simplest level, the burden is to explain why the relevant pragmatic fact makes a difference in Bribe but not in Risks. For the explanation to be a good one, this must involve more than simply pointing to a difference between Bribe and Risks, of which there are several that are no doubt already occurring to you. The explanation must also be (i) theoretically compelling, and (ii) generalizable to a number of other cases (which I'll be mentioning shortly) that moderate pragmatists have typically wanted to include, or not to include, among the group of cases where pragmatic facts can make a difference to what you ought to believe.

Explicit attempts to discharge this explanatory burden are surprisingly rare – and, where they do appear, often brief. In this paper, I will reconstruct what I take to be the five most promising explanations of how pragmatism can be kept moderate – drawing on suggestions made and hinted at in existing literature – and argue that none of them ultimately succeed. Thus, I will argue, it is harder for pragmatists to keep their view moderate than many have assumed. Moderate pragmatism may be an unstable stopping point between an entirely anti-pragmatist view on one hand, and hard pragmatism on the other. Of course, it is possible that there is some other explanation of the differences between the cases that can vindicate moderate pragmatism. But until we have such an explanation, I believe, it is not rational for us to accept moderate pragmatism.

1. Pragmatic encroachment, and reasons vs. threshold pragmatism

As I mentioned, moderate pragmatism tends to be the favored view of advocates of the “pragmatic encroachment” thesis. According to this thesis, pragmatic considerations make a difference to what one ought to believe tout court *via* making a difference to what one *epistemically* ought to believe, or is justified in believing.⁵ As a matter of logic, there's nothing inconsistent about combining pragmatic encroachment with what I'm calling hard pragmatism. But pragmatic encroachers face strong pressure to be moderate pragmatists. For it seems wild to suggest that in a case like Bribe, my offer of money to you to withhold belief makes a difference to what you're *epistemically* justified in believing. Thus,


³ Hard pragmatism, as I've defined it, is consistent both with the view that *only* pragmatic considerations supply reasons for belief (Rinard 2017), and with the view that both evidential and pragmatic considerations supply reasons for belief (Reisner 2008; Howard 2016).

⁴ This is often implicit in the view, but for more explicit statements of it, see e.g. Owens (2000: 24, 29-31); Fantl & McGrath (2002: 81-83); Schroeder (2012a: 266, 268).

⁵ My formulation of pragmatic encroachment makes it a view about justified belief, whereas many of its advocates (especially its early advocates) make it a view about knowledge (Hawthorne 2004; Stanley 2005). I find the view about justified belief more instructive for my purposes here, where I want to compare moderate pragmatism with hard pragmatism, where neither of those views are about knowledge. Most advocates of pragmatic encroachment on knowledge also embrace pragmatic encroachment on justified belief (see, e.g., Fantl & McGrath 2002, 2009; Schroeder (2012a: 268)). Moss (2018) constitutes an exception to this, at least when it comes to credences.

pragmatic encroachers incur the explanatory burden of explaining how to keep their pragmatism moderate.⁶

It may seem initially puzzling how pragmatic considerations could make a difference to what one *epistemically* ought to believe, as pragmatic encroachers hold. If pragmatic considerations are reasons for or against belief, the character of those reasons does not seem to be epistemic in any good sense. A natural and popular answer to this challenge is to hold that pragmatic considerations do not make a difference to what one ought to believe by constituting *reasons* for or against belief at all (as they would on a view that we can **reasons pragmatism**), but rather by serving as what we might call “threshold-shifting considerations” (let’s call this latter view **threshold pragmatism**). In particular, pragmatic considerations determine *how much* evidence is required for (epistemically) justified belief.⁷ They thus play a different role to evidential considerations, which serve as reasons. I like to think of this using an illustration that I’ll call the “beaker of reasons”:




The beaker of reasons is filled up by – unsurprisingly – reasons. The shaded area represents how much reason there is to believe p , where this is a function of the reasons for and against believing p , and their weight. The higher up in the beaker that the shaded area goes, the more reason there is to believe p . But to know whether there is *sufficient* reason to believe p , we need to know not just how much reason there is, but how much reason for belief is needed in order to be all-out justified in believing p . The black dotted line represents the answer to the latter question – the threshold for justified belief. When the reasons exceed the black dotted line, there is sufficient reason to believe. In the above example, there is.

On a threshold pragmatist version of pragmatic encroachment, evidential considerations are what *fill up* the beaker – the reasons. By contrast, pragmatic considerations are what help to determine where the dotted line is – bar-shifting considerations. Thus, the thought is, in at least some cases where believing p carries significant pragmatic costs, those costs move the dotted line upward, like so:

⁶ One can imagine a view that accepts pragmatic encroachment in cases like Risks, *as well as* accepting pragmatism (but *not* pragmatic encroachment) in cases like Bribe. This would square pragmatic encroachment (about some cases) with hard pragmatism, as I’ve defined it. Still, such a view would face the same explanatory burden that moderate pragmatists face, since it would need to explain why there is pragmatic *encroachment* in cases like Risks, but not in cases like Bribe.

⁷ Views broadly along these lines are endorsed by Owens (2000: ch. 2), Fantl & McGrath (2002: esp. 88), Ganson (2008), Pace (2011), Morton & Paul (2018), Basu (2019a), and Basu & Schroeder (2019: §4.1) (though Schroeder had earlier expressed skepticism about this sort of picture in his 2012b: 479-481).



This has the effect of making it the case that there is no longer sufficient reason to believe p . But it does this not by providing reasons *against* believing p , but rather by changing how much evidence is required for believing p – shifting the epistemic “standards”. Thus, threshold pragmatism is consistent with evidentialism about reasons for belief – that is, the view that all reasons for belief are evidential.

Threshold pragmatism also gives us an explanation of the pragmatic encroacher’s contention that pragmatic considerations affect whether you are *epistemically* justified, or *epistemically* ought to believe p – since it holds that pragmatic considerations affect how much graded epistemic justification is required to count as being “on-off” epistemically justified. It is because *the considerations that justify* (the reasons) are still evidential than we can call this justification epistemic. This allows the pragmatic encroacher to maintain the orthodox view that one ought (*tout court*) to believe p iff one *epistemically* ought to believe p . Since the pragmatic considerations already affect what one epistemically ought to believe, they don’t need to separately trade off against epistemic considerations to produce an all-things-considered verdict that might differ from the epistemic verdict.

The threshold version of pragmatic encroachment is elegant and attractive. Moreover, one could, in one sense, call this view “moderate”, given its consistency with evidentialism about reasons for belief. Nevertheless, on its own, the threshold model provides no explanation of how pragmatists can be moderate in the sense in which I defined “moderate” pragmatism. Whereas the distinction between reasons and threshold pragmatism concerns the *way* in which pragmatic considerations affect what one ought to believe, the distinction between hard and moderate pragmatism concerns the *breadth of cases* in which they do so. On its own, the threshold model says nothing about *when* pragmatic considerations shift the threshold for epistemically justified belief. In particular, it does not say anything that precludes the possibility that in Bribe, the costs of believing p raise the bar for epistemically justified belief in p . Thus, as attractive as the threshold model is, on its own it is not the right sort of thing to provide a solution to our central question. We will have to look at potential views about *when* the bar is raised, and why, in order to get anything that is even a candidate solution.

2. Motivating pragmatism in moderate cases

Moderate pragmatists, as I’ve defined the view, accept that pragmatic factors make a difference to what one ought to believe in cases like Risks, but not in cases like Bribe. However, some might be initially puzzled by this ambition, thinking that it is actually *harder* to see why pragmatic factors would make a difference in the former cases than the latter.

The line of thought would be this. In *Bribe*, it is relatively obvious what the theory that underlies the putative relevance of pragmatic factors is. The theory is just this: considerations that pertain to the pragmatic value of believing p bear on whether one ought to believe p . In *Bribe*, believing that the game starts at 3pm results in your missing out on \$1,000,000. This is a disvaluable outcome, and so (on this theory) it bears on whether one ought to believe that the game starts at 3pm. While this theory is of course highly controversial, it is nothing if not simple and clear. By contrast, it might be thought, it is more obscure why pragmatic factors should make a difference to what one ought to believe in *Risks*. For, in *Risks*, the pragmatic costs come about not from *believing* that the interview starts at 3pm, but rather from *acting* in ways that take it for granted that the interview starts at 3pm (i.e., showing up at 3pm without any further checking). As such, they might seem only to pertain to whether you ought to act in those ways, not to whether you ought to *believe* anything.⁸

A related challenge comes from those who think we can do without the state of outright belief, and do everything in terms of credence. In *Risks*, even given a very high credence that the interview starts at 3pm, the expected utility calculations do not favor just showing up at 3pm without engaging in any further checking. For showing up at 3pm without checking, after all, *risks* a very bad outcome if it turns out that the interview starts at a different time. Provided one's credence that the interview starts at 3pm falls short of credence 1, these risks will get their due weight in the expected utility calculus. Consequently, there is no need to say that such risks make any difference to what your *credence* ought to be in order to explain the verdict that you should not show up at 3pm without further checking. Indeed, to adjust one's credence in response to these risks would result in their being "double-counted" in the final expected utility calculus about how to act.

These are important challenges. The leading response, which attempts to address both challenges in one fell swoop, appeals to what I'll call the **reliance-involving conception of belief**. This conception can come in either a descriptive or a normative form. On the descriptive version of the view, roughly, it is part of what it is to believe p that one relies on p (or, equivalently, takes p for granted) in one's reasoning.⁹ To rely on p in one's reasoning is to reason and act on the assumption that p is true, ignoring the chance that it is not. On the normative version of the view, believing p is the state that, when justified, *licenses* relying on p in one's reasoning.¹⁰ Thus, the normative version of the view entails that you're justified in believing p only if you're justified in relying on p .

Let me explain how these views seek to meet both of the challenges I just outlined. To begin to address the first challenge, note that if you're justified in relying on the proposition that the interview starts at 3pm proposition in your reasoning, then you are justified in showing up at 3pm without further checking. For conditional on the interview starting at 3pm, showing up at 3pm without further checking does not risk any bad outcome. Thus, the first challenger must concede not only that you shouldn't show up at 3pm without further checking, but also that you shouldn't rely on the proposition that the interview starts at 3pm in your reasoning. But given that, plus either the descriptive or normative version of the reliance-involving conception of belief, you also shouldn't

⁸ See, e.g., Jackson (2019).

⁹ Views at least broadly in the ballpark of this include Weatherson (2005, 2016), Ganson (2008), Fantl & McGrath, and (2009: ch. 5). Ross & Schroeder (2014) accept a view that is in some respects similar, but also has some important differences.

¹⁰ Views at least broadly in the ballpark of this include Fantl & McGrath (2002: 78); Bolinger (forthcoming).

believe that the interview starts at 3pm. On the descriptive version of the view, it isn't even *possible* to believe this proposition without relying on it: to cease to rely on it is *ipso facto* to cease believing it. On the normative version of the view, while such a combination is possible, it can't be justified: if you shouldn't rely on p , you shouldn't believe p either. Thus, we get an explanation of how pragmatic considerations make a difference not just to action but to belief itself in cases like Risks.

The second challenge can be seen as, more radically, rejecting the idea that we should “rely on propositions” (at least, those propositions that we're not justified in having credence 1 for) in our reasoning, where this involves ignoring the possibility that they are false. It says that we should simply reason with credences, and calculate expected utilities on that basis – and that once we do so, we do not need to allow for any influence of pragmatic considerations on our doxastic states themselves – at least not in cases like Risks. In response, pragmatic encroachers typically concede that (at least in cases like Risks), there is no pragmatic encroachment on justified credences.¹¹ Indeed, it's part of the very point of the reliance-involving conception of belief that what distinguishes belief (or justified belief) in p from a high credence (or high justified credence) in p is that the former involves (or licenses) relying on p , whereas the latter does not necessarily do so. Thus, the case for pragmatic encroachment on belief that goes via the reliance-involving conception of belief does not extend to credence.

However, the proponent of the reliance-involving conception of belief also contends that, as agents with finite processing power, we cannot reason exclusively with credences. We need (reliance-involving) full belief as a simplifying device for reasoning, one that allows us to ignore low-probability possibilities.¹² That said, when the risks of ignoring these low-probability possibilities start to mount up, we should stop ignoring them or taking their negations for granted. And that just is, on the reliance-involving conception of belief, to stop believing their negations (or to be required to do so). Again, we get a (putative) explanation of how pragmatic considerations make a difference to what one ought to believe. It is because full belief (unlike credence) involves ignoring small chances of error that it must be sensitive to the costs of error in a way that credence need not be.

Now, both the reliance-involving conception of belief and its success in meeting the two challenges just mentioned can be disputed.¹³ But in this paper, I am just going to grant it to the pragmatic encroacher for the sake of argument.

One further clarification about the reliance-involving conception of belief is important.¹⁴ I have talked of relying on p in one's reasoning, rather than on relying on p in one's actions. The proponent of the reliance-involving conception of belief might also say that you're justified in believing p only if you're justified in relying on p in your *actions* (or in “acting on p ”). But we should be

¹¹ Cf., e.g., Ganson (2008: 450, esp. n. 8); Pace (2011: 255); Ross & Schroeder (2014: 260); Bolinger (forthcoming: 16-17). Clarke (2013) dissents, holding that there is pragmatic encroachment on justified credence too; I criticize his view in Worsnip (2016: 554-7). Moss (2018: 183-190) accepts that there is no pragmatic encroachment on *justified* credence, but holds that credences can constitute (or fail to constitute) knowledge, and that there *is* pragmatic encroachment on whether one's credences constitute knowledge.

¹² See Weatherston (2012: 77-82); Ross & Schroeder (2014: 265-7); Bolinger (forthcoming: 5). The point, at its most sophisticated, is not that we cannot do sophisticated expected utility calculations. Rather it is that, doing manageable expected utility calculations requires taking certain things for granted: in identifying what the different possible states of the world are, what the outcomes in each of them are, and so on.

¹³ For a very interesting discussion, see Staffel (2019).

¹⁴ The point I make in what follows is an exact structural analogue of the one that Ichikawa (2012) makes in defending *knowledge* norms on action, specifically, against supposed counterexamples.

careful about how we interpret these kinds of claims. Everyone, even the most radical pragmatic encroachers and the most committed proponents of the reliance-involving conception of belief, should admit that sometimes, people perform unjustified actions on the basis of justified beliefs.

For example, suppose that Donald believes that enacting draconian immigration policies will galvanize his base of supporters. And suppose that this motivates him to enact draconian immigration policies: he does so on the basis of his belief that doing so will galvanize his base of supporters. Clearly, in order to say that Donald's action is unjustified, we do not need to say that Donald's belief that enacting draconian immigration policies will galvanize his base of supporters is epistemically unjustified (or, indeed, itself morally problematic). The problem is not with Donald's belief at all, but rather with the fact that the content of his belief does not actually morally justify his action.

So if we say that you are justified in believing p only if you are justified in relying on p in your actions, this should not be interpreted as meaning that you are justified in believing p only if you are justified in doing *just anything* on the basis of p . Rather, it should be interpreted as meaning that you are justified in believing p only if you are justified in doing the things that p actually supports doing, or that are (substantively) rational *conditional* on p , or that can be *correctly* reasoned to using p as a premise, or similar. To put the point another way, being justified in believing p implies that p is eligible as something on the basis of which you can, in principle, act – but in order for you to be justified in performing some *specific* action Φ on the basis of your belief in p , it must also be the case that p is actually a sufficient reason to Φ . The norm that we are permitted to act on the basis of epistemically justified beliefs does not replace the difficult business of ethics.

That said, it's possible to overplay this point. Suppose someone tried to extend the strategy that we used to explain how Donald's belief can still be epistemically justified to Risks. The claim would have to be this: perhaps, in Risks, the proposition that the interview starts at 3pm *doesn't* justify turning up at 3pm without checking first. For not much is lost by checking, and you ought to do that *just in case* the interview doesn't start at 3pm. So, the thought continues, even though you ought not to turn up at 3pm without checking, you can still be justified in believing that the interview starts at 3pm. However, at least so long as we are holding on to the reliance-involving conception of belief, this line of thought is not coherent. Relying on a proposition p in your reasoning, by its very nature, precludes making decisions on the basis of thoughts like "I ought to check just in case not- p ." To make a decision on that sort of basis is *ipso facto* to cease to rely on p . Conversely, if one really were justified in relying on the proposition that the interview starts at 3pm, showing up at 3pm without further checking *would* be the right decision to reach.

Thus, Risks is fundamentally disanalogous to the case of Donald. In the case of Donald, there is nothing wrong with Donald's belief, or with his using it in his reasoning as such; it is just that he reaches the wrong conclusions on the basis of this belief. In Risks, conversely, since the decision to arrive at 3pm without checking *would* be the right decision conditional on the interview starting at 3pm, the problem has to be located in your using that proposition as a premise in your reasoning in the first place (and thus, by extension (on the reliance-involving conception of belief), in your believing it). This disanalogy is good for the pragmatic encroacher who motivates her view via the reliance-involving conception of belief. Such a view is implausible if it predicts pragmatic encroachment even in Donald's case, but its success as an explanation of pragmatic encroachment is in doubt if it does

not predict pragmatic encroachment in paradigm cases like Risks. So it is good news that it predicts pragmatic encroachment in the latter but not the former case.

3. Further cases

We are on our way to considering strategies for keeping pragmatism moderate. However, before we can do so, we need one more thing in place. Above, I presented moderate pragmatism with reference to just two cases: moderate pragmatists accept the influence of pragmatic considerations in cases like Risks but not in cases like Bribe. But as I also said, a good explanation of how pragmatism can be kept moderate will generalize to a number of other cases that moderate pragmatists either do, or don't, want to include among the set of cases in which pragmatic considerations can make a difference. (Or, at the very least, it won't preclude the moderate pragmatist from giving her desired answers about these cases.) So I need to introduce these cases.

A. Pascal's Wager

First, a case that moderate pragmatists have typically¹⁵ wanted to *exclude* from this set: Pascal's Wager. Let us focus on a simplified, determinate version of the case:

Pascal's Wager. If God exists, then you will go to heaven iff you believe God exists, and to hell otherwise. (If God doesn't exist, then there is no afterlife.)

For many moderate pragmatists, Pascal's Wager is, just as much as Bribe, the paradigm of the sort of case in which pragmatic factors cannot make a difference to what one ought to believe: just as the prospect of gaining money for one's belief (or lack of) cannot change whether you ought to believe, so too the prospect of going to heaven or hell for one's belief (or lack of) cannot change whether you ought to believe. And as with Bribe, this is especially important for pragmatic encroachers, who want to say that pragmatic considerations make a difference to what you ought to believe by making a difference to what you're *epistemically* justified in believing. For example, the pragmatic encroacher Mark Schroeder writes:

"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. [...] 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."¹⁶

¹⁵ Though there is at least one exception, namely Benton (2018).

¹⁶ Schroeder 2012a: 266.

Throughout his paper, Schroeder continues to use the adjective ‘Pascalian’ to refer to the sort of consideration that *can’t* make a difference to what one ought to believe (at least, not via making a difference to what one is epistemically justified in believing).

Pascal’s Wager is worth introducing because, although it seems to intuitively and pretheoretically pattern with Bribe rather than Risks, it will turn out that several of the most initially promising accounts of the difference between Bribe and Risks turn out to group it (or variants of it) with Risks, thus committing themselves to the view that pragmatic considerations *do* make a difference to what one ought to believe in Pascal’s Wager. At least for moderate pragmatists with Schroeder’s ambitions, this will be an important mark against such theories.

B. *Forthcoming evidence and costs of inquiry*

The next group of cases are ones that moderate pragmatists have typically wanted to *include* among the set of cases where pragmatic considerations make a difference to what one ought to believe. These are cases where there is pertinent information about how easy or difficult it is to gather future evidence.¹⁷ It might be that further evidence about whether *p* is shortly forthcoming, or whether it will be easy to gather without any further cost to oneself. Here’s an example of such a case:

Expected News. You have strong (but not utterly infallible) evidence that next Monday’s meeting starts at 3pm. But you know that your department chair will confirm the time of the meeting later today.

In Expected News, it’s not exactly that believing that the meeting starts at 3pm, or even relying on this proposition, is costly as such. Rather, it’s that it seems silly or pointless to make up your mind about this proposition now when you’ll be getting much stronger evidence about it later via your chair. Although the fact that you will hearing from the chair later is a fact *about* (forthcoming) evidence, it isn’t itself evidence for or against the proposition that the meeting starts at 3pm. Thus, in a broad sense, it is a pragmatic consideration; at the very least, it’s a non-evidential one. Since Expected News at least *prima facie* differs from Risks in its underlying structure, however, it is an important test for theories of why pragmatic considerations do make a difference in Risks (but not in Bribe) that they don’t preclude pragmatic considerations from making a difference in Expected News.

C. *Moral encroachment cases*

There is a recent, burgeoning literature on *moral* encroachment.¹⁸ I use the term ‘pragmatic’ broadly to include moral considerations, so on my usage, moral encroachment is a kind of pragmatic

¹⁷ Schroeder (2012b) discusses such cases at length.

¹⁸ Notable contributions include Pace (2011), Fritz (2017), Moss (2018), Gardiner (2018), Schroeder (2018), Basu (2019a), and Basu & Schroeder (2019).

encroachment.¹⁹ Some cases of putative moral encroachment are exact structural analogues of cases like Risks, with costs to someone else substituted for the costs to the believing agent.²⁰ Consider:

Moral Risks. You have strong (but not utterly infallible) evidence that the train leaves at 3pm. But if you take it for granted that the train leaves at 3pm, and you're wrong, you will miss the train, and consequently miss your student's dissertation defense. Without your presence, she won't be able to receive her PhD this year, which will have serious consequences for her career.

This case is so structurally similar to Risks that it shouldn't be hard to extend any theory of why pragmatic considerations make a difference in Risks (but not Bribe) to cover Moral Risks too. So I'll set Moral Risks aside in what follows. However, there is another widely-discussed kind of case of putative moral encroachment that – at least *prima facie* – seems to have a rather different structure. An example of this kind of case is as follows:

Wedding. Andrew is at the wedding of two college friends. The wedding is black-tie, so the waiters are dressed the same as the guests. Andrew knows these two college friends pretty well, and he knows that (unfortunately and for whatever reason) nearly all their friends are white. Andrew also knows that the wedding is taking place in a city where service-industry workers are, statistically, mostly black – and he has already interacted with several black servers at the wedding. Andrew sees a black man in a tuxedo (let's call him John) walking past. On the basis of the fact that John is black (together with the background statistical information, just mentioned, that he has encoded), Andrew forms the belief that John is a waiter, and asks him to bring him a drink. In fact, John is not a waiter, but the wedding's only black guest.²¹

Intuitively, Andrew is criticizable in Wedding; indeed, it is very natural to describe his behavior as racist. Moreover (and even before we appeal to anything as theoretically loaded as the reliance-involving conception of belief), there is a strong pull to the claim that it's not just Andrew's *action*, but also his *belief*, that is criticizable: he shouldn't believe that John is a waiter on the basis of his race.

Wedding is a complex case with many features that our previous cases lack. This opens up a multiplicity of explanations of what is wrong with Andrew's belief. One broad group of explanations

¹⁹ While it's true that the original pragmatic encroachers tended to focus on prudential examples, they also didn't have any well-developed account of the relevant property of those prudential costs, in virtue of which they result in pragmatic encroachment. (Negatively impacting the putative knower's own well-being as such? Negatively impacting something the putative knower cares about? Negatively impacting *someone's* well-being? Being a disvaluable outcome of any kind?) Thus, it was left unresolved whether the view applied to moral cases. For example, though Stanley's book is entitled "Knowledge and Practical *Interests*" (my italics), his stated view is that knowledge depends on practical facts, which are "the costs of being right or wrong about one's beliefs" (Stanley 2005: 6). This does not seem to rule out moral costs. For my part, I find it unmotivated to claim that costs redounding to the agent's well-being affect the agent's epistemic justification but that other costs that carry normative weight for her actions do not affect her epistemic justification.

²⁰ Cf. Fritz (2017: 650); Basu (2019a: 16).

²¹ This case is adapted from one introduced by Gendler (2011). Though Gendler herself doesn't discuss it in connection with moral encroachment, a number of other authors (Moss 2018; Gardiner 2018; Basu 2019a; Basu & Schroeder 2019; Bolinger forthcoming) do.

attempts to locate some *purely epistemic* defect in Andrew's belief that is not connected with its moral status.²² Initially, one might worry that such a purely epistemic explanation will be hard to give, since it's a result of the stipulations of the case that, given the base rates of which Andrew is aware (regarding the high percentage of guests who are white, and the high percentage of servers who are black), the fact that John is black *is*, in this case, fairly strong probabilifying evidence that he is waiter. But in fact, even conceding this, there are still a range of purely epistemic explanations of the defect in Andrew's belief that one might pursue. For example, one might claim that the evidence here is of a purely statistical or "non-individualized" kind that in general cannot warrant outright belief,²³ or that the belief fails a sensitivity or safety condition on justified outright belief (compare lottery cases).²⁴ Or, one might take issue with the stipulation that in any realistic version of the case, Andrew's belief could really be *based* in the relevant statistical evidence, as opposed to being formed automatically in response to a socially-encoded stereotype – in which case the belief would be at least *doxastically*, even if not propositionally, irrational. Or, one might claim even if Andrew's belief *is* formed on the basis of the relevant statistical evidence, he has higher-order evidence that should make him *suspect* that his belief was formed on the basis of a stereotype and thus to become less confident in it.²⁵ This is only a sampling of ways that one might try to locate a purely epistemic defect in Andrew's belief.

On the other end of the spectrum, one might argue that the defect in Andrew's belief is *purely* moral, and is orthogonal to its epistemic status. Perhaps, in forming the belief that John is a waiter on the basis of a race, Andrew fails to treat John with the respect that personhood requires, treating him instead as an object to be predicted and classified on the basis of general, non-individualized facts.²⁶ This might ground a purely moral requirement not to form the belief that John is a waiter regardless of whether such a belief would be epistemically justified.²⁷

Both of these groups of explanations are orthogonal to the moderate pragmatist's attempt to differentiate cases like Bribe from cases like Risks, since neither view seems to commit its proponent to any particular view about either Bribe or Risks. However, some philosophers have argued that cases like Wedding should be explained as cases of moral encroachment, co-opting the broad kind of model that is used to explain why pragmatic considerations make a difference to the epistemic status of one's beliefs in Risks to likewise explain how they make a difference to the epistemic status of Andrew's belief in Wedding.²⁸ For *these* philosophers, it will be an important test of any theory of the difference between Bribe and Risks that it doesn't preclude pragmatic considerations from making a difference

²² Examples of such explanations of this or similar cases include Gardiner (2018), Munton (2019), and to some degree Fritz (forthcoming: §4).

²³ Cf. Thomson (1986).

²⁴ Cf. Enoch, Spectre & Fisher (2012).

²⁵ These last two possibilities were both suggested to me by Kara Boschert.

²⁶ Cf. Basu (2019b), though Basu herself also accepts moral encroachment and thus holds that the belief is not epistemically justified either.

²⁷ This need not necessarily involve posing a dilemma between epistemic and moral norms. The epistemic norms might be permissive with respect to whether you believe or don't believe in cases where you have strong, but not utterly infallible, evidence.

²⁸ Moss (2018), Schroeder (2018), Basu (2019a), Basu & Schroeder (2019), and Bolinger (forthcoming) all pursue moral encroachment diagnoses of the epistemic impropriety of beliefs in cases like Wedding. It's worth noting, though that these authors do vary somewhat in how tight they want the structural analogy between Risks and Wedding to be.

in Wedding. That is why I have introduced the case here, and why I consider it in what follows. However, let me be clear that even if I am right in my eventual contention that there is no good explanation of how to keep one’s pragmatism moderate, it doesn’t follow that there is nothing wrong with Andrew’s belief in Wedding – or even that the only way to contend that there is something wrong with this belief will commit one to pragmatism about Bribe. For the other explanations of what goes wrong with Andrew’s belief in Wedding remain open.²⁹

Let’s now summarize the verdicts that I take most moderate pragmatists to want in each of the cases we’ve discussed so far:

Cases where pragmatic considerations <i>shouldn’t</i> make a difference	Cases where pragmatic considerations <i>should</i> make a difference
Bribe	Risks
Pascal’s Wager	Expected News
	Moral Risks
	Wedding [<i>if</i> one is pursuing a particular explanation of what goes wrong]

This table categorizes the cases *from the point of view of the moderate pragmatist’s ambitions*. I’m not treating the moderate pragmatist’s intuitive verdicts about all of these cases as gospel. (Though, for what it’s worth, I think they are quite plausible. I find moderate pragmatism a *prima facie* compelling view, and I would genuinely like for there to be a way of vindicating it.) I am going to assume, for the sake of argument, that these are the results we want, as would-be moderate pragmatists. The question is whether we can get them all at once.

4. Strategies for keeping pragmatism moderate

In what remains, I will consider five strategies for keeping pragmatism moderate, which I have picked because they seem the most promising.

A. Reasons to believe vs. reasons to withhold

One potential strategy appeals to the following principle, endorsed by Mark Schroeder (2012b):

Schroeder’s Principle. Pragmatic, or non-evidential, factors can constitute reasons to *withhold* belief, but not positive (right-kind³⁰) reasons to *believe*.

²⁹ One might ask whether one could accept the possibility of robust, non-epistemic *moral requirements* on belief – as the purely moral explanation posits – without allowing for *prudential* reasons for belief that would apply in cases like Bribe. This is an interesting question worth pursuing, but it is not within the scope of my task here. I am interested in asking whether pragmatists can be moderate, where that necessarily involves allow pragmatic considerations to make a difference to what one ought to believe in Risks, but not in Bribe. The view that pragmatic considerations make a difference *only* in moral cases like Wedding, and not in prudential cases like Bribe or Risks, is outside of my target.

³⁰ Schroeder leaves open the possibility that non-evidential considerations can constitute *wrong kinds* of reasons for withholding belief that are nevertheless genuine reasons. He himself stays neutral on whether wrong kinds of reasons are genuine reasons (*ibid.*: 461).

Schroeder gives a case very similar to Expected News, and claims that in this case, the fact that evidence relevant to p is forthcoming – though not itself evidence either for or against p – is a (right-kind) reason to withhold belief for the moment. He contrasts this case with Pascal's Wager, expressing sympathy with the view that the positive expected utility of believing in God is *not* a (right-kind) reason to so believe, but noting that in this case, the pragmatic factor would – if it were a reason at all – be a positive reason to *believe* and not merely a reason to withhold. The upshot, he seems to think, is that we can explain why non-evidential factors make a difference in Expected News but not in Pascal's Wager by appealing to Schroeder's Principle. One might wonder if this can be generalized to explain the differences between all the relevant cases.

The most obvious problem with this strategy is that, even if it rules out pragmatic factors from making a difference in Pascal's Wager, it does not rule out pragmatic factors from making a difference in Bribe. I deliberately set Bribe up so that you are offered money to *withhold* belief, rather than to believe. Thus, the difference between belief and withholding cannot explain why pragmatic considerations do not make a difference in Bribe.

This makes it clear that Schroeder's Principle cannot be a complete explanation of the differences between the relevant cases, and thus of how pragmatists can be moderate (a point Schroeder himself may acknowledge³¹). Additionally, in my view, the principle itself is undermotivated. It might be thought that it gets support from the difference between Expected News and Pascal's Wager. But these cases do not constitute a minimal pair. Rather, a minimal pair would consist of Pascal's Wager together with:

Pascal's Wager (Modified). If God exists, then you will go to heaven iff you do not positively believe that God doesn't exist, and to hell otherwise. (If God doesn't exist, then there is no afterlife.)

Just as the moderate pragmatist should want to hold that expected utility considerations make no difference to whether you ought to believe that God exists in the original Pascal's Wager, she should also want to hold that expected utility considerations make no difference to whether you ought to believe that God *doesn't* exist in Pascal's Wager (Modified). But Schroeder's Principle can deliver only the former, and not the latter, result, since for all it says, expected utility considerations constitute a reason to *withhold* from believing that God doesn't exist in Pascal's Wager (Modified). Now, it's possible that something *else* explains why expected utility considerations don't give you a reason to withhold in Pascal's Wager (Modified). But then we might expect that factor to also explain why the expected utility considerations don't give you a reason to believe in the original case. If so, Schroeder's Principle does not do any work in explaining the difference between Expected News and Pascal's Wager either, and so this difference gives it no support.

We can challenge Schroeder's Principle even more directly by looking at a minimal pair involving Expected News together with:

³¹ See Schroeder 2012b: 472.

No More News. You have strong (but not utterly infallible) evidence that next Monday’s meeting starts at 3pm. You know for sure that there won’t be any more evidence forthcoming about when the meeting starts.

In Expected News, Schroeder thinks, you have reasons to withhold from believing a proposition p based on your knowledge that more evidence concerning p is shortly forthcoming. This counts against making up your mind now, and hence in favor of withholding. But if that is so, why can’t the fact that you know that there will be *no* evidence forthcoming regarding p – as in No More News – constitute a reason to make up your mind, and hence believe p ? (At least, against a background where you already have at least some good evidence for p .) It seems to me that if one allows non-evidential factors to give one reasons to withhold in Expected News, one should likewise allow non-evidential factors to give one reason to believe in No More News. So even Schroeder’s paradigm case, Expected News, fails to support Schroeder’s Principle. In fact, if one agrees with Schroeder that one has non-evidential reasons to withhold in Expected News, one should positively *reject* Schroeder’s Principle, in order to allow analogous non-evidential reasons to believe in No More News.

Thus, not only does Schroeder’s Principle fail to save the letter of moderate pragmatism as I defined it; it is also implausible on its own merits, and undersupported by the intuitive difference between Pascal’s Wager and Expected News. I conclude that it can’t explain how pragmatists can be moderate.

B. *Costs vs. risks*

Pragmatic encroachers often say that the epistemic standards depend on the “costs of being wrong”. If that is interpreted to refer to *any* cost of being in a state whereby (i) one believes p and (ii) p is false, then it yields pragmatic encroachment in Bribe. After all, in Bribe, if one believes that the game starts at 3pm, and this is false, one loses one’s shot at \$1,000,000. However, the pragmatic encroacher might reasonably suggest that this is not what she meant by “costs of being wrong”. In Bribe, the relevant costs occur *solely* in virtue of one’s believing p – whether or not p is false. But perhaps it is only the costs of specifically believing *falsely* that make a difference to the epistemic standards.

Sarah Moss (2018) gives us some helpful stipulative terminology here, reserving the term ‘costs of belief’ for costs that accrue whether the belief is true or false, and using ‘risks of belief’ for those costs that accrue specifically when the belief is false. Using Moss’s terminology, the putative principle is as follows:

Risks/Costs Principle. The risks of believing p make a difference to whether one ought to believe p , but the costs of believing p (in Moss’s sense of ‘costs’) do not make a difference to whether one ought to believe p .³²

³² I originally read Moss (2018) as herself endorsing Risks/Costs Principle, but this was too hasty. Moss describes the moral *costs* of beliefs, in contrast to moral *risks*, as “not the right sort of moral features to ground moral encroachment on the epistemic status of those beliefs” (*ibid.*: 195). And in supporting this contention about the moral case, Moss

This principle has considerable theoretical appeal. Many, following James (1896/2000), think that there are two fundamental epistemic goals: believing truth, and avoiding error. When one has a moderate amount of evidence for a proposition, one has to weigh those two goals. How much evidence is needed to justify belief will depend on the relative priority of the two goals.³³ It's natural to think that the importance of avoiding error depends on how costly it is to make an error, i.e. to believe falsely. (Similarly, the importance of believing truly depends on how costly it is to fail to believe a true proposition.) Thus, we have a natural picture of how pragmatic considerations affect the relative priority of the two distinctively *epistemic* goals, believing truth and avoiding error. When pragmatic considerations influence our beliefs in these ways, we are still being guided by those two goals: we are just changing how we weigh them against each other. By contrast, if we are simply trying to avoid beliefs that are *costly* in Moss's sense – that is, costly *whether true or false* – then we are no longer (solely) in the game of aiming to believe truth and avoid error. We are specifically aiming to avoid costly beliefs, *whether or not* they are false. This cannot be plausibly thought of as an *epistemic* goal. Thus, by embracing pragmatic encroachment in cases of *risky* belief, but not in cases of *costly* belief, Risks/Costs Principle sits nicely with the pragmatic encroacher's claim to still be discussing *epistemic* justification.

Risks/Costs Principle also, clearly, distinguishes Bribe and Risk nicely. Whereas Bribe is a case of costly belief, Risk is a case of risky belief. So the principle explains why pragmatic considerations could help to settle what you ought to believe in Risk without doing so in Bribe.

A natural extension of the view can also plausibly explain why the relevant non-evidential considerations make a difference in Expected News. The Jamesian logic brings out how another set of risks is also relevant, namely the risks of failing to believe a true proposition. What determines the relative priority of believing truth and avoiding error is the *relative* risk of believing p (i.e., the costs of believing p specifically when p is false) as compared with the risk of failing to believe p (i.e., the costs of failing to believe p specifically when p is true). Thus, it is natural for the pragmatic encroacher to say that what makes the epistemic standards high is the risks of believing being much *greater* than the risks of failing to believe.³⁴ This is consistent with Risks/Costs Principle as stated – it is still the case that, *ceteris paribus*, the risks of believing p make a difference to whether one ought to believe p . We

appeals by analogy to the claim that financial costs (that are not risks) cannot make a difference to the epistemic status of a belief (*ibid.*: 196). This presentation seems to implicate the more general claim that the costs of a belief (in Moss's sense) cannot make a difference to the epistemic status of that belief. However, it does not commit Moss to claiming that the risks of a belief always do make a difference to its epistemic status, as Risks/Costs Principle also claims. Moreover, in personal correspondence, Moss clarified that she does not accept this principle in full generality. Still, the principle is still a natural one for us to consider here, since it is a plausible interpretation of the pragmatic encroacher's slogan that the epistemic standards depend on the costs of being wrong, and it seems an initially promising way of explaining the differences between some of our central cases.

³³ See also Owens (2000: 26); Riggs (2003); Pace (2011: 253-4).

³⁴ Cf. Pace (2011: 257); Weatherson (2011: 591-2); Schroeder (2012a: 277-8); Worsnip (2015: 311-2). Both Weatherson and Schroeder point out that there is some oddness to referring to the *balance* of the risks of believing against the risks of failing to believe as “the stakes”, since on the ordinary meaning of “stakes”, a high-stakes situation is one where the risks in general are high, not one in which a certain kind of risk is much higher than another kind of risk. On the mismatch between the intuitive notion of ‘stakes’ and the notion that pragmatic encroachers need, see also Anderson & Hawthorne (2019).

just add that risks of *failing* to believe p also make a difference to whether one ought to believe p . As the former set of risks drive the epistemic standards up, the latter drive them down.³⁵

Now, in Expected News, the risk of failing to believe *now* – that is, the cost of failing to believe p now in the event that p is true – is very small. Since you are going to find out when the meeting starts for sure later today, there is simply no need to form a belief now. We didn't stipulate anything about what the risks of believing are, but even if they're not especially high, most beliefs carry at least *some* risk if they're false. For example, maybe if you form a false belief about when the meeting starts now, you might organize some other commitments for the actual time of the meeting, which you'll later have to reorganize. That's not a huge deal, but the point is that since the risks of *failing* to believe are nearly zero, these moderate costs of believing are still high *relative* to the risk of failing to believe. And, all other things equal, that drives the epistemic standards up, on the present view. So the present view can accommodate the claim that there is pragmatic encroachment in Expected News as well.

What about Wedding? As Moss points out, not all of the potential ways of diagnosing what's wrong with Andrew's belief in Wedding show it to be *risky*, rather than *costly*. In particular, at least on some accounts, it is morally wrong to believe that John is a waiter on the basis of his race, irrespective of whether, as a matter of fact, it is true that he is a waiter.³⁶ If that is so, then the fact that the belief commits a moral infraction is, in Moss's terms, a *cost* of this belief, not a *risk* of it. Thus, given Risks/Costs Principle, the fact that the belief is morally wrong should not make a difference to whether one ought to have it. However, Moss also holds that "there are moral harms that bear a distinctive connection to *false* racial profiling" (Moss 2018: 196; my italics). This is not to deny that *some* moral harms accrue whether or not the belief is false, but to say there are *additional* moral harms that accrue only when the belief is false.

What does it mean to say that there are "additional moral harms" that accrue only when Andrew's belief is false? One interpretation would have it that Andrew's belief is actually *more gravely morally wrong* in the case where it turns out to be false. This would commit us to a fairly robust kind of moral luck for belief – such that there is a pair of cases (the case where Andrew's belief is true, and that where it is false), where everything seems precisely identical to Andrew (both regarding his evidence for his belief and regarding the belief's potential practical effects), and yet the moral status of his belief differs. A second, different interpretation would have it that Andrew's belief has worse consequences (or harmful effects) in the case where it is false than the case in which it is true, where these consequences are morally salient. For example, it may be that when Andrew's belief that John is a waiter is false, his behavior based on that belief (viz., asking John to get him a drink) will insult or

³⁵ That said, it's not simple to devise cases where there are risks of *failing* to believe p . One candidate might be "forced choice" cases (Schroeder 2012a: 278; Worsnip 2015: 310, 313, 317-8), where (for example) you just have to act as if p because p is your best guess and you don't have any time to do any further checking. However, such cases aren't conclusive, because it can always be protested that in such cases, there's no risk of failing to believe p , even given a belief-reliance connection, since it suffices to rely on a proposition like *probably* p (cf. Hawthorne & Stanley 2008: 581-2). A perhaps better set of candidate cases are those where we might morally wrong someone by failing to believe them, such as testimony of victims of some kinds of crimes, or of persons from marginalized groups (cf. Crewe & Ichikawa forthcoming).

³⁶ Basu (2019b); Bolinger (forthcoming: 9). Schroeder (2018), conversely, holds that a belief can wrong someone only when it is false.

stigmatize John more than it will do when this belief is true. However, on this second interpretation, Andrew's belief is equally morally *wrong* whether it is true or false. For in both cases Andrew risks the bad consequences that come about when the belief is false; it's a matter of luck in the case where the belief is true that these bad consequences don't come about.³⁷

Either of these ways of going, if it can be made to work, squares Risks/Costs Principle with the verdict that there is pragmatic encroachment in Wedding. For my own part, I find the second way of going more attractive than the first, since I find the kind of moral luck that the former posits unpalatable. However, we should note that if we go the second way, we cannot say that there is pragmatic encroachment in Wedding *because Andrew's belief is morally wrong*. For, on the second way of going, the fact that the belief is morally wrong is a cost, not a risk, since the belief is equally wrong whether it is true or false. The risk involved is not the moral status of the belief, but its harmful consequences (such as, though not limited to, its insulting or stigmatizing John); and more particularly, the harmful consequences that arise *specifically when it is false*. So, on this view, the pragmatic encroachment in Wedding occurs not because the belief is wrong, but because (and only to the extent that) it has worse harmful consequences when it is false than when it is true. While this may seem like a subtle difference, there are some moral encroachers who think that the latter claim doesn't capture the full extent, range and grounds of moral encroachment.³⁸ These moral encroachers must either reject Risks/Costs Principle or accept the kind of robust moral luck entertained by the first way of squaring this principle with moral encroachment in Wedding.

That is a rather subtle objection to Risks/Costs Principle, which will only move people with fairly specific theoretical commitments. However, there is also a much simpler objection to Risks/Costs Principle. As we saw, Risks/Costs Principle can explain why pragmatic considerations do not make a difference in Bribe. However, it cannot explain why they do not make a difference in Pascal's Wager (Modified). In the setup of that case, believing that God does not exist results in considerable costs (namely, your going to hell) only if that belief is false. So this belief is a "risky" rather than a "costly" belief in Moss's sense – and so, Risks/Costs Principle counts Pascal's Wager (Modified) as a case where pragmatic considerations *can* make a difference to whether you (epistemically!) ought to have this belief (as opposed to withholding). Moreover, given the extension of the Risks/Costs Principle to risks of failing to have a true belief, a parallel result holds for the original Pascal's Wager case, where failing to believe that God does exist results in your going to hell only when God does exist.³⁹

³⁷ See also Bolinger (forthcoming: 9).

³⁸ In particular, Basu (2019a, 2019b).

³⁹ It should be admitted that whether Risks/Costs Principle will say that there is *actually* pragmatic encroachment in Pascal's Wager cases will depend on the correct characterization of the force of our evidence in such cases. On most pragmatic encroachment models, pragmatic considerations only help to determine whether you ought to believe some proposition *p* in cases where you have pretty good, but not utterly infallible, evidence for *p*. Thus, we'll only get pragmatic encroachment in Pascal's Wager-Modified if we have pretty good, but not utterly infallible, evidence that God doesn't exist, such that raising the threshold for justified belief means that we are no longer justified in believing this. (If we have utterly infallible evidence that God doesn't exist, we'll still be justified even given the pragmatic factors; but if we don't have at least *pretty* good evidence that God doesn't exist, then we wouldn't be justified even if these pragmatic factors weren't present). Likewise, for pragmatic encroachment in the original Pascal's Wager case, we'll need at least a minimal amount of evidence that God *does* exist; otherwise it still won't be enough even if pragmatic considerations

Perhaps a fan of the present view could bite the bullet on Pascal’s wager, content that her account at least rules out *some* cases, like Bribe, out.⁴⁰ Still, it bears stressing that one can also devise bribe cases that are structurally similar to Pascal’s Wager (Modified), in which the Risks/Costs Principle likewise entails that pragmatic considerations can make a difference to what one ought to believe.⁴¹ Suppose I make you the following offer: I will pay you \$1,000,000 if you withhold from believing p and p turns out to be false. (Call this **Modified Bribe**.) In this case, there *are* significant risks, in Moss’s sense, to believing p – namely, the opportunity cost of losing \$1,000,000 – and so the Risks/Costs Principle says that my bribe can make a difference to whether you ought to believe p . To me, Modified Bribe doesn’t feel very intuitively different from Bribe, and so I think it is a significant cost (at least from the point of view of the ambitions of moderate pragmatism) that the Risks/Costs Principle has to say that that pragmatic considerations make a difference in cases like Modified Bribe.

Moreover, Risks/Costs Principle delivers other verdicts that, in combination, are highly implausible, regardless of whether one is a moderate pragmatist or what one’s theoretical ambitions are. Consider the following sequence of cases:

Betting Sequence.

- *Case 1.* If you believe p truly, you win \$1,000,000. If you believe p falsely, you lose \$1,000,000. [If you don’t believe p (where this includes either withholding or believing not- p), you don’t win or lose anything.]
- *Case 2.* If you believe p truly, you win \$1. If you believe p falsely, you lose \$1,000,000. [If you don’t believe p , you don’t win or lose anything.]
- *Case 3.* If you believe p truly, you lose \$300,000. If you believe p falsely, you lose \$1,000,000. [If you don’t believe p , you don’t win or lose anything.]

Now, here are the risks of believing p (i.e., the costs of believing p that arise specifically when p is false), and the risks of not believing p (i.e. the costs of failing to believe p that arise specifically *when* p is true), for each case:

- Case 1:
 - Risks of believing p : lose \$1,000,000
 - Risks of failing to believe p : lose \$1,000,000 [forgone benefit, i.e. “opportunity cost”]
- Case 2:
 - Risks of believing p : lose \$1,000,000
 - Risks of failing to believe p : lose \$1
- Case 3:

lower the threshold. Arguably, then, no set of evidence will give us pragmatic encroachment in *both* Pascal’s Wager *and* Pascal’s Wager-Modified. But this is fine for the purposes of the objection, which doesn’t require pragmatic encroachment in *both* cases. Moreover, I think it’s objectionable enough if a view says that there *can* be pragmatic encroachment in Pascal’s Wager cases, depending on one’s evidence, even if this possibility doesn’t obtain in the actual world. Thanks to Melanie Sarzano for prompting this clarification.

⁴⁰ As mentioned earlier, Benton (2018) embraces pragmatic encroachment in Pascal’s Wager cases.

⁴¹ Thanks to Matt Kotzen for pointing this out to me.

- Risks of believing p : lose \$700,000⁴²
- Risks of failing to believe p : none

On the present view, recall, the epistemic standards depend on the *relative* risk of believing p compared with the risk of failing to believe p . In case 1, these risks cancel each other out, so the standards are normal. In case 2, the risks of believing p dwarf the risks of failing to believe p , so the standards go up as compared with case 1. So far, so good. But now consider what happens in case 3. In case 3, the difference between the risks of believing p and the risks of failing to believe p is *smaller* than it is case 2. Thus, the standards go *down*, compared with case 2.

But this is bizarre. The bet in Case 1 is a decent one: high-risk, but high-reward.⁴³ The bet in Case 2 is a bad one: high-risk, but very low reward. But the “bet” in Case 3 is a true stinker: taking it (i.e. believing p) guarantees a loss either way (i.e., whether p is true or false), so that the best you can hope for is that you “only” lose \$300,000. The present view says that, when the bet of believing goes from being reasonable to unfavorable, the standards for belief go up, but then when it goes from unfavorable to *incredibly* unfavorable, the standards for belief go back down, comparatively speaking. Indeed, it seems that there will be some level of evidence for p such that, if we hold that level of evidence for p fixed across all three cases, the present view recommends believing p in Case 1 (i.e., taking the reasonable bet), and withholding from p in Case 2 (i.e., not taking the unfavorable bet), but believing p in Case 3 (i.e., taking the incredibly unfavorable bet).⁴⁴

Indeed, this result can be extended.⁴⁵ The *worse* we make the bet in Case 3, the further the epistemic standards will go down. So, if we change the case so that you lose \$400,000 if you believe p truly, rather than just \$300,000, the risks of believing p will correspondingly go down from \$700,000 to \$600,000, driving the epistemic standards down. And so on.⁴⁶ Perversely, the worse the bet of believing p gets, the easier it is (according to the present view) to be justified in believing it.

This, I think, is an unacceptable result. Perhaps the monetary effects of believing just never make a difference to what one ought to believe. In that case, one should deny that the shift from Case 1 to Case 2 raises the epistemic standards. But the present view does not yield this result: it allows that the shift from Case 1 to Case 2 *does* raise the epistemic standards. My claim is conditional: that *if* one thinks that the shift from Case 1 to Case 2 raises the epistemic standards, then one should also think

⁴² Remember that risks of believing p have to be costs of believing p that arise *specifically when p is false*. Since you already lose \$300,000 when you believe p and p is true, there are only \$700,000 of further costs of believing p that arise specifically when p is false.

⁴³ Though I talk of “bets” in describing the three cases, it is important to remember that in all three cases, it is the mere *having* of the belief that amounts to making/taking the bet, since the rewards and risks accrue merely in virtue of having the belief. The cases thus differ from conventional betting scenarios, in which the choice as to whether to take the bet on p 's truth is a *further* decision that one might take, distinguishable from having the belief itself.

⁴⁴ Incidentally, the problem is not solved if we revert to a view where *only* the risks of believing p matter, and the risks of failing to believe p don't. As you can see from the chart of risks above, such a view would entail that the standards are the *same* in Cases 1 and 2, and then go down in Case 3. This is even worse.

⁴⁵ Thanks to Jim Joyce for pointing this out to me.

⁴⁶ Eventually, we reach the case where you lose \$1,000,000 if you believe p , whether or not it is true or false (a case that is just like *Bribe*, except that it involves a “stick” for believing, rather than a “carrot” for withholding). In this case there are no risks to believing p at all (in Moss's sense – no costs of believing p that arise *specifically when p is false*), and so we will be back at a situation, like that of Case 1, where the risks of believing p and the risks of failing to believe p are equal – and so, by the lights of the Risks/Costs Principle, where the standards are back to normal.

that the shift from Case 2 to Case 3 further raises the epistemic standards (or, at the very least, doesn't lower the epistemic standards). The present view violates this conditional claim.

C. *Costs that make a difference to what it's rational to prefer vs. those that don't*

I next want to discuss a view suggested by Jeremy Fantl & Matthew McGrath, who explicitly consider how to avoid pragmatic encroachment in cases like Bribe. As we'll see, Fantl & McGrath's view has strikingly similar results to the view that we just considered in the previous subsection.

Fantl & McGrath endorse the following principle:

Fantl-McGrath Principle. S is justified in believing p only if S is rational to prefer as if p .
(2002: 77)

Essentially, what this means is that you are justified in believing p only if what it's rational to prefer *in fact* is the same as what it would be rational to prefer conditional on (certainty in) p .

Suppose one is in a case where one is initially justified in believing p and rational to prefer as if p , and we then introduce a practical cost that makes it the case that one is no longer rational to prefer as if p . Fantl & McGrath Principle entails that the introduction of this cost also makes it the case you are no longer justified in believing p . Moreover, though this isn't a strict entailment of their principle, Fantl & McGrath also suggest that if some cost doesn't make a difference to whether it's rational to prefer as if p , it doesn't make a difference to what one is justified in believing p .

As Fantl & McGrath show, their principle handles the Bribe case nicely. Suppose I've offered you \$1,000,000 not to believe that the game starts at 3pm. Now, even conditional on the game starting at 3pm, it is still rational for you to *prefer that you not believe* that the game starts at 3pm, assuming you care (or should care) a lot more about getting \$1,000,000 than about having a true belief about this matter.⁴⁷ Given what Fantl & McGrath mean by 'prefer as if p ', this means that in this particular case, to 'prefer as if the game starts at 3pm' is (though this sounds slightly paradoxical) to prefer that you not believe that the game starts at 3pm. And, given my monetary offer, you are in fact rational to prefer that you not believe that the game starts at 3pm. Thus, in this case, it is rational to prefer as if the game starts at 3pm: what it's rational for you to prefer in fact is the same as what it would be rational to you to prefer conditional on p . And thus, the necessary condition on justification that Fantl-McGrath Principle posits is met. My monetary offer thus does not preclude you from being justified in believing that the game starts at 3pm.

Unfortunately, however, Fantl-McGrath Principle has all the same problematic results as Risks/Costs Principle. First, it delivers pragmatic encroachment in Pascal's Wager (Modified). Recall that in this case, if God exists, then you will go to heaven iff you do not positively believe that God doesn't exist, and to hell otherwise. Let p be the proposition that God *doesn't* exist. In this case, you are in fact rational to prefer *not* to believe p , since the belief that God doesn't exist risks sending you to hell. However, conditional on p – conditional on God *not* existing – it is rational to prefer that you

⁴⁷ Even if bribes don't change what it's rational to believe, it should be uncontroversial that they can change what it's rational to *prefer* to believe.

do believe p . After all, if God doesn't exist, there's no cost to accurately believing that this is so. Thus, to prefer as if p , in this case, is to prefer to believe p (or at the very least, to be indifferent as to whether you do so). So, since you aren't *in fact* rational to prefer to believe p (given the setup of the case), you aren't rational to prefer as if p . Thus, by Fantl-McGrath Principle, you aren't justified in believing p either. So their principle commits them to pragmatic encroachment in Pascal's Wager (Modified).

The same thing goes for Modified Bribe. In this case, where I'll pay out \$1,000,000 if you withhold from p and p is false, it's in fact rational for you to prefer to withhold from p . But conditional on p 's truth, it's rational for you to prefer to believe p (since you can't get any money by withholding from p). Thus, again, you're not rational to prefer as if p in this case, which means by Fantl-McGrath Principle that you're not justified in believing p .

Finally, Fantl-McGrath Principle gives the same problematic combination of verdicts in Betting Sequence that Risks/Costs Principle did:

- In Case 1, recall, if you believe p truly you win \$1,000,000 and if you believe p falsely you lose \$1,000,000. Assuming you have pretty good evidence for p , you're rational to prefer to believe p in this case. And conditional on p , it's rational to believe p . So what it's rational for you to prefer in fact is the same as what it would be rational to prefer conditional on p : you are rational to prefer as if p . Thus, by the lights of Fantl-McGrath Principle, there's nothing about your practical situation that precludes you from being justified in believing p .
- In Case 2, if you believe p truly you win \$1 and if you believe p falsely you lose \$1,000,000. In this case, it's rational for you to prefer not to believe p . But conditional on p , it's rational to believe p . So what it's rational for you to prefer in fact is not the same as what it would be rational to prefer conditional on p : you are not rational to prefer as if p . Thus, by the lights of Fantl-McGrath Principle, your practical situation precludes you from being justified in believing p .
- In Case 3, if you believe p truly you lose \$300,000 and if you believe p falsely you lose \$1,000,000. In this case, it's rational for you to prefer not to believe p . And, conditional on p , it's still rational for you to prefer not to believe p . So what it's rational for you to prefer in fact is the same as what it would be rational to prefer conditional on p : you are rational to prefer as if p . Thus, by the lights of Fantl-McGrath Principle, there's nothing about your practical situation that precludes you from being justified in believing p .

Assume that you meet whatever other conditions are required for being justified in believing p in all three cases. Then, Fantl-McGrath Principle yields the verdict that you're justified in believing p in Case 1, but not in Case 2 – and yet that in Case 3, you go back to being justified in believing p . Again, this seems to me highly implausible.

D. *W-stakes vs. A-stakes*

One way to take the upshot of the last two sections is that it is hard to draw a satisfying dividing line between the costs (and benefits) of believing p that do make a difference to what we should believe, and those that don't. So a natural question to ask is whether there is a version of the pragmatic

encroachment view that is not centered around the costs and benefits of believing p at all. And in a sense, I think there is. Pragmatic encroachers sometimes say that the stringency of the epistemic standards depends upon how high the “stakes” are. In previous work (Worsnip 2015), I distinguished two interpretations of the notion of ‘stakes’ that is employed in this slogan. The A-stakes interpretation focuses on what is at stake as to what *attitude* – viz. belief, withholding, etc. – the subject takes toward a proposition p . The W-stakes interpretation focuses on what is at stake as to whether that proposition is *true or false* – i.e. how the world is. So, it might be thought, a W-stakes version of the pragmatic encroachment view might be in a better position to avoid the slide into hard pragmatism.

In a way, the view considered in §4B was a sort of hybrid of the A-stakes and W-stakes views, since it held that the costs that matter are the costs of believing p that arise *only when p is false*. In the light of the difficulties with this account, though, we might investigate a purer W-stakes view, according to which what matters are not the costs *of believing p* at all, but merely the costs of p being false.

Now, we have to be careful in how we state this view. The idea cannot be that in order for the stakes to be high (so as to raise the epistemic standards), it must be costly for p to be false *irrespective* of whether one believes p or not. For if that is so, the view cannot capture the idea that there is pragmatic encroachment in cases like Risks or Wedding. In Risks – where p is the proposition that the interview starts at 3pm – it isn’t, in and of itself, costly for p to be false. For disaster to strike, one must also believe (and rely on) p . Similarly, in Wedding – where p is the proposition that Andrew is a waiter – it isn’t, in and of itself, costly for p to be false. For the relevant costs to obtain, one must also believe (and, possibly, rely on) p . So, if the W-stakes view is to allow for pragmatic encroachment in these cases, it cannot demand that it must be costly for p to be false *both* when one believes p and when one doesn’t. Rather, the view is something like the following: for the stakes to be high (in a way that drives up the epistemic standards), it must be the case that there is at least one attitude toward p such that, holding fixed that one has that attitude, it makes a big difference whether p is true or false.^{48, 49}

With that clarification, the account *does* say that Risks and Wedding involve high stakes.⁵⁰ But it *also* allows that cases where it makes a big difference whether p , irrespective of what you believe, involve high stakes. For example, consider a case in which you will be given \$1,000,000 if p , and you will lose \$1,000,000 if not- p – regardless of what attitude you take toward p . Here, it is true that there is at least one attitude such that, holding fixed that you have that attitude toward p , it makes a big difference to you whether p . Holding fixed that you believe p , it makes a big difference to you whether p – and likewise for any other attitude you might have toward p . So, on the W-stakes view, the stakes are high, and so correspondingly are the epistemic standards. This verdict is a distinctive feature of

⁴⁸ This sort of idea might be suggested by Stanley’s (2005: 6) talk of “the costs of being right or wrong about one’s beliefs”.

⁴⁹ Is this just the view from section 4B, stated differently? I see why it’s tempting to think that, but it isn’t, as can be illustrated by its verdicts about cases. See n. 51 & 53 below.

⁵⁰ In the case of Wedding, this relies on pulling the same kind of move that we made on behalf of the view considered in section 4B.

the W-stakes view,⁵¹ but I don't have a clear intuition about whether it is an advantage or a disadvantage.⁵²

Like the last two views, the W-stakes view nicely explains why there is no pragmatic encroachment in Bribe. In Bribe, holding fixed that you believe that the game starts at 3pm, it makes no difference whether this proposition is true or false – and likewise if you withhold from it (or disbelieve it). However, also like the last two views, the W-stakes view cannot say the same thing about Pascal's Wager cases. Holding fixed that you believe that God exists, it does make a big difference whether God exists or not: it is the difference between your going to heaven and death being the end of everything.

Moreover, this view is also subject to an objection that did not apply to the last two views. Return to the first two cases of Betting Sequence:

Case 1. If you believe p truly, you win \$1,000,000. If you believe p falsely, you lose \$1,000,000. [If you don't believe p (where this includes either withholding or believing not- p), you don't win or lose anything.]

Case 2. If you believe p truly, you win \$1. If you believe p falsely, you lose \$1,000,000. [If you don't believe p , you don't win or lose anything.]

In Case 1, holding fixed that you believe p , the difference between p being true and p being false is \$2,000,000 (since the former involves *winning* \$1,000,000 and the latter involves *losing* \$1,000,000). By contrast, in Case 2, holding fixed that you believe p , the difference between p being true and p being false is \$1,000,0001 (since the former involves winning \$1 and the latter involving losing \$1,000,000). (In neither case does it matter whether p is true if you do not believe p .) Thus, the W-stakes view says that the stakes are *higher* in Case 1 than in Case 2, and correspondingly, that the epistemic standards are higher.⁵³ I do think that there is a good, ordinary sense in which the stakes *are* higher in Case 1 than in Case 2, but it is bizarre to say that the epistemic standards are higher in Case 1 than in Case 2, or that (potentially, for some body of evidence held fixed across the two cases) you should believe p

⁵¹ Here the W-stakes view contrasts not just with the straightforward A-stakes view but also with the view considered in section 4B. On that view, recall, the epistemic standards are high when believing p is "risky", where that means that if (but only if) p is false, there are significant costs to believing p . This is not true in the present case: if p is false, disaster occurs *whether or not* you believe p . In no good sense is believing p risky: whether or not p is true, one's believing p makes no difference.

⁵² In Worsnip (2015), I said that the W-stakes account makes the right prediction about our intuitions about *knowledge* in cases like this one. I do think that we are, *ceteris paribus*, less inclined to self-attribute knowledge of propositions when it makes a big difference to us whether they are true, even when there are no actual costs to believing or relying on them. Still, there are other explanations of these intuitions that don't allow that high W-stakes genuinely preclude knowledge (as I noted in the original paper; see Worsnip 2015: 314). Moreover, even if we were to accept that high W-stakes preclude knowledge, this is not to concede that they preclude justified belief: this might be a case where knowledge and justified belief come interestingly apart.

⁵³ By contrast, recall, the view considered in section 4B yielded the more intuitive result that the epistemic standards are higher in Case 2 than Case 1.

in Case 2 but not Case 1.⁵⁴ Once again, the verdict that all these costs and benefits have no effect of any kind is a respectable view, but if they *do* have an effect, intuitively it should be the other way round.

E. Direct costs of belief vs. costs of reliance

I will consider one final attempt to keep pragmatism moderate. In §2, I considered the worry for moderate pragmatists that it is actually *harder* to explain the relevance of pragmatic considerations to what one ought to believe in cases like Risks than in cases like Bribe, because whereas the latter involve costs of *belief*, the former only involve costs of *relying* on what one believes in one's reasoning and actions. In brief, the response on behalf of the pragmatic encroacher was to invoke a reliance-involving account of belief that explains why the costs of reliance *are* relevant to what one ought to believe. However, the moderate pragmatist might try to go on the offensive on this point, holding that in fact the costs of relying on one's beliefs are the *only* costs that make a difference to what one ought to believe; the direct costs of the beliefs themselves do not make a difference.

This proposal does fairly well in our central cases. It rules out the effect of pragmatic considerations in both Bribe and (the aforementioned versions of) Pascal's Wager, where the costs in question are direct costs of believing, and not of relying on one's belief. And it allows pragmatic considerations to have an effect in Risks, where the costs are of relying on one's belief, rather than directly of believing. Similarly in Expected News, where (as we saw in §4B) the costs of relying on the proposition that the meeting will start at 3pm, even if not very high in absolute terms, are high *relative to* the costs of failing to rely on that proposition.

As with the view that embraced Risks/Costs Principle, Wedding is more complicated. Some have wanted to say that the moral defect in Wedding is in the *belief itself*, and not merely in relying on the belief: Andrew's belief that John is a waiter (on the basis of his race) itself *wrongs* John, independently of his treating him any particular way.⁵⁵ The picture here is a deontological one where, much as a lie can (on some views) wrong a hearer independently of whether the hearer discovers that it is a lie, a belief can wrong someone independently of whether the person whom the belief is about discovers the belief through the believer's words or actions. Basu & Schroeder (2019) make the intuitive case for this view particularly vivid, pointing out that when your belief morally wrongs someone (e.g. by being based on a racial generalization), it often seems inadequate to apologize merely for *manifesting* your belief in your actions; it seems that you owe an apology for the *belief itself*. (Put tentatively: "sorry that I let my racism slip" is not an adequate apology!) If that's so, then it seems that the present proposal cannot account for the thought that the (central) moral defect of Andrew's belief in Wedding makes it the case that he ought not have that belief. For the present proposal says

⁵⁴ Thus, if the W-stakes account is the correct account of the ordinary meaning of the English word 'stakes' – as it may well be – then the view that whether one ought to believe depends on how high the stakes are, in this ordinary sense, is not very plausible. This reinforces the point made by Schroeder and Weatherston (see n. 34 above) about the misleading tendencies of the pragmatic encroacher's slogan that the epistemic standards "depend on the stakes". Arguably, "A-stakes" accounts are more perspicuously formulated using entirely different terminology than that of stakes.

⁵⁵ Most prominently, Basu (2019a, 2019b); Basu & Schroeder (2019).

that the *direct* costs of *belief* make no difference to what you ought to believe; it is only the costs of relying on your belief in action that do that.⁵⁶

Additionally, it is somewhat hard to see what the theoretical rationale for the present view is: why would it be only costs of belief that are generated *via* costs of reliance that make a difference to what you ought to believe, and not also the direct costs of belief? Perhaps returning to the explanation from §2 of why the former make a difference – *via* appeal to a reliance-involving conception of belief – might help. Recall that the reliance-involving conception of belief came in two versions. The first version was descriptive: believing *p* involves relying on *p*. On its own, this does not seem to help much with explaining why costs of reliance would make a difference while costs of belief itself would not. On the contrary: this view's explanation of why the costs of reliance make a difference seems to go *via* showing that costs of reliance just are, *ipso facto*, costs of belief (since belief involves reliance); and that seems to presuppose that costs of belief, more generally, *do* make a difference to what one ought to believe.

The normative version of the reliance-involving conception of belief might seem more promising here. On this view, you're justified in believing *p* only if you're justified in relying on *p*. This might seem to give a direct explanation of why costs of reliance impact your justification for believing *p*, that does not obviously generalize to the costs of belief itself. Now, it is questionable how deep this explanation goes, for the normative principle that delivers it itself presumes that practical factors that impact whether it's rational to rely on *p* can make a difference to whether one is epistemically justified in believing *p*, and we might wonder *why* that is so if practical factors of other kinds don't make a difference to epistemic justification in general. But suppose we have been satisfied on that score. Still, it seems that a suitable *generalization* of the relevant normative principle will let the direct costs of believing make a difference in at least some of cases.

Fantl-McGrath Principle, which we already met with in section 4C, and which is a version of the normative version of the reliance-involving view, illustrates this nicely. Fantl & McGrath see that it would be *ad hoc* to *only* endorse a principle that applies solely to what it's rational to rely on in one's reasoning about how to act. Rather, such a principle falls out of a more general principle connecting epistemic justification with what's it's rational to *prefer*, where the objects of preferences might be acts, but alternatively might be propositional attitudes or emotional responses. But once we say you're justified in believing *p* only if you're rational to prefer as if *p*, as we saw in section 4C, we get pragmatic encroachment in cases like Pascal's Wager (Modified) and Modified Bribe. As I said earlier, I think that's bad news for the strategy of using Fantl-McGrath Principle to explain how pragmatism can be kept moderate. But the point for our purposes now is just that since Pascal's Wager (Modified) and Modified Bribe are cases where the relevant costs attach directly to one's *believing* *p* itself, and not merely to one's relying on *p* in one's reasoning, Fantl-McGrath Principle does not vindicate the current strategy we're considering, on which only the costs of relying on *p* in one's reasoning can make a difference, and the direct costs of believing do not.

⁵⁶ This is similar to the situation we saw for the view considered in §4B, but the issues for the two views are not quite identical. That view could only acknowledge costs of the belief *that obtain only when the belief is false*. The present view can only acknowledge costs of the belief *that are generated via relying on it*. These features are not the same.

Are there any other candidate explanations of why that would be so? In a brief but suggestive passage, Schroeder (2012a: 277) makes a suggestion that could be read as giving such an explanation. Schroeder stresses that, on his view, it's part of the *nature* or *function* of belief that it involves reliance (or the guidance of action).⁵⁷ He concludes from this that the costs of relying on *p* when *p* turns out to be false are costs of the belief in *p* that are “due to its playing its normal role as a belief – the sort of cost that is intrinsic to the nature of belief.”⁵⁸ The implicit contrast, I take it, is that direct costs of belief, for example those that arise in Bribe and Pascal's Wager, are not in the same way intrinsic to the nature of belief. And only costs of belief that are intrinsic to the nature of belief, on his view, can make a difference to the epistemic standards for belief.

However, I find the idea here a little obscure. Again, it seems clear that if it is part of the nature of belief that belief involves reliance, then costs of reliance are themselves costs of believing. But I do not see how it follows that the costs of reliance are more “intrinsic to the nature of belief” than the (direct) costs of *belief itself* are. After all, it is part of the nature of belief – albeit a vacuous and trivial part – that belief involves *belief*, and yet Schroeder doesn't want it to follow from *this* that (all) costs of belief are “intrinsic to the nature of belief” in his sense. It is true that, if it's in the nature of belief that belief involves reliance, then it is part and parcel of the *sort of thing belief is* that beliefs are capable of incurring pragmatic costs and benefits connected with relying on them. But it is also part and parcel of the sort of thing belief is that beliefs are capable of incurring pragmatic costs and benefits that arise merely by *having* the belief: beliefs are the sort of things that *can* make us feel comforted, threatened, insecure, hopeful, and so on – and that can, in far-out but possible cases, carry more tangible (e.g. financial) costs and rewards. Of course, that any one particular belief has any one of these *particular* costs and benefits is a contingent matter that is not intrinsic to the nature of belief; but the same goes for the particular costs and benefits of *relying* on particular beliefs.

Moreover, putting great stress on the difference between direct costs of belief and costs of reliance can produce unintuitive asymmetries between cases. Recall that I earlier distinguished:

Pascal's Wager. If God exists, then you will go to heaven iff you believe God exists, and to hell otherwise. (If God doesn't exist, then there is no afterlife.)

Pascal's Wager (Modified). If God exists, then you will go to heaven iff you do not believe that God doesn't exist, and to hell otherwise. (If God doesn't exist, then there is no afterlife.)

The present view yields the verdict that pragmatic considerations do not make a difference in either of these cases, since the costs here are of *belief itself*, and not of relying on one's belief. But now consider a further case:

Pascal's Wager (Modified-2). If God exists, you will go to heaven iff you pray nightly, and to hell otherwise. (If God doesn't exist, then there is no afterlife.)

⁵⁷ For more on this sort of idea, see e.g. Nolfi (2019).

⁵⁸ Cf. also Bolinger (forthcoming: 6).

I assume that (at least given many plausible ways of developing the case), conditional on God's non-existence, it is rational not to pray nightly.⁵⁹ Thus, in Pascal's Wager (Modified-2), going to hell is a potential cost of *relying* on the proposition that God does not exist. Thus, the present view says that the (potential) cost in Pascal's Wager (Modified-2) *does* make a difference to the epistemic standards for believing that God doesn't exist, potentially making it the case that I should not believe that God doesn't exist in this case.

This compares oddly with the present view's verdicts in the original Pascal's Wager case and, especially, in Pascal's Wager (Modified) – namely that the relevant pragmatic considerations do *not* make a difference in these cases. Intuitively, the considerations against believing that God doesn't exist in Pascal's Wager (Modified-2) are just as “Pascalian” as the considerations against believing that God doesn't exist in Pascal's Wager (Modified). If the fact that my believing that God doesn't exist might send me to hell *via licensing reasoning that would lead to certain actions* counsels in favor of erring on the side of caution and setting my evidential standards for believing this proposition high (in Pascal's Wager (Modified-2)), why shouldn't the fact that my believing that God doesn't exist might *directly* send me to hell counsel in favor of just the same thing (in Pascal's Wager (Modified))? The present view is committed to an implausible asymmetry here.

5. Conclusion

In this paper, I've reviewed different ways for a pragmatist to keep her pragmatism “moderate”, and argued that none of them look promising. Some will take that as a reason to conclude that pragmatism must be rejected outright. Others will see it as an argument for hard pragmatism, which simply embraces the thought that pragmatic considerations make a difference to what one ought to believe in all the cases we've discussed. I leave this choice open: the arguments of this paper, as I think of them, provide support for the disjunction of these two more extreme positions. Put another way: they provide reason to think that we will have to *either* revise the moderate pragmatist's verdict about cases like Bribe and Pascal's Wager – namely that pragmatic considerations do not affect what we ought to believe in these cases – *or* to revise the moderate pragmatist's verdict about cases like Risks – namely that pragmatic considerations *do* affect what we ought to believe in *these* cases.

Moreover, as I noted earlier, there is considerable pressure on pragmatic encroachers to be moderate pragmatists. The idea that pragmatic considerations make a difference to what you're *epistemically* justified in believing even in cases like Bribe is extreme. To the degree that it is unacceptable, the argument of this paper suggests that the pragmatic encroachment thesis must be rejected – even by those willing to accept hard pragmatism.

References

Anderson, C. & Hawthorne, J. (2019). “Knowledge, Practical Adequacy, and Stakes,” *Oxford Studies in Epistemology*, 6: 234-257.

⁵⁹ For reasons analogous to those given for the Risks case toward the end of §2.

- Basu, R. (2019a). "Radical Moral Encroachment: The Moral Stakes of Racist Beliefs," *Philosophical Issues*, 29: 9-23.
- (2019b). "The Wrongs of Racist Beliefs," *Philosophical Studies*, 176/9: 2497-2515.
- Basu, R. & Schroeder, M. (2019). "Doxastic Wronging," in Kim & McGrath (eds.), *Pragmatic Encroachment in Epistemology*. London: Routledge.
- Benton, M. (2018). "Pragmatic Encroachment and Theistic Knowledge," in Benton, Hawthorne & Rabinowitz (eds.), *Knowledge, Belief, and God: New Insights in Religious Epistemology*. Oxford: Oxford University Press.
- Berker, S. (2018). "A Combinatorial Argument against Practical Reasons for Belief," *Analytic Philosophy*, 59/4: 427-470.
- Bolinger, R. (forthcoming). "The Rational Impermissibility of Accepting (Some) Racial Generalizations," *Synthese*, DOI: 10.1007/s11229-018-1809-5. Page references are to the Online First version.
- Clarke, R. (2013). "Belief is Credence One (In Context)," *Philosophers' Imprint*, 13/11.
- Crewe, B. & Ichikawa, J.J. (forthcoming). "Rape Culture and Epistemology," in Lackey (ed.), *Applied Epistemology*. Oxford: Oxford University Press.
- Enoch, D., Spectre, L. & Fisher, T. (2012). "Statistical Evidence, Sensitivity, and the Legal Value of Knowledge," 40/3: 197-224.
- Fantl, J. & McGrath, M. (2002). "Evidence, Pragmatics, and Justification," *Philosophical Review*, 111/1: 67-94.
- (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Fritz, J. (2017). "Pragmatic Encroachment and Moral Encroachment," *Pacific Philosophical Quarterly*, 98/S1: 643-661.
- (forthcoming). "Moral Encroachment and Reasons of the Wrong Kind," *Philosophical Studies*.
- Ganson, D. (2008). "Evidentialism and Pragmatic Constraints on Outright Belief," *Philosophical Studies*, 139/3: 441-458.
- Gardiner, G. (2018). "Evidentialism and Moral Encroachment," in McCain (ed.), *Believing in Accordance with the Evidence*. Dordrecht: Springer.
- Gendler, T.S. (2011). "On the Epistemic Costs of Implicit Bias," *Philosophical Studies*, 156/1: 33-63.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. & Stanley, J. (2008). "Knowledge and Action," *Journal of Philosophy*, 105/10: 571-590.
- Howard, C. (2016). "In Defense of the Wrong Kind of Reason," *Thought*, 5/1: 53-62.
- Ichikawa, J.J. (2012). "Knowledge Norms and Acting Well," *Thought*, 1/1: 49-55.
- Jackson, E. (2019). "How Belief-Credence Dualism Explains Away Pragmatic Encroachment," *Philosophical Quarterly*, 69/276: 511-533.
- James, W. (1896/2000). "The Will to Believe," in his *Pragmatism and other Writings* (ed. Gunn). New York: Penguin.
- Morton, J. & Paul, S. (2018). "Grit," *Ethics*, 129/2: 175-203.

- Moss, S. (2018). "Moral Encroachment," *Proceedings of the Aristotelian Society*, 118/2: 177-205.
- Munton, J. (2019). "Beyond Accuracy: Epistemic Flaws with Statistical Generalizations," *Philosophical Issues*, 29: 228-240.
- Nolfi, K. (2019). "Another Kind of Pragmatic Encroachment," in Kim & McGrath (eds.), *Pragmatic Encroachment in Epistemology*. London: Routledge.
- Owens, D. (2000). *Reason without Freedom: The Problem of Epistemic Normativity*. London: Routledge.
- Pace, M. (2011). "The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe'," *Noûs*, 45/2: 239-268.
- Reisner, A. (2008). "Weighing Pragmatic and Evidential Reasons for Belief," *Philosophical Studies*, 138/1: 17-27.
- (2009). "The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem," *Philosophical Studies*, 145: 257-272.
- Riggs, W.D. (2003). "Balancing Our Epistemic Goals," *Noûs*, 37/2: 342-352.
- Rinard, S. (2017). "No Exception for Belief," *Philosophy and Phenomenological Research*, 94/1: 121-143.
- (2018). "Believing for Practical Reasons," *Noûs*, 53/4: 763-784.
- Ross, J. & Schroeder, M. (2014). "Belief, Credence, and Pragmatic Encroachment," *Philosophy & Phenomenological Research*, 88/2: 259-288.
- Schroeder, M. (2012a). "Stakes, Withholding, and Pragmatic Encroachment," *Philosophical Studies*, 160/2: 265-285.
- (2012b). "The Ubiquity of State-Given Reasons," *Ethics*, 122/3: 457-488.
- (2018). "When Beliefs Wrong," *Philosophical Topics*, 46/1: 115-127.
- Staffel, J. (2019). "How Do Beliefs Simplify Reasoning?," *Noûs*, 53/4: 937-962.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Thomson, J.J. (1986). "Liability and Individualized Evidence," *Law and Contemporary Problems*, 49: 199-219.
- Weatherson, B. (2005). 'Can We Do Without Pragmatic Encroachment?,' *Philosophical Perspectives*, 19: 417-443.
- (2011). "Defending Interest-Relative Invariantism," *Logos & Episteme*, 2/4: 591-609.
- (2012). "Knowledge, Bets, and Interests," in Brown & Gerken (eds.), *Knowledge Ascriptions*. Oxford: Oxford University Press.
- (2016). "Games, Beliefs and Credences," *Philosophy and Phenomenological Research*, 92/2: 209-236.
- Worsnip, A. (2015). "Two Kinds of Stakes," *Pacific Philosophical Quarterly*, 96/3: 307-324.
- (2016). "Belief, Credence, and the Preface Paradox," *Australasian Journal of Philosophy*, 94/3: 549-562.