

PRACTICAL CERTAINTY

Dustin Locke
Claremont McKenna College

Penultimate draft. To be published in *Philosophy and Phenomenological Research*.

Please cite published version.

Abstract. When we engage in practical deliberation, we sometimes engage in careful *probabilistic* reasoning. At other times, we simply make *flat out* assumptions about how the world is or will be. A question thus arises: when, if ever, is it rationally permissible to engage in the latter, less sophisticated kind of practical deliberation? Recently, a number of authors have argued that the answer concerns whether one *knows* that *p*. Others have argued that the answer concerns whether one is *justified in believing that one knows* that *p*. Against both of these, this paper argues that the answer concerns whether *p* is ‘practically certain’—that is, whether the actual epistemic probability that *p* differs from epistemic certainty that *p* only in ways that are irrelevant to the decision one currently faces.

Keywords. practical rationality, practical reasoning, knowledge and action, pragmatic encroachment, motivating reasons, knowledge norm, justification norm

1. INTRODUCTION

This paper is concerned with something I call ‘premissing mistakes’.¹ It will be helpful to begin by contrasting premissing mistakes with two more familiar kinds of mistakes. Here are some illustrative examples.

Mike is a compulsive gambler. Today he has studied all of the evidence available to him concerning the 6th race, and that evidence strongly supports the proposition that Zenyatta will *not* win. Moreover, given the high epistemic probability that Z will not win, and given the odds at which a win bet on Z is paying (and given anything else you think is relevant here), Mike has most reason to *not* bet on Z. Now Mike understands all of this, and he does not believe that Z will win, nor does he form the intention to bet on Z. However, unable to resist the urge to bet on Z, Mike soon finds himself at the pari-mutuel window. In betting on Z, Mike is not completely rational: although he does not believe that Z will win—as he ought not—he performs at least one *action* that he ought not to perform. Call this a *practical mistake*—the mistake of performing an action one ought not to perform.²

John is a compulsive *non*-gambler. Today John has studied all of the evidence available to him concerning the 6th race, and that evidence strongly supports the proposition that Z will not win. Given the probability and the payout, John has most reason to *not* bet on Z. However, John (irrationally) believes that Z will win and has made up his mind to bet on her. Nonetheless, when John gets to the pari-mutuel window, he just can’t go through with it. Is John then completely rational? Of course not: although he does not bet on Z—as he ought not—he has one belief that he ought not to have—namely, the belief that Z will win. Call this a *doxastic mistake*—the mistake of believing something one ought not to believe.

¹ I owe many people many thanks for their help in the construction of this paper. Most notable among these is Amy Kind, my faculty mentor at CMC. Both intellectually and professionally, Prof. Kind has played a tremendous role in my transition from graduate student to faculty member, and I can’t say enough to thank her for this. Others to whom I owe special thanks include Yuval Avnur, Alex Bundy, Vanessa Carbonell, Andy Cullison, Kenny Easwaran, Paul Hurley, Jonathan Jenkins-Ichikawa, Clayton Littlejohn, Heather Lowe, Neil Mehta, Eliot Michaelson, David Plunkett, Alex Rajczi, Drew Schroeder, Dion Scott-Kakures, Declan Smithies, and anonymous referee for this journal, and all the participants in the Claremont Colleges Works in Progress Group (WIP), the Circle of Los Angeles Philosophers (CLAP), and a session at the 2012 Pacific APA.

² Some philosophers use the phrase ‘practical mistake’ in a narrower sense than I am using it here. On my usage, one commits a practical mistake if and only if one performs an action that one ought not to have performed, regardless of whether the reasons one ought not to have performed it are prudential, moral, or some other sort.

Suzy is neither a compulsive gambler nor a compulsive non-gambler. Like John and Mike, Suzy has studied all of the evidence available to her concerning the 6th race. In Suzy's case, that evidence strongly supports the proposition that *Z will win*. Accordingly, Suzy believes that *Z will win* and, deliberating on the basis of that belief, forms the intention to bet on *Z*. More specifically, Suzy reasons as follows.

Z will win.

If *Z will win*, then betting on *Z* is the 'optimal' thing to do.

So I'll bet on *Z*.³

Let us say that in reasoning like this, Suzy *premises* that *Z will win*. I'll have quite a bit to say about the notion of *premissing* below, but for now I trust the idea is familiar enough.

In a moment, we'll consider whether and under what conditions it is appropriate for Suzy to engage in this kind of reasoning. Before proceeding, however, let me be clear on how I intend the term 'optimal' to be understood. First, 'optimal' should not be understood as (analytically) equivalent to 'the thing to do'. Rather, 'optimal' should be understood as a mere placeholder for whatever purely descriptive features Suzy ought to be thinking about in her practical deliberation.⁴ Just to fix ideas, suppose we thought that Suzy ought to be thinking as a hedonistic egoist. In that case, the second premise of her reasoning would simply say that if *Z will win*, then betting on *Z* will bring Suzy a greater balance of 'hedons' than not betting on *Z*. So that we might focus our attention

³ The conclusion 'So I'll bet on *Z*' should be understood to express Suzy's *intention* to bet on *Z*, not a (mere) *belief* about what Suzy will in fact do. See John Broome and Christian Piller (2001) for a useful discussion of the structure of practical reasoning.

⁴ The point of using 'optimal' in this way is to make what I say here compatible with a wide-range of theories about which sorts of facts provide reasons to act in which ways. For the hedonistic egoist, the only relevant facts are facts about pain and pleasure for the agent. Other (more plausible) theories suggest otherwise. In defense of the principle I call 'KR' below, Jonathan Jenkins Ichikawa (2012) has cautioned us against moving too quickly from claims about what it is rational to *do* to claims about what it is rational to *premise* (or, as he says, 'treat as a reason'). Such moves, notes Ichikawa, rely on implicit claims to the effect that such and such is a (decisive) reason to act in thus and so way. Ichikawa seems right about at least this much—whether he is right that KR can successfully be defended on these grounds is another matter—and thus I here use 'optimal' in a way that allows the reader to adjust Suzy's reasoning so that she is responding to whatever sorts of facts the reader deems relevant. Thanks to an anonymous referee for pressing me on this point.

on Suzy's epistemic position with respect to the first premise in her reasoning, let us simply stipulate that the case is such that Suzy is justified in being *as certain as you like* that the second premise—the conditional premise—is true.⁵

Now let us suppose that despite the unlimited certainty of the second premise, and despite Suzy's justified high degree of confidence in the first premise, betting on Z is not the thing to do. Specifically, suppose that although Suzy is justified in being as confident as you like that if Z will win, then betting on Z is the optimal thing to do, and although she is justified in being *highly* confident that Z will win, she is not justified in being so confident that Z will win that it makes sense to bet on Z—that is, that betting on Z is the thing to do. Suppose, for example, that Suzy is required to stake ten dollars for the possibility of winning three cents, but that she is justified in having no more than a .97 degree of confidence that Z will win. In that case, Suzy makes a practical mistake when she bets on Z.

But is the practical mistake the only mistake that Suzy makes? No, prior to making the practical mistake, Suzy has *already* made a mistake—indeed, a mistake that

⁵ Importantly, this stipulation does not settle whether Suzy is justified in believing the following, importantly different conditional.

(TTD) If Z will win, then betting on Z is the thing to do.

(TTD) may very well be false, even if the conditional premise in Suzy's reasoning is true. For the (sophisticated) hedonistic egoist, the thing to do is not simply what will *in fact* bring the greatest balance of hedons, but rather what has the highest *expected hedonic value*, where the latter is to be calculated in accordance with one's favorite decision theory. My favorite decision theory is causal decision theory, and so (simplifying things a bit) I would calculate expected hedonic value as follows.

$$EV(A) = \sum_O P(A|O)V(O\&A)$$

where $EV(A)$ is the expected hedonic value of action A,

$P(A|O)$ is the degree of confidence it is rational for the agent to have in the subjunctive (or causal) conditional $A|O$, and

$V(O\&A)$ is, for the hedonistic egoist, the balance of hedons for the agent in the event that $O\&A$.

On this account, the thing to do depends, not on the fact of whether Z will win, but (in part) on Suzy's justified degree of confidence that Z will win. Hence, (TTD) may be false, even if the conditional premise in Suzy's reasoning is true.

lead to her making the practical mistake. The prior mistake was that of simply *premissing* that Z would win—more specifically, the mistake of premissing that Z would win when deciding whether to take this specific bet. Call this a *premissing mistake*—the mistake of premissing something one ought not to premiss. This paper is about premissing mistakes.

Are premissing mistakes really distinct from doxastic mistakes? They are. Suzy does not make a doxastic mistake, at least not with respect to the proposition that Z will win. As we said, the evidence strongly supports the proposition that Z will win, and so Suzy makes no mistake in *believing* that Z will win.⁶ Her mistake is allowing this belief to figure into her practical deliberation in the way that it does. Are premissing mistakes really distinct from practical mistakes? They are. In Suzy’s case, the premissing mistake leads to, but is distinct from, the practical mistake. Suppose that Suzy hadn’t had the chance to commit the practical mistake, or even to form the *intention* to commit the action that would be the practical mistake. Still, as soon as she allows the belief that Z will win to figure in her deliberation over whether to take the bet at the designated odds, she makes a mistake—for purposes of deciding whether to take that bet, she ought not take it for granted that Z will win.

Note that it’s not always a mistake to premiss. Suppose Lucy looks outside and sees that it’s raining. She then reasons as follows.

It’s raining.

If it’s raining, then taking an umbrella is the optimal thing to do.

So I’ll take an umbrella.

If we fill in the details of Lucy’s situation in a normal way, there’s nothing wrong with her simply premissing that it’s raining.

The central question of this paper is this: when is it appropriate—that is, rationally permissible—to premiss that *p* in one’s practical deliberation? Call this the ‘premissing question’. In terms of the examples above, why is it appropriate for Lucy to

⁶ Philosophers fond of the idea of ‘pragmatic encroachment’ might balk at this suggestion, thinking that the rationality of *believing* in part depends upon the rationality of *letting that belief figure in one’s practical deliberation*. I consider this objection in section five.

premise that it is raining (in the way that she does), while it is not appropriate for Suzy to premise that Z will win (in the way that she does)?

A number of philosophers have recently argued that one ought not premise that p in one's practical deliberation unless one *knows* that p .⁷ This suggestion seems to make good sense of the above cases: it is appropriate for Lucy to premise that it is raining because she *knows* that it is raining; it is not appropriate for Suzy to premise that Z will win because she does *not know* that Z will win. This view is also seemingly supported by the prevalence of knowledge-based criticisms in ordinary discourse. When someone premises something they ought not to have premised, we often criticize them by saying something of the form 'But you didn't *know* that...'.

In contrast to the knowledge-based view, more traditional philosophers have argued that as long as one's belief that p is sufficiently *justified*, it is appropriate to premise that p in one's practical deliberation.⁸ In what follows I argue in favor of a sophisticated version of the justification-based view. Roughly put, my thesis is that it is appropriate to premise that p in one's practical deliberation if and only if p is *epistemically certain enough*, where what counts as 'enough' is determined by the particular nature of the decision problem one faces.

The next two sections are concerned with laying some important groundwork for the rest of the paper. In section two I distinguish several types of premising and isolate the particular kind that we are here interested in. In section three I position the premising question in relation to more traditional questions for the theory of practical reasoning. Unfortunately, epistemologists have so far not done a very good job positioning what I am calling the 'premising question' in relation to the large body of work by ethicists on the nature of practical reasoning. As we will see, answers to the premising question are best seen as supplementary to, rather than competitors of, the primary views that have been carved out in the ethics literature.

The remaining sections of the paper work towards answering the premising question. Section four considers the view offered by the most often-cited champions of the knowledge-based approach to the premising question, John Hawthorne and Jason Stanley (2008), and contrasts this with a simple justification-based approach. I argue

⁷ John Hawthorne (2004), Jason Stanley (2005), and Hawthorne and Stanley (2008).

⁸ See Mark Kaplan (1985) and (2003) for a defense of the general idea that when it comes to rationality, it's what you *justifiably believe*, and not what you know, that counts.

that both views face problems. Sections five and six develop a novel view. According to the view defended here, the appropriateness of premising that p depends on (1) the epistemic probability of p and (2) the nature of the particular decision problem one faces. Finally, section seven offers an account of why ordinary discourse frequently contains knowledge-based criticisms despite the fact that knowledge is not the norm of practical deliberation. Collectively, the various sections of this paper make a robust case for a sophisticated justification-based answer to the premising question.

2. PREMISING

We should begin by saying something more about the particular phenomenon we are here interested in. First, it is important to distinguish reasoning aimed at settling what *to do* (practical deliberation) from reasoning aimed at settling what *to believe* (theoretical deliberation). Some philosophers have maintained that the norm of premising—that is, the standard of when it is appropriate to premise—is the same with respect to both kinds of deliberation.⁹ While I cannot do this issue justice here, there is compelling reason to think otherwise. Consider again our case from the introduction. We already noted that it is not appropriate for Suzy to premise that Z will win in the particular piece of practical deliberation in which she so premises. But suppose that, rather than deliberating over whether to make a bet, Suzy were engaged in a piece of theoretical deliberation over whether Z will be running in this year’s Breeder’s Cup Classic. Here Suzy is not trying to decide what to *do*, but simply what to *believe*. Suppose that upon learning that today’s 6th race is a Breeder’s Cup Classic “Win and You’re In” race, Suzy reasons as follows.

Z will win.

If Z wins, then she will run in this year’s Breeder’s Cup Classic.

So Z will run in this year’s Breeder’s Cup Classic.

If the epistemic probability of Z winning is high enough, there doesn’t seem to be anything inappropriate about Suzy reasoning in this way. And yet the payout odds on today’s race might be so low that even this high epistemic probability that Z will win does not make it appropriate for Suzy to premise that Z will win when deliberating

⁹ Hawthorne and Stanley (2008: p. 577).

whether to bet. In any event, in this paper we will be concerned with premising in practical deliberation, and one should not assume that what we have to say here applies equally and to all cases of theoretical deliberation.¹⁰

In addition to distinguishing premising in practical deliberation from premising in theoretical deliberation, we also need to distinguish premising that involves *belief* from premising that involves other kinds of mental states. Here we are interested exclusively in the former. An example of a kind of premising we are *not* interested in can be seen in a case from Ram Neta (2009: p. 685). Suppose that I set my mind to *either* going to the movies tonight or staying home. I might then reason as follows

1. Either I will go to the movies tonight or I will stay home.
2. If I go to the movies, then...
3. If I stay home, then...
4. So I will stay home.

In a case like this, there is a clear sense in which I premise that *either I will go to the movies or I will stay home*. However, it is not a *belief* that either I will go to the movies or I will stay home that figures in my deliberation. It is, rather, the *intention* or *plan* to either go to the movies or stay home that figures in the deliberation. Spelling out the distinction in any sort of detail will no doubt be difficult, but I trust the distinction is familiar enough for our purposes: what is going on in Neta's case is something different from what is going on in the case of Suzy's deliberation over whether to bet on Z. In this paper, we will be interested in cases like Suzy's—cases where a *belief* that *p* figures

¹⁰ There is at least one type of deliberation that is arguably both practical and theoretical—namely, normative deliberation, e.g., deliberation over what one *ought* to do. On the one hand, successful deliberation over what one all-things-considered ought to do arguably settles what to do (Allan Gibbard 2003), and so seems to be practical deliberation. On the other hand, successful deliberation over what one ought to do arguably settles what to *believe* about what one ought to do, and so seems to be theoretical deliberation. If normative deliberation is genuinely both practical and theoretical, then in this case at least, the norm of premising in practical deliberation must not conflict with the norm of premising in theoretical deliberation. All of this is compatible with the above argument, which merely aims to put pressure on the thesis that the norms that govern these two types of reasoning (in general) are one and the same. Thanks to an anonymous referee for pressing this point.

in one's practical deliberation.¹¹ Henceforth, this is what I shall mean by 'premissing that *p*'.

3. GOOD REASONS AND PREMISING MISTAKES

When one acts on the basis of a deliberation in which one premises that *p*, one will have acted (in part) for the reason that *p*. But there is considerable controversy over what acting for the reason that *p* amounts to. On some views, this amounts to treating the *proposition* that *p* as a reason for acting.¹² On other views, this amounts to treating the *fact* that *p* (if it is a fact) as a reason.¹³ On still other views, this amounts to treating the *state of affairs* that *p* (which may or may not obtain) as a reason.¹⁴ We should not try to settle this dispute here. My hope is that the notion of 'premissing that *p*' is neutral enough that proponents of any of these views will find it acceptable.

While we can avoid taking a stance on what it is to act for a reason, we cannot avoid the obligation to explain how the premissing question relates to more traditional questions being asked in the literature on practical reasons. Recent work by epistemologists on the premissing question has drawn both comment and concern from ethicists interested to know just how this new work in epistemology connects to more traditional work on practical reasons.¹⁵ It is a shame, I think, that we epistemologists have not done a better job of explaining that connection.

In the ethics literature, philosophers commonly distinguish between normative reasons and motivating reasons. Here is Michael Smith.

[It] seems to me that we operate with quite different concepts of a reason for action depending on whether we emphasize the explanatory dimension and downplay the justificatory, or *vice versa*. The claim 'A has a reason to ϕ ' is thus ambiguous. It may be a claim about a *motivating* reason A has, when we

¹¹ Because we are interested solely in cases where a *belief* that *p* figures into one's deliberation, we are also not interested in cases where one merely assumes that *p* for the sake of reductio, or other cases where one assumes that *p* merely for the sake of seeing what follows.

¹² T. M. Scanlon (1999: p. 57).

¹³ John Hyman (2011).

¹⁴ Jonathan Dancy (2000).

¹⁵ Commenting on Hawthorne and Stanley (2008), Jussi Suikkanen writes: "Usually, when I read epistemologists writing about [practical] reasons I feel like entering a strange new world where things look slightly peculiar." <http://peasoup.typepad.com/peasoup/2007/07/hawthorne-and-s.html>

emphasize the explanatory dimension and downplay the justificatory, or a claim about a *normative* reason A has, when we emphasize the justificatory dimension and downplay the explanatory... To say that someone has a normative reason to ϕ is to say that there is some normative requirement that she ϕ 's, and is thus to say that her ϕ -ing is justified from the perspective of the normative system that generates that requirement... Motivating reasons are, however, different. The distinctive feature of a motivating reason to ϕ is that, in virtue of having such a reason, an agent is in a state that is *explanatory* of her ϕ -ing, at least other things being equal... (1994: pp. 95 – 96)

While Smith goes on to make certain claims about the *ontology* of both sorts of reasons, we need not endorse any particular ontology here.¹⁶ All we need is the general and generally accepted distinction between normative reasons, on the one hand, and motivating reasons, on the other. Once we have drawn this distinction, we might wonder whether motivating reasons have any place in a theory of practical rationality. It seems to many that they do, and this for the simple fact that it is possible to do what one has most normative reason to do, but to do so for *bad* motivating reasons. In such cases, one will be in some sense *rationaly criticisable*. A complete theory of practical rationality thus ought to account tell us something about what it is to act for *good* motivating reasons.¹⁷

On the standard view, to act for good motivating reasons is to act for motivating reasons that are normative reasons for one to so act.¹⁸ But if this is all that there is to

¹⁶ There is considerable controversy over just what it is to have a normative reason to ϕ . Relatedly (but not quite equivalently) there is considerable controversy over just what normative reasons *are*. For some philosophers, normative reasons are psychological states—e.g., beliefs, desires, or combinations thereof. For other philosophers, normative reasons are facts or states of affairs typically external to the agent's own psychology—according to some such philosophers, the fact that the bus is heading your way is a normative reason to move. Similarly, there is considerable controversy over just what motivating reasons are. For some, motivating reasons are psychological states—e.g., beliefs, desires, or combinations thereof. For others, motivating reasons are facts or states of affairs typically external to the agents own psychology—according to some such philosophers, the fact that the bus was heading your way might have been the motivating reason that you moved.

¹⁷ The distinction I am making here is not simply the distinction between doing (a) what one has normative reason to do and (b) what is the optimal thing to do, given the actual facts about how the world is. In addition to that distinction, there is also the distinction between (a) doing what one has normative reason to do and (c) doing what one does for good motivating reasons.

¹⁸ See Dancy (2000: Ch. 5).

acting for good motivating reasons, then it seems that Suzy acts for a good reason: Suzy acts for the reason that Z will win the race, and that Z will win the race is (part of) a normative reason for Suzy to make the bet.¹⁹ Still, it seems that Suzy proceeds in a way that is in at least one respect rationally criticisable: when she ought to be reasoning in a way that is sensitive to the more or less precise *probability* that Z will win, instead she simply premises that Z will win. As those who think of reasons as *facts* might put the point, the fact that Z will win is a normative reason for Suzy to make the bet, but Suzy's epistemic position with respect to that fact is too *weak* to make it appropriate for her to treat that fact as a reason.

Cases like Suzy's suggest that either (1) there is more to acting for good motivating reasons than acting for motivating reasons that are normative reasons, or else (2) even when agents act for good motivating reasons (*and* perform the act that they have most normative reason to perform), there is another way that they may fail to satisfy the demands of practical rationality—namely, they may premise propositions that they ought not premise. While it is an interesting question which of these lessons is the correct one to draw, I here remain agnostic. Either way, we need to address the question of when is it appropriate to premise that *p* in practical deliberation. This is the premising question.

4. KNOWLEDGE VS. JUSTIFICATION

Under what conditions is it appropriate to premise that *p*? At a minimum, *p* must be *relevant* to the agent's decision. Following others, let us understand 'relevance' as *p*-dependence, where a choice between options $x_1...x_n$ is *p*-dependent if and only if the most preferable of $x_1...x_n$ conditional on the proposition that *p* is not the same as the most preferable of $x_1...x_n$ conditional on not-*p*.²⁰

Besides the relevance of *p* to one's decision, what else must be the case for it to be appropriate to premise that *p*? Recently, certain philosophers have suggested that one ought not to premise that *p* unless one knows that *p*. The champions of such a view are John Hawthorne and Jason Stanley (2008), who offer

¹⁹ If you're bothered by the fact that Z's winning the race is in the *future*, feel free to change the case accordingly.

²⁰ Hawthorne and Stanley (2008).

The Knowledge-Reasons Principle. When p is relevant to S's decision, it is appropriate for S to treat p as a reason if and only if S knows that p .²¹

As I noted above, it is controversial whether agents treat *propositions* as reasons. Let us thus reformulate KR in terms of the more neutral notion of *premising*.

The Knowledge-Reasons Principle. When p is relevant to S's decision, it is appropriate for S to premise that p in her practical deliberation if and only if S knows that p .

While KR handles certain cases quite well—e.g., Suzy's case from the introduction—it is easy to get oneself in a position where KR starts looking implausible. In (2009), Jeremy Fantl and Matthew McGrath run what they call a 'subtraction argument' on a similar principle—more precisely, the right-to-left direction of a similar principle. I will run an analogous argument here, without pausing to detail the perhaps important ways in which their argument differs from my own. The subtraction argument proceeds in two steps. The first step centers on a case of justified true belief that does not amount to knowledge—a Gettier case—and the second step centers on a case of justified *false* belief.

Begin with the Gettier case. John has been in a coma for six months and awakes to see the digital clock next to his hospital bed reading '5:00 a.m.'. As it happens, it is 5:00 a.m., but the clock has been broken for several weeks and so it is just by chance that what the clock says is correct. John has no reason to think that the clock is broken

²¹ It is important not to confuse KR with similar-sounding principles. For example, Peter Unger (1975), John Hyman (1999), and Jennifer Hornsby (2007) argue that S treats (the fact that) p as a reason only if S knows that p . Importantly, Unger, Hyman, and Hornsby's principle is not an answer to the premising question—their principle merely addresses the conditions under which it is *possible* to treat (the fact that) p as a reason, while leaving aside the question of when it is *rationally permissible* to premise that p . Jeremy Fantl and Matthew McGrath (2009) also defend a principle that is similar to, but distinct from KR. According to Fantl and McGrath, if S knows that p , then p is warranted enough to justify S in ϕ -ing, for any ϕ . Fantl and McGrath also defend the stronger principle that if S has knowledge-level justification that p , then p is warranted enough to justify S in ϕ -ing, for any ϕ . As Fantl and McGrath make clear, these are principles about the conditions under which p is a normative reason to ϕ . As such, these principles, like Hyman, Hornsby and Unger's principle, do not address the premising question. I will therefore simply set them aside. I do, however, address the Unger/Hyman/Hornsby thesis elsewhere (Locke Manuscript A).

and so he forms the belief that it is 5:00 a.m.. Acting on this belief, he decides not to press the button for a nurse, because he figures that a nurse will be coming by soon anyway. Here it seems rationally permissible for John to premise that it is 5:00 a.m. Yet, since John is Gettiered, he does not *know* that it is 5:00 a.m.. This suggests that we ought to “subtract” from our principle whatever it is that separates knowledge from justified true belief.

The Justified-and-True Principle. When p is relevant to S’s decision, it is appropriate for S to premise that p if and only if S has a justified true belief that p .²²

The second stage of the argument is designed to motivate the subtraction of *truth*. Let’s change our case a bit. Suppose that, as it happens, it’s not 5:00 a.m., but midnight. Still, John remains justified in believing that it is 5:00 a.m., and it seems that it is rationally permissible for him to premise that it is 5:00 a.m. in his practical deliberation. This suggests that we should “subtract” truth from our principle as well.

Justified-Belief Principle. When P is relevant to S’s decision, it is appropriate for S to premise that p if and only if S is justified in believing that p .

It is considerations like these, I think, that motivate proponents of a justification-based approach to the premising question: whatever there is to knowledge that goes beyond justified belief is simply irrelevant to whether it is rationally permissible to act on one’s belief.²³

²² Clayton Littlejohn (2009).

²³ Not everyone finds the subtraction argument convincing, especially not proponents of KR. There is something of a ‘standard move’ in response to the argument. The standard move involves an appeal to excuses: while John does not know it is 5:00 a.m. (either because this is not true or because he is Gettiered), he has a good excuse for premising that it is 5:00 a.m.—and so is not *blameworthy* for doing so. His good excuse, allegedly, is that he *reasonably believes* that he knows that it is 5:00 a.m.. Hawthorne and Stanley (2005) make this move in defense of KR specifically, while Keith DeRose (2002) and Matt Weiner (2005) make this sort of move in defense of the knowledge account of assertion (which faces an analogous subtraction argument). I do not have space to do justice to this issue in this essay, but here are some considerations to keep in mind. First, as we will see below, there appear to be counterexamples to KR—e.g., Brown’s *Surgeon* case—for which the appeal to excuses will not help. Second, there is reason to be suspicious of this sort of appeal to excuses. While the idea of something’s being *wrong but excused by non-culpable ignorance* is a familiar one, an appeal to excuses here requires the idea of something’s being *not rational but excused by non-culpable ignorance*—the idea being that John

In light of considerations like the above, why would anyone be tempted by KR? Proponents of KR make their case mainly in two parts. The first part involves apparent counterexamples to JB, such as the case of Suzy from the introduction. Suzy is justified in believing that Z will win, but it is not appropriate for Suzy to premise that Z will win in deliberating whether to make the bet. It's not implausible to think that this is because Suzy does not *know* that Z will win: since Suzy does not know that Z will win, she ought to be reasoning in a way that is sensitive to the more or less precise probability that Z will win.

The second part of the case for KR over JB involves appeal to a wide range of cases in which we find it natural to criticize someone for acting on a belief that did not amount to knowledge. Here are Hawthorne and Stanley.

[O]ur ordinary folk appraisals of the behavior of others suggest that the concept of knowledge is intimately intertwined with the rationality of action. Suppose, for example, that Hannah and Sarah are trying to find a restaurant, at which they have time-limited reservations. Instead of asking someone for directions, Hannah goes on her hunch that the restaurant is down a street on the left. After walking for some amount of time, it becomes quite clear that they went down the wrong street. A natural way for Sarah to point out that Hannah made the wrong decision is to say, "You shouldn't have gone down this street, since you did not know that the restaurant was here." (Hawthorne and Stanley 2008: p. 571)

Obviously a *hunch* falls short of even justified belief, and so one might suggest that Hannah's mistake is that of acting on a belief that was not even justified. Be that as it

is non-culpably ignorant of the fact that he does not know that it is 5:00 a.m.. The notion of *not rational but excused by non-culpable ignorance*, however, isn't clearly coherent. Why, after all, does non-culpable ignorance excuse? A plausible answer is that non-culpable ignorance excuses because non-culpable ignorance can make it *rational* to do something that is (in some other sense) *wrong*. If this answer is correct, then to say that someone has done something *not rational but excused by non-culpable ignorance* is to say that he has done something that is *not rational but also rational due to non-culpable ignorance*. That looks incoherent. As I said, I cannot here give the appeal to excuses the attention it deserves. A more complete answer would have to take into account the large and growing body of literature on the nature of excuses. I do hope, however, that these considerations at least put some pressure on proponents of the move—pressure, that is, to explain how it is that ignorance can excuse irrationality, and to do so without falling into contradiction. In any event, as I said above and as I will discuss below, there are apparent counterexamples to KR for which an appeal to excuses is simply irrelevant. I am thankful to an anonymous referee for encouraging me to discuss the appeal to excuses.

may, Hawthorne and Stanley seem right in so far as they are merely claiming that in ordinary discourse we often find it natural to criticize people for premising mistakes in terms of *knowledge*.²⁴ This fact about us—about what we find natural to say—stands in need of an explanation. The simplest account would seem to be that we are tacitly assuming KR. In so far as one takes there to be *pro tanto* reason to prefer the more conservative philosophical theory, we thus have here a *pro tanto* reason to prefer KR.

As I've just rehearsed, there are two main components to the case for KR over a justification-based alternative. First, there are apparent counterexamples to at least the most obvious form that a justification-based view might take. And second, there is the prevalence and naturalness of knowledge-based criticism in ordinary discourse. A proponent of a justification-based view thus has two main tasks. First, she must formulate a principle that can withstand the counterexamples leveled at the simple justification-based approach. And second, she should offer an account of why knowledge-based criticisms are prevalent and natural in ordinary discourse, despite the fact that knowledge is not the norm of practical deliberation. These are the two main tasks to be taken up in the remainder of this paper.

5. AN EARLIER PROPOSAL: JBK

In light of considerations like those offered in the previous section, at least two authors have offered something of a hybrid between the justification-based approach and the knowledge-based approach. Both Ram Neta (2009) and Declan Smithies (2011) have defended variants of the

Justified-Belief-in-Knowledge Principle. When P is relevant to S's decision, it is appropriate for S to premise that *p* if and only if S is justified in believing that she knows that *p*.²⁵

²⁴ Experimental philosophers will of course be concerned that Hawthorne and Stanley make this claim without providing carefully-collected empirical data. I propose to set that objection aside. I will simply grant, for purposes of this paper, that people do in fact find knowledge-based criticisms natural in the described circumstances.

²⁵ Following Hawthorne and Stanley, both Neta and Smithies formulate their principle in terms of the more controversial notion of 'treating *p* as a reason'.

According to JBK, the norm of premising does not concern what you know, but what you are justified in believing you know. As such, JBK is a sophisticated version of the justification-based approach.

JBK handles the cases used in the subtraction argument quite nicely: in neither variant of the case of John in the coma does John know that it is 5:00 a.m., but in both cases he is *justified in believing* that he knows it is 5:00 a.m.. Hence, according to JBK, it is appropriate for John to premise that it is 5:00 a.m.. JBK also handles the case appealed to in the introduction: while Suzy is justified in believing that Z will win, it seems odd (to say the least) to say that Suzy is justified in believing that she *knows* that Z will win. Hence, according to JBK, it is not appropriate for her to premise that Z will win. Finally, it is also plausible that JBK can explain the prevalence of knowledge-based criticisms in ordinary discourse. As Ram Neta (2009: p. 693 – 695) argues, we often criticize someone by saying ‘not *p*’, when we really mean to criticize them on the basis of *not being justified in believing that p*. Hence, when we explicitly criticize someone by simply saying ‘you didn’t know’, we might really mean to be criticizing them on the basis of *not being justified in believing that they knew*.

Is JBK then the sophisticated justification-based view we’ve been looking for—the answer to the premising question? Unfortunately, it is not. In (2008), Jessica Brown offered a set of counterexamples to KR. As it happens, some of these examples *also* work against JBK. Here is one.

SURGEON

A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is lying anaesthetized on the operating table. The operation hasn’t started as the surgeon is consulting the patient’s notes. The student is puzzled and asks one of the nurses what’s going on:

Student: I don’t understand. Why is she looking at the patient’s records? She was in clinic with the patient this morning. Doesn’t she even know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn't operate before checking the patient's records. (Brown 2008: p. 1144 – 1145)

This case is a counterexample to the right-to-left direction of JBK: the Surgeon knows and *is justified in believing that she knows* that the patient's left kidney is diseased, and yet it is not appropriate for her to act on that belief without double-checking the patient's records. Imagine if she were to, without checking the records, simply premise that *it's the left kidney that is diseased*, and on that basis form the intention to remove the left kidney rather than the right. Given how easy it would be to double-check the records, this would clearly be inappropriate. It's worth mentioning that *Surgeon* is also a counterexample to the right-to-left directions of both KR (as Brown points out) and JB: not only is the surgeon justified in believing that she knows that the left kidney is diseased, but she is justified in believing that the left kidney is diseased, and she knows that the left kidney is diseased.

Brown's *Surgeon* case is an instructive one. It points to the idea that there is something beyond the strength of one's evidence with respect to *p* that matters for the permissibility of acting on a belief that *p*. This something else is the *stakes* at play in the situation. Of course, so much is obvious, and proponents of JB, JBK, and KR won't deny this. Their standard solution is to allow stakes to play a role in determining whether the agent is justified in believing/knows/is justified in believing she knows that *p*.²⁶ This idea is known as 'pragmatic encroachment'.²⁷ If pragmatics encroach, then it might be that, contrary to what I said above, the high-stakes at play in *Surgeon* make it such that the surgeon is not justified in believing/does not know/is not justified in believing that she knows that the patient's left kidney is diseased. If so, then the case really isn't a counterexample to RK/JB/JBK.

I find the pragmatic encroachment defense implausible in the case of *Surgeon*: of course the surgeon is justified in believing that the patient's left kidney is diseased, of course she knows that the patient's left kidney is diseased, and of course she is justified in believing that she knows the patient's left kidney is diseased. But if you don't find the surgeon case convincing, try this one, again from Brown (2008):

²⁶ Hawthorne and Stanley propose this solution in defense of KR (2008: p. 588). Declan Smithies has proposed this solution, with some reservations, in defense of JBK (personal communication).

²⁷ See *inter alia* Fantl and McGrath (2007).

Suppose that Liz knows that she was born in England. Nonetheless, it would seem irrational for her to accept a bet with the following pay-offs:

Liz was born in England: Liz gains £1.

Liz was not born in England: Liz loses her home [and is tortured for the next thirty years]. (p. 1144)

If Liz justifiably believes/knows/justifiably believes that she knows that she was born in England, then, according to JB/KR/GBK, it is appropriate for her to premise that she was born in England, and if it is appropriate for her to premise that she was born in England, then it is appropriate for her to take the bet. Hence, since it's not appropriate for her to take the bet, then either she isn't justified in believing/doesn't know/isn't justified in believing that she knows that she was born in England, or else JB/KR/GBK is false. I for one am more confident that Liz is justified in believing, knows, and is justified in believing she knows that she was born in England than I am that any of those principles is true. If pragmatics encroach, they don't encroach *that* far.

The lesson we should draw is this. Yes, stakes play a role in determining whether it is appropriate to premise that p , but not, or at least not always, by playing a role in determining whether you are justified/know/are justified in believing that you know that p . Hence, the rational permissibility of premising that p must depend at least in part on something other than whether you are justified/know/are justified in believing that you know that p . In the next section I say what that something is.

6. PRACTICAL CERTAINTY

While proponents of KR typically approach the premising question by reflecting on *what people say*—in particular, the kinds of criticisms people make of one another in ordinary discourse—we can also take a bit more theoretical approach. Let us consider the notion of premising more directly. By thinking a bit more carefully about what premising *is*, we can gain some insight into *when it is appropriate to do it*.

Roughly, to premise that p in a particular piece of deliberation is to deliberate in a way that ignores, at least for the purposes of that deliberation, the possibility that not- p . When Suzy premises that Z will win, she ignores—or perhaps “sets aside”—the possibility that Z will not win. But to ignore the possibility that not- p in one's practical

deliberation is to deliberate as if there is no epistemic possibility that not- p —that is, to deliberate as if p is epistemically certain. Putting this together, we can say that to premise that p in one's practical deliberation is to deliberate as if p is epistemically certain. Hence the premising question becomes this: when is it rational to deliberate as if p is epistemically certain?

It is rationally permissible to deliberate as if p is epistemically certain just in case p is epistemically certain *enough*. But how certain is certain enough? Certain enough that the difference between the actual epistemic probability, on the one hand, and epistemic certainty, on the other, is not a difference that makes a difference to what one has most reason to do. In such a case, let us say that p is 'practically certain'.

p is **practically certain** for S relative to her decision D if and only if the actual degree of certainty of p for S is such that the act which S has most reason to do is the act that S would have most reason to do were p to be epistemically certain.²⁸

Must someone who deliberates as if p is epistemically certain actually *believe* that p is practically certain? No, for one need not have even considered the matter. However, I suggest that it is rationally permissible to deliberate as if p is epistemically certain if and only if p is in fact practically certain. Hence I propose

The Practical Certainty Principle. When p is relevant to S 's decision D , it is appropriate for S to premise that p if and only if p is practically certain for S relative to D .

PC gives stakes a role in determining whether it is rationally permissible to act on a given belief. Consider *Surgeon*. Given how easy it is to double-check the patient's records, and given just how bad it would be to remove the wrong kidney, the surgeon *in*

²⁸ Two notes are in order here. First, decisions are to be understood as choices between a fixed set of options. Hence, p might be practically certain relative to one set of options, but not relative to another. Second, the subjunctive mood used in this definition is not quite right. As any philosopher knows, it is possible to design cases in which all manner of things might have been the case had something else been the case. Here I must rely on the good faith of my reader to interpret this definition in the way it is intended. Roughly, the idea is simply that p is practically certain if and only if the difference between the actual epistemic probability that p and epistemic certainty that p is not a difference that makes a difference with respect to what one has most reason to do.

fact has most reason to double-check the patient's records. But if it were epistemically certain for the surgeon that it was the left kidney that was diseased, then given the (small) cost of double-checking the patient's record, the surgeon would have most reason to proceed with the surgery without double-checking.²⁹ Hence, the stakes are such that it is not practically certain—for the surgeon, relative to her current decision—that the patient's left kidney is diseased. Thus, by PC, it is not appropriate for her to premise that *p*.

Contrast *Surgeon* with the case of Lucy from the introduction. Recall that Lucy sees that it is raining and then premises that it is raining in the following piece of reasoning.

It's raining.

If it's raining, then taking an umbrella is the optimal thing to do.

So I'll take an umbrella.

Provided that the details of the case are filled-in in the usual way, it is rationally permissible for Lucy to premise that it is raining in this piece of deliberation. PC accounts for this fact. In Lucy's case, it is practically certain that it is raining. Given that she has just seen the rain, the epistemic probability that it is raining is near one. Given such a high epistemic probability, and given what is at stake, taking an umbrella is what Lucy has most reason to do. But this is exactly what Lucy would have most reason to do were it epistemically *certain* that it is raining. Thus, by PC, it is appropriate for Lucy to premise that it is raining.

Note that PC gives the stakes at play in a particular decision an independent role in determining whether it is appropriate to premise—independent, that is, of any alleged effect that stakes may have on whether the agent's belief is justified or amounts to knowledge. Consider Suzy's case from the introduction. Suppose that it is as epistemically certain (for Suzy) that *Z* will in as it is epistemically certain (for Lucy) that it is raining. Still, if the stakes are *extreme enough*, betting on *Z* might not be the thing

²⁹ One might worry that if *p* were epistemically certain, then the surgeon would have most reason to do something else entirely—say, place some bet at extraordinarily low odds on whether the patient's left kidney is diseased. To handle this worry, we need only keep in mind that the implicit quantifier in 'what she would have most reason to do' is to be understood as ranging over a fixed set of available options—namely, those that the agent is currently deliberating over. Please see the previous footnote.

to do. Hence, because betting on Z would be the thing to do, were it epistemically certain for Suzy that Z would win, it follows that, according to PC, it is not appropriate for Suzy to premise that Z will win.

It is also worth noting that, by denying that the rational permissibility of premising that p depends on whether one knows that p , and instead insisting that it depends on whether p is practically certain, PC undermines Jeremy Fantl and Matthew McGrath's (2009, ch. 7) argument that we must choose between infallibilism and pragmatic encroachment. Fantl and McGrath note that no matter how strong someone's evidence is, provided that it falls short of establishing epistemic certainty, we can always design a case such that it is not rationally permissible for the agent to premise the proposition in question—we need only make the stakes extreme enough. If KR were true, this would mean that either it is impossible to have knowledge without having epistemic certainty (infallibilism) or else the standards for knowledge are sensitive to the stakes (pragmatic encroachment). Rejecting KR in favor of PC undermines this argument, and thus is quite friendly to those who wish to defend fallibilism but reject pragmatic encroachment. Note, however, that PC is also consistent with infallibilism and pragmatic encroachment, should one want to accept either. The point here is simply that, as far as PC is concerned, one is not forced to choose between just these two options.³⁰

With some development, PC also gives us the resources to handle slightly more complicated cases. Consider Kathryn's case, which is just like Suzy's, except that in Kathryn's case the evidence *does* support the proposition that Z will win strongly enough to make betting on Z the thing to do. More precisely, the epistemic probability of Z's winning—that is, the probability of Z's winning given Kathryn's evidence—is such that the expected value of betting on Z is *slightly* higher than the expected value of not betting on Z, and hence (we'll suppose)³¹ Kathryn has most reason to bet on Z. However, Kathryn, reacting to the *gross* fact that Z is more likely to win than not, reasons in precisely the same way that Suzy does.

Z will win.

³⁰ I owe thanks to an anonymous referee for encouraging me to highlight this feature of PC.

³¹ More precisely: we'll suppose that the weight of all *other* reasons to bet/not bet balances out, and so the reasonableness of Kathryn's betting depends exclusively on whether betting has a higher expected monetary value for her than not betting.

If Z will win, then betting on Z is the optimal thing to do.

So I'll bet on Z.

Does Kathryn commit a promising mistake? There is a sense in which the answer is 'yes', and a sense in which the answer is 'no'.

According to PC, it is appropriate for Kathryn to premise that Z will win. However, it does not follow that Kathryn *appropriately* premises that Z will win. The distinction here is between what some philosophers call 'ex ante evaluation' and 'ex post evaluation'.³² On the standard way of drawing the distinction, it is appropriate for an agent to ϕ if and only if she has most normative reason to ϕ , whereas an agent ϕ 's appropriately if and only if she has most normative reason to ϕ and her ϕ -ing is *based on* or *sensitive to* the fact that she has most normative reason to ϕ .³³ According to PC, Kathryn's normative reason to premise that Z will win is that it is practically certain that Z will win. However, Kathryn's promising that Z will win is not sensitive to whether it is practically certain that Z will win. I said that the epistemic probability of Z's winning is such that the expected value of betting on Z is *slightly* higher than the expected value of not betting on Z. Hence, if Z's winning had been slightly less probable, not betting on Z would have been what Kathryn had most reason to do, and so betting on Z would not have been practically certain. But I also said that when Kathryn premises that Z will win, she is merely reacting to the *gross* fact that Z is more likely to win than not—that is, she is not reacting to any *fine grained fact* about the more or less precise probability that Z will win. Hence, if the probability that Z will win had been slightly lower, and so it was not practically certain that Z would win, Kathryn still would have premised that Z will win. Hence, Kathryn's promising that Z will win is not sufficiently sensitive to the

³² The terms 'ex post' and 'ex ante' were originally applied to this distinction by Alvin Goldman (1979). The *ex post/ex ante* distinction is a generalization of the familiar epistemological distinction between propositional justification (*ex ante*) and doxastic justification (*ex post*). It is controversial, however, just how far this distinction generalizes—in particular, it is controversial whether it generalizes to the case of action. J. S. Mill (2001, p. 18) and W.D. Ross (1988, ch. 2) would argue that *ex post* evaluation makes little sense—according to these theorists, agents merely have duties to perform the right acts; they do not also have duties to perform those acts *for the right reasons*. Kant, on the other hand, seems to accept both kinds of evaluations, at least on one standard interpretation of his view—for Kant, agents ought to act in accordance with the moral law, and they also ought to do so *out of respect* for the moral law. Thanks to an anonymous referee for reminding me that it is controversial just how far the *ex ante/ex post* distinction can be generalized.

³³ For more on drawing the distinction through the use of adjectives/adverbs, see Matthew Hanser (2005).

practical certainty of Z's winning, and hence, according to PC, Kathryn does not *appropriately* premise that Z will win. This verdict seems intuitively correct.

Is there a way for Kathryn to premise such that her so premising *is* sensitive to whether it is practically certain that Z will win? One way that Kathryn could do so is to first calculate the threshold for practical certainty—that is, she could calculate the degree x such that Z's winning is practically certain if and only if the epistemic probability that Z wins is greater than or equal to x . She could then do her best to determine the epistemic probability y that Z will win and then premise that Z will win if and only if y is greater than or equal to x . If that is the process by which Kathryn premises that Z will win, then her premising that Z will win is sensitive to her reasons for believing that Z's winning is practically certain, and hence she appropriately premises that Z will win. But such a process would be silly: if she is going to go through these calculations, it would be easier to simply use y to determine whether betting on Z has higher expected value than not betting on z , and then on the basis of *that determination*, decide whether to bet on Z. There's no need to go through the further trouble of premising that Z will win and thereby deciding what to do.

Is going through these explicit calculations the only way for Kathryn to make her premising that Z will win sensitive to whether it is practically certain that Z will win? I don't know. Certainly there are *logically possible* creatures whose premising practices are sensitive to fine-grained differences in epistemic probabilities and stakes, even when they don't engage in the kind of explicit calculations described above. Perhaps Kathryn is one of these creatures, contrary to what I've been assuming all along. Nonetheless, it seems that for any given creature, there will be a limit to just how sensitive their non-calculative premising practices are. Hence, for any given creature, there will be possible cases where it is impossible for that creature to premise p appropriately without making the explicit calculations.

In this section I have offered an answer to the premising question that succeeds where previous proposals have failed. PC accounts for our intuitive judgments about all of the cases considered in this paper, and indeed all of the cases of which I am aware. Moreover, PC accounts for these cases without appeal to any controversial auxiliary assumptions, such as the assumption that pragmatics encroach so far into knowledge/justified belief that many of the cases that we would ordinary count as knowledge/justified belief (e.g., *Surgeon*) are actually not. This makes for quite a strong case for PC over its rivals. However, there remains one more issue with which we must

deal: if PC, and not KR or JBK, is the norm of practical deliberation, then why do we so often criticize people for premising that p by pointing out that they do not *know* that p ? I take up this issue in the next and last section.

7. ORDINARY DISCOURSE RECONSIDERED

Consider the following dialogue, paraphrased from a recent episode of AMC's "Hell on Wheels".

Cammie: Why did you shoot that man?!

Bridget: Because he was going to kill me.

Cammie: But you didn't know that!

It's easy to imagine scenarios where Cammie's criticism seems quite appropriate. For example, the man was carrying a weapon and approaching Bridget with a menacing look on his face, but there was no further reason that Bridget had to think that the man was going to kill her. The naturalness of this discourse is a big part of why principles like KR and JBK initially seem correct. If these principles are false, why does Cammie's knowledge-denial at least *seem* to the point?

The answer, I suggest, concerns what is pragmatically implicated by assertions of 'S does not know that p '. To know is, in part, to believe. I argue elsewhere (Locke Manuscript B) that to believe that p is, in part, to be disposed to premise that p in at least *some* decision contexts. A similar view has been defended by Jacob Ross and Mark Schroeder (2012). On my view, there is no set range of decision contexts where one must be disposed to premise that p in order to count as believing that p . In particular, it is not required that an agent be disposed to premise that p in *every* decision context where p is relevant. However, if part of what it is to believe that p is to be disposed to premise that p in *some* decision contexts, then we should expect that assertions of 'S believes that p ' often pragmatically implicate (but do not semantically entail) that S is currently in one of the decision contexts in which she is disposed to premise that p . Hence, since part of what it is to know is to believe, we should also expect that assertions of 'S knows that p ' often implicate that S is currently in one of the decision contexts in which she is disposed to premise that p .

Knowing also requires that one's belief be justified. It follows from my partial account of belief, together with some plausible auxiliary assumptions, that to be justified

in believing that p , it must be appropriate for one to premise that p in those decision contexts where one is disposed to do just that. Hence, we should expect that assertions of ‘S justifiably believes p ’ implicate that S is in a decision context in which it is appropriate for her to premise that p . Hence, since part of what it is to know is to justifiably believe, an assertion of ‘S knows that p ’ will often implicate that S’s current decision context is one of those in which she is disposed to premise that p and that it is appropriate for her to premise that p in that decision context. Hence, in contexts where it is not appropriate for S to premise that p , an assertion of ‘S knows that p ’ will often seem inappropriate. Similarly, assertions of ‘S does not know that p ’ will often seem appropriate, for they will often pragmatically implicate that it is not appropriate for S to premise that p .³⁴

Applying this account to our example, once Cammie has asked Bridget why she shot the man, and Bridget has replied, ‘Because he was going to kill me’, it is now common ground between them that Bridget premised that the man was going to kill her. Hence, when Cammie replies ‘But you didn’t know that’, she implicates that it was not appropriate for Bridget to premise that the man was going to kill her. Cammie’s knowledge-denial implicates this, not because knowledge is the norm of premising, but because to know is in part to justifiably believe and to believe is in part to be disposed to premise in practical deliberation.

8. CONCLUSION

In section four I set out two tasks for a defender of the justification-based approach to the premising question. The relatively minor task was to explain why it is that knowledge-based criticisms are prevalent in ordinary discourse, despite the fact that knowledge is not the norm of premising in practical deliberation. In the previous section I gave such an explanation. While my explanation did appeal to a particular conception of the nature of belief, it did not appeal to PC or any other principle discussed above. Hence, that explanation is in principle open to proponents of views other than PC, and I invite them to take it up.

The major task described in section four was to offer a plausible justification-based principle that withstands the counterexamples to the simple justification-based

³⁴ The account offered here is a version of what Keith DeRose (2002) calls a ‘warranted assertability maneuver’ or ‘WAM’. However, as I explain elsewhere (Locke, Manuscript C), my account differs in an important way from earlier WAMs, in that it proceeds in terms of the nature of belief, rather than in terms of the nature of knowledge *per se*.

view—that is, JB. In section six I did this and more: PC handles not only the cases that cause problems for JB, but also those cases that provide counterexamples to RK and JBK. In fact, in so far as I know, PC handles all of the cases that have been discussed in the literature on the premising question. Admittedly, the price we have to pay is that of endorsing a principle that is not nearly as “neat” as JB, KR, or JBK. But the notions employed in PC are all familiar from other areas of philosophical work. Moreover, the “neat” principles just don’t seem to get the job done. The slight complexities of PC therefore look to be just what we need to correctly describe when it is appropriate to premise in practical deliberation.

WORKS CITED

- Broome, John and Christian Piller. 2001. ‘Normative Practical Reasoning’. *Proceedings of the Aristotelian Society, Supplementary Volumes* 76: 175 – 193 + 195 – 216
- Brown, Jessica. 2008. ‘Knowledge and Practical Reasoning’. *Philosophy Compass* 3: pp. 1135 – 1152
- DeRose, Keith. 2002. ‘Assertion, Knowledge, and Context.’ *The Philosophical Review* 111: 167 – 203
- Dancy, Jonathan. 2000. *Practical Reality*. Oxford: Oxford University Press
- Fantl, Jeremy and Matthew McGrath. 2009. *Knowledge in an Uncertain World*. Oxford University Press
- Fantl, Jeremy and Matthew McGrath. 2007. ‘On Pragmatic Encroachment in Epistemology’. *Philosophy and Phenomenological Research* 75: 558 - 589
- Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge: Harvard University Press

- Goldman, Alvin. 1979. 'What is Justified Belief?' in *Justification and Knowledge*. George Pappas (ed.). Boston: D. Reidel
- Hanser, Matthew. 2005. 'Permissibility and Practical Inference'. *Ethics* 115: 443 – 470
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press
- Hawthorne, John and Jason Stanley. 2008. 'Knowledge and Action'. *The Journal of Philosophy* 105: 571 - 590
- Hornsby, Jennifer. 2007. 'Knowledge, Belief, and Reasons for Acting' in *Explaining the Mental*. Carlo Penco, Michael Beaney, Massimiliano Vignolo (eds). Cambridge Scholars Publishing: 88 – 105
- Hyman, John. 1999. 'How Knowledge Works'. *The Philosophical Quarterly* 49: 433 – 451
- Hyman, John. 2011. 'Acting for Reasons: Reply to Dancy'. *Frontiers of Philosophy in China* 6: 358 – 368
- Ichikawa, Jonathan Jenkins. 2012. 'Knowledge Norms and Acting Well'. *Thought* 1: 49 – 55
- Kaplan, Mark. 1985. 'It's Not What You Know that Counts'. *The Journal of Philosophy* 82: 350 – 363
- Kaplan, Mark. 2003. 'Who Cares What You Know?' *Philosophical Quarterly* 53: 105-116
- Littlejohn, Clayton. 2009. 'Must we Act Only on What We Know?' *The Journal of Philosophy*. 106: 463 - 473
- Locke, Dustin. Manuscript A. 'Knowledge, Explanation, and Motivating Reasons'.

- Locke, Dustin. Manuscript B. 'The Decision-Theoretic Lockean Thesis'.
- Locke, Dustin. Manuscript C. 'Pragmatic Encroachment and the Nature of Belief'.
- Mill, J. S.. 2001. *Utilitarianism*. Hackett Publishing
- Neta, Ram. 2009. 'Treating Something as a Reason for Action'. *Nous* 43: 684 – 689
- Parfit, Derek. 2011. *On What Matters* Vol 1 and 2. USA: Oxford University Press.
- Ross, W. D. 1988. *The Right and the Good*. Hackett Publishing Company.
- Scanlon, T. M.. 1999. *What we Owe to Each Other*. Cambridge, Mass.: Harvard University Press
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press
- Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell
- Smithies, Declan. 2011. 'The Normative Role of Knowledge'. *Nous* 46: pp. 265 – 288
- Unger, Peter. 1975. *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press
- Weiner, Matthew. 2005. 'Must We Know What We Say?' *Philosophical Review* 114: 227 – 251