How I Learned to Stop Worrying and Love Probability 1

Daniel GReco

Penultimate Draft, Final Version in *Philosophical Perspectives*

Much recent work in epistemology concerns the relationship between binary notions like belief, and graded notions like degree of confidence or probability, also called "credence".[1] In particular, while earlier work tended to be broadly reductionist in one or the other direction, it's now common to adopt what might be called a "pluralist" approach to the two families of notions, holding that each is of independent epistemological interest and worth studying in its own right, not merely as an epiphenomenon of the other.[2] These defenses of pluralism tend to take place against the backdrop of the thought that a certain simple view of the relationship between the two notions is off the table, because it is inconsistent with an extremely plausible, broadly fallibilist approach to epistemology. The simple view is the view that binary belief is maximal degree of belief—it is the endpoint of the scale of degreed belief. This view is typically rejected on the grounds that good post-Cartesians like us aren't maximally confident of *anything*, so if belief is going to be of any interest at all, it must not involve maximal confidence. My central aim in this paper is to defend the simple view. My main target will be the pluralist who accepts the importance of binary belief, but rejects the simple view on the broadly

[1]While I'll focus in this paper on belief and credence, I hope that much of what I say in this paper will also concern the relationship between the binary notion of knowledge, and graded notions like degree of evidential support.

[2]For some paradigm examples, **?** elevates graded notions, holding binary notions to be philosophically uninteresting. **?** is only slightly more moderate, arguing that binary notions can be reduced to graded ones. In the other direction, both **?** and **?** (Forthcoming) hold that graded notions can be reduced to binary ones. Recent examples of pluralism include **?**, **?**, **?**, and **?**. Isaac **??** is in some ways a tricky case, though should probably be classed as a pluralist.

fallibilist grounds just mentioned. I'll argue that this view faces strong internal tensions; some of the most attractive motivations for taking binary belief seriously ultimately require some version of the simple view.

I'll proceed as follows. In §1 I'll fix ideas, clarifying my use of the terms "belief", and "pluralism". In §2 I'll offer some prima facie motivations for the simple view. In §3 I'll offer a more general theoretical framework from which the simple view follows. In §4 I'll introduce the fallibilist objections to the simple view, along with a "sensitivist" response explicitly defended by Roger **?**, and implicitly defened (so I'll suggest) by **?**. My main aim in that section will be to argue that this response is not ad hoc maneuver, but is instead independently motivated, especially for the pluralist. Lastly, in §5, I'll consider a natural fallback position for the opponent of the simple view, on which avoiding maximal confidence is a rational *ideal*, even if it's one we cannot ourselves live up to. While I won't try to show that this position is untenable, I will argue that it has some stronger commitments than have generally been appreciated.

1

The idea that we are maximally confident in what we believe can seem absurd if we focus on certain natural language constructions. If you ask me whether the bank is open on Saturdays, "yes, I believe" is a much more cautious answer than "yes".[3] "I believe" is a device used for making guarded or "hedged" assertions, rather than outright ones. So if by "belief" we mean whatever mental state is typically reported by locutions like "I believe", then the idea that belief involves *maximal* confidence is obviously false; the point of the locution is to convey that one is *less* than maximally confident.

For this reason, if the claim that belief involves maximal confidence is to be worth taking seriously at all, we cannot be working with a conception of belief closely tied to

[3]See, e.g., **?**.

natural language constructions involving "belief" and "believe". Much work in epistemology suggests an alternative conception of belief, more closely related to knowledge. Call this conception "strong belief" (or "full" or "outright" belief). In a nutshell, the idea is that believing, in the strong sense, feels just like knowing. Depending on one's views about knowledge and belief, there are various strategies one might use to flesh out this connection. Someone optimsitic about the traditional project of analyzing knowledge in terms of necessary and sufficient conditions might hold that belief is the mental component of knowledge; knowledge involves being in a certain mental state (strong belief), *and* various additional conditions holding (e.g., the belief's being true, justified, etc.)[4] But even if we are skeptical about that analytical project, there are various ways we might get at a strong notion of belief, closely tied to knowledge. **?** discusses a conception of strong belief constrained by the following principle: if you strongly believe that p, then you strongly believe that you know that p. While **?** doesn't endorse this conception of strong belief, he does suggest a related view on which "believing p is conceptualized as being in a state sufficiently like knowing p 'from the inside' in the relevant respects" (p. 2). When belief is understood along some version of these lines, the simple view that (strong) belief involves credence 1 is once again a live option.[5]

But *is* there an epistemologically significant notion of strong belief, along anything like the above lines? Much work in the Bayesian epistemological tradition argues (or presupposes) that there is not. Richard Jeffrey was famously happy to discard the traditional conceptions of belief and knowledge in favor of that of degree of belief, suggesting that Frank Ramsey "sucked the marrow out of the ordinary notion, and used it to nourish a more adequate view." (**?**, p. 171-2) And while David **?** never makes quite such strong claims, one of the central themes in *Putting Logic in its Place* is that the interesting principles to be found in epistemology are at the level of degree of belief—epistemological

[4]See **?**, who attributes this view to Stephen Stich and Jaegwon Kim, though himself argues against it.

[5]In what follows, when I write "belief" without qualification, the reader should assume I'm referring to strong belief unless otherwise specified.

generalizations about belief are rules of thumb, whose truth (or approximate truth) is always explained by more fundamental principles governing rational degrees of belief. For instance, he holds that the rule that beliefs should be consistent is subject to exceptions, and its approximate truth is explained by the more fundamental rule that degrees of belief should be probabilistically coherent. And while he doesn't explicitly endorse this, a natural corrolary involves holding that any rules connecting belief and rational action will be subject to exceptions, and their approximate truth will be explained by the more fundamental rule that one should act so as to maximize expected utility.

The *pluralist* rejects this picture; she holds that there are fundamental epistemological and/or decision theoretic principles governing belief, as well as credence.[6] For example, the pluralist might endorse a belief rule of assertion, according to which one should only assert what one believes, and hold that the rule is not reducible to any rule concerning credence. She might support her case by pointing to the example of lottery propositions—e.g., the proposition that one's ticket won't win a large, fair lottery—as propositions one is highly confident in, but which one does not strongly believe, and which one should not assert. The pluralist might also hold that there are genuine principles of rationality concerning the consistency of belief and the relationship between one's actions and one's beliefs, rather than merely rules of thumb explained by more fundamental principles concerning rational degrees of belief.

One of the main routes to pluralism, especially in recent work, goes via the idea

[6]Of course, one could reject pluralism by holding that all the fundamental epistemological principles concern only belief, and not credence. This is suggested by a remark by Gilbert **?**, p.22:

> How should we account for the varying strengths of explicit beliefs? I am inclined to suppose that these varying strengths are implicit in a system of beliefs one accepts in a yes/no fashion. My guess is that they are to be explained as a kind of epiphenomenon resulting from the operation of rules of revision. For example, it may be that P is believed more strongly than Q if it would be harder to stop believing P than to stop believing Q, perhaps because it would require more of a revision of one's view to stop believing P than to stop believing Q.

My focus in this paper, however, will be on pluralists in the first place, and those who would reduce belief to credence in the second. I will not discuss the sort of reduction Harman contemplates, which would go in the other direction. See also **?** (Forthcoming).

that belief is a necessity for cognitively limited agents like us. Jacob Ross and Mark Schroeder capture the idea nicely:

> If we had infinite cognitive resources, then we'd have no need for an attitude of outright belief by which to guide our actions, for we could reason in an ideal Bayesian manner on the basis of our credences and preferences alone. But such reasoning isn't feasible for cognitively limited agents like us, and so we need an attitude of *outright belief* or of *settling on the truth of propositions*, so as to limit what we consider in our reasoning to possibilities consistent with what we have settled on. (**?**, pp.30-1)

For most of this paper, I'll presuppose that pluralism is correct and that this anti-idealizing motivation for it is on the right track, and will argue for the simple view that we have credence 1 in what we believe, given that starting point. Later on I'll do some work to defend those presuppositions. But before moving on, I want to make clear that my target is not a straw man. It's true that much recent work holds that believing that P involves somehow ignoring possibilities in which P is false, which might sound a lot like assigning such possibilities probability 0. In addition to Ross sand Schroeder, see **?**, **?**, and **?** (Forthcoming). But a common theme in all these works is that even if one treats one's beliefs *as if* one were fully confident of them *in some respects*, one generally is *not* fully confident in one's beliefs—it is possible, indeed typical, for one to have less than maximal credence in propositions one believes. Ultimately, I'll argue that when credences are as independent from beliefs as these views would have it, they cannot play the roles they must.

2

Why think that we have maximal credence in our (outright) beliefs? In a nutshell, the idea is that various principles and norms that plausibly govern belief—at least, assuming

pluralism is correct, and there are *any* principles and norms that govern belief—are easy to explain if we have credence one in what we believe, and harder to explain otherwise. Let's start with the example of norms of assertion, mentioned above. We can distinguish outright, unqualified assertions from assertions that are "hedged" in various ways. Consider the following schematic examples:

1. More likely than not, P

2. Very probably, P

3. With at least 99% probability, P

4. P

Whatever sentence we plug in for P, 1-4 will naturally be heard as expressing increasing levels of confidence. Lottery cases are notorious examples where 1-3 would be appropriate, but 4 would not. The pluralist who accepts the simple view has an easy time explaining this. Suppose, e.g., in addition to the simple view, we accept a knowledge rule of assertion.[7] Then we can offer the following explanation. For an outright, unqualified assertion that P to be appropriate, one must know that P. To know that P, one must (strongly) believe that P. And to believe that P, one must have credence 1 in P—any lesser credence precludes belief, therefore precludes knowledge, therefore precludes proper assertion. So whenever one is aware that one has less than maximal credence that P, if one is attempting to comply with the norm of assertion, one will refrain from making unqualified assertions that P. But assertions like 1-3 are plausibly governed by different norms—e.g., 1 is plausibly governed by a norm that requires only that one have credence >0.5 that P, 2 governed by a norm that requires a higher credence, and 3 governed by a norm that requires a credence of 0.99.[8] None of these norms

[7]See, e.g., **?**, Ch. 11.

[8]Actually, I suspect the most elegant account available will offer a unified explanation of the norms governing 1-4, and that in each case, the norm will involve knowledge. E.g., asserting 1 will require

require that one have maximal confidence for one's assertion to be appropriate. The desired contrast is captured.[9]

But if the simple view is false, and assertion does not require maximal credence, then it's harder to see how to explain the contrast between hedged, probabilistic assertions, and unqualified, outright ones. I don't claim that it's impossible, only that competitor views need to do serious work to explain data that is utterly unsurprising, given the simple view.[10]

The dialectic is similar concerning norms of consistency. Many writers have thought that there's something wrong with having inconsistent (full) beliefs.[11] This claim is even more plausible if we accept some tight connection between belief and assertion, along the lines spelled out in the above paragraphs. After all, it's plausible that in making inconsistent assertions, one violates some kind of norm—if beliefs are subject to a norm of consistency, and appropriate assertion requires, *inter alia*, that one believe what one asserts, then we'd have an explanation of the impropriety involved in making inconsistent assertions.[12]

Given the simple view, a consistency norm on belief is entirely to be expected. If a subject's beliefs get credence 1, and that subject's credences form a probability function, then her beliefs will be consistent. This claim is subject to This is a straightforward

knowing that it's more likely than not that P, asserting 2 will require knowing that it's very likely that P, etc.. To have an account of what this involves, we'll need a theory of probabilistic knowledge. But luckily, such theories are available. See **?** (Forthcoming) for a unified account of the norms governing assertions like 1-4, which will still support essentially the same argument I offer above.

[9]See **?** for similar motivation for the view that beliefs require credence 1.

[10]For example, **?** develops a competitor to the simple view designed to capture intuitions about lottery cases, and consistency norms for belief. But while his treatment certainly handles *some* versions of lottery cases—in particular, lotteries with a guaranteed winner, where each ticket has the same chance of winning—it falters when applied to lotteries with no guaranteed winner, or where some tickets have better chances than others, as **?** (Forthcoming) convincingly argues. Such lotteries provide no special challenges for the simple view. So while Staffel dismisses the simple view as "not a serious contender", it seems to me that her arguments do a nice job of showing how alternatives to the simple view have a much harder time capturing its virtues than might initially be apparent.

[11]This is among the data about outright belief that **?** and **?** try to capture.

[12]I don't mean for this explanation to necessarily contrast with explanations involving knowledge. E.g., a norm of consistency on belief might be derived from a norm to the effect that beliefs should be known—see, e.g., **?** **?**. From the factivity of knowledge, and the fact that truths are consistent with truths, it would follow that normatively appropriate bodies of belief are consistent.

consequence of the probability axioms. But in the absence of the simple view, a consistency norm on belief is surprising, and requires serious work to explain. In general, there's nothing wrong with having high credence in each of the members of some set of inconsistent propositions. Lottery cases provide an example—in such examples, one may perfectly appropriately have high credence in each proposition of the form ⌜ ticket n won't win ⌝, as well as in the proposition that some ticket will win. If belief is some species of high credence, then the probability calculus alone provides no guarantee that it should be subject to a consistency norm. Of course, there are pluralists who hold that belief doesn't require maximal credence but is nevertheless subject to a consistency norm, and set themselves the task of explaining how this can be.[13] My point is just that the theorist who accepts the simple view has a much easier time here than the theorist who rejects it.

3

In the previous section I offered two prima facie motivations for the simple view—two examples of plausible theses (at least to pluralists) that are easy to explain if we accept the simple view. In this section I'll offer a more general view about the relationship between belief and credence, that would explain *why* the simple view must be true.

I've already mentioned the idea that belief is a necessity for cognitively limited creatures like us. Why should this be? The basic idea has been nicely expressed by Richard **?** (Forthcoming), drawing on Michael Bratman's work in the philosophy of action:

> There are very many things that, consistent with our wants, beliefs and values, we could do. We could keep all of these, or at least, many of them, as open possibilities. But if we foreclose our options we arrive at a simpler picture which enables us to achieve things that would otherwise be beyond

[13] As already mentioned, **?** and **?** are two recent examples.

> us. Perhaps there are creatures who can coordinate any number of open possibilities. But for limited creatures like us, reducing the options to a number of fixed points is essential. It enables us to coordinate with our (equally limited) peers, and indeed to coordinate with ourselves: to keep track of what we need to do to implement our own plans. It also provides a way to keep the time spent on deliberation in check. Once we have deliberated long enough we form our intentions and then can turn our attention elsewhere...
>
> The idea I explore here is that much the same considerations govern belief. Just as we can imagine creatures who do not foreclose their practical possibilities, so we can imagine creatures who do not foreclose their epistemic possibilities. By this I mean that they do not resolve their uncertainties into all-out belief... A creature that had credences would benefit from the ability to keep many options open at once: such a creature would be able, for instance, to keep in play four mutually incompatible outcomes, regarding each with a credence of 0.25. But such abilities come at a price... Unless their powers of memory and reasoning are very great, those who employ credences risk being overwhelmed by the huge mass of uncertainty that the approach generates. (p.2)

As the above quote suggests, not only does Holton think we need outright belief, but he also thinks we do *not* need credences—he is no pluralist. But I think we can accept the basic motivation that Holton gives for thinking that limited creatures like us need beliefs, while still holding that we *also* need credences.

For now, let's take on board the idea, motivated by Holton, Ross and Schroeder, and others, that a central function of belief is to let us ignore the vast majority of logically possible scenarios, and instead focus our deliberation and planning on a manageable number of serious possibilities.[14] Where, if at all, does credence fit in once we accept

[14]See also **?**.

this idea?

Credence comes in because often, our beliefs don't settle what we should do. There may be no course of action open to me such that I have an outright belief that, if I pursue it, things will go for the best. Rather, we often lack outright beliefs about these matters; we don't take ourselves to know which choice would turn out best. Even if I've foreclosed the vast majority of epistemic possibilities, and am focusing on just two possible states of the world—e.g., rain, or shine—if one action would turn out best in one state, and another action in the other state, then the rule: "do what you believe will be for the best" will not provide me with any actionable advice. But if I have *credences* about these states, then a rule that tells me to maximize expected utility may still be useful. E.g., if a meteorologist has reported a 40% chance of rain, and I trust her, then I may decide to carry an umbrella.

So even agents like us, who are computationally limited, cannot get by *just* with outright belief; we need something like credences too.[15] The picture I'm suggesting is that credences come into the picture only *after* beliefs have done their job. Beliefs let us focus on a small number of possibilities, and then credences let us make finer grained distinctions *among the unignored possibilities.* If we think of a credence function as a function from some event space to real numbers (probabilities), the claim is that the relevant event space is given by the possibilities compatible with the agent's beliefs. And because we accept the anti-idealizing picture of belief, this will be a small, tractable space.

Given the above, the simple view that beliefs get credence 1 trivially follows. The reason is that any probability function defined over an event space Ω must assign Ω itself—and any proposition true throughout Ω—probability 1. So if an agent only outright believes a proposition if it is true throughout the set of her doxastic possibilities—

[15] I'm not sure whether this claim is really inconsistent with the picture Holton defends—I'm not sure whether what I call credences here he'd think of as, instead, outright beliefs about probabilities. Also, I haven't said anything to motivate the idea that the credences we need must form a probability function, as is typically assumed in Bayesian contexts.

contrapositively, if she only fails to believe a proposition if there is some doxastic possibility for her in which it is false—and this set of doxastic possibilities gives us the event space for her credence function, then any proposition she believes must get credence 1.

This picture of the relationship between belief and credence is inspired by a similar view in the philosophy of language. It starts with the idea that conversation takes place against a "common ground" of mutual presuppositions. These presuppostiions can be modeled as a set of possibilities—the set taken seriously for the purposes of the conversation.[16] Assertion, on this picture, can be understood as a proposal to change the common ground—a proposal that we take seriously only those possibilities compatible with what's asserted. Suppose we are discussing what to do this afternoon. I assert that it will rain. My assertion can be understood as a proposal to change the common ground so that in the rest of the conversation, it be taken for granted that it will rain. If my interlocutors accept my assertion, then conversation will evolve in a certain way—e.g., since they'll be taking rain for granted, they probably won't suggest that we go for a walk in the park.

This picture can be naturally extended to handle the sort of hedged, probabilistic assertions mentioned earlier. First, we think of the common ground as involving presuppositions not just about *which* possibilities are taken seriously, but also about *how likely* those possibilities are. Probabilistic assertions can be understood as proposals to change this part of the common ground.[17] E.g., if I say "probably, it will rain", and my interlocutors accept my assertion, then while they won't presuppose that it will rain, they will treat this possibility as likely. Maybe they'll still suggest a walk in the park, but only if this suggestion can be defended even given a high probability of rain—e.g., if there are indoor activities close to the park that can serve as an attractive fallback.

How does this view about the relationship between outright and probabilistic asser-

[16]This picture is due primarily to Robert Stalnaker, originally in **?** **?** represents his most recent thinking on the topic.

[17]See **??**, **??**, **??**.

tions inspire the simple view about the relationship between outright belief and credence? We saw earlier that if we think of credences as probabilities defined over an event space given by doxastic possibilities, then any proposition that is true throughout the set of doxastic possibilities—any proposition that is believed outright—must get credence 1. Similarly, if we can think of the probabilities manipulated by probabilistic assertions as defined over an event space given by the common ground, then any proposition true throughout the common ground—any proposition commonly presupposed by the parties to the conversation—must get a conversational probability of 1. One reason this close parallel is surprising is that, while the view that outright beliefs get credence 1 is widely regarded as implausible in epistemology, the view that conversational contexts can be modeled as involving a set of possibilities along with probabilities defined over that set is comparatively mainstream in the philosophy of language.

4

However initially plausible the preceding considerations in favor of the simple view may seem, the view faces obvious objections. Very generally, the idea is that we are maximally confident in nothing, or almost nothing. Even if we don't take skeptical scenarios all that seriously, we're never 100% sure that that we're not in one. But for all that, we still have some beliefs. So belief must not require maximal credence.

This worry is often fleshed out by appeal to high-stakes decisions. If you ask me what time it is, I'll look at the lower right corner of my computer monitor and report, without hesitation, that it's 9:55 AM. That it's 9:55 seems like a pretty good candidate for something I (strongly) believe. But now suppose you ask me if I'm willing to bet my life that it's 9:55 for the chance to win a penny. Even while I continue to look at the monitor, which continues to display 9:55, I will demur. I'll allow that the computer *might* be wrong—perhaps it's a bit fast, or slow, or failed to update for daylight savings

time, etc. If we rationally reconstruct my choice in decision theoretic terms, I'm clearly *not* treating the proposition that it's 9:55 as if it has probability 1.[18]

How should the defender of the simple view respond to this challenge? One answer, considered and rejected by **?**, explicitly defended by **?**, and arguably implicit in the work of **?**, involves holding that what one believes may change in response to whether one has been offered such a bet.[19] In particular, one believes that it's 9:55 before being offered the bet—and thus, has credence 1 that it's 9:55—but afterwards, one doesn't. Why should this be? A natural answer is that which possibilities one takes seriously—which possibilities one treats as "live"—is sensitive to a wide range of situational factors, including practical ones.[20] Clarke calls the resulting package—the view that combines the claim that what we believe gets probability 1, with the claim that what we believe is sensitive, *inter alia*, to what bets we've been offered—"sensitivism", and I'll adopt that label in the rest of the paper.[21]

Adopting sensitivism can seem like a desperate, ad hoc maneuver to save the simple view of the relationship between belief and credence, but it needn't be. In the remainder of this section I'll provide three independent motivations—independent in the sense

[18]**?** calls this argument agains the simple view the "betting worry", and finds a wide range of authors appealing to it.

[19]Why do I say that the view is implicit in Hawthorne and Stanley's work? For the following reason. They defend a view on which one's knowledge is sensitive to practical considerations, as well as the norm that one should act on one's beliefs only if they constitute knowledge. If both of these views are right, and we generally are capable of doing a reasonably good job of following the knowledge norm of action, we must be able to abandon a strong belief—to stop acting as if we know it—when situational factors change in such a way that it no longer constitutes knowledge. That is, if we have a somewhat systematic tendency to at least try to conform to the knowledge norm of action, as they understand it, then our beliefs must be somewhat sensitive to practical considerations, including whether high stakes bets are available.

[20]While the picture of belief defended here is very much in the spirit of much of Isaac Levi's work, the idea that which possibilities one takes seriously should be sensitive to practical considerations in this way is one he explicitly rejects. See the dagger footnote in (**?**, pp.3-4).

[21]This is somewhat of a simplification. As Clarke puts things, outright belief *just is* credence 1. I've preferred to put things in terms of belief "requiring" or "involving" credence 1, but the difference only matters in infinitary cases. E.g., if I know that some process will randomly determine a real number between 0 and 1, I might have degree of belief 1 that the number won't be 0.5, without outright believing this. That is, all the real number possibilities between 0 and 1 might be live possibilities for me, even though each gets probability 0. I won't discuss such cases in this paper though, and I doubt Clarke meant to be taking a stand on them in equating belief with credence 1, rather than taking credence 1 to be a necessary but not sufficient condition for belief.

that they have nothing to do with the relationship between belief and credence—for sensitivism. The first comes from work in the philosophy of mind which independently motivates the idea that one's doxastic possibilities are sensitive to one's practical situation, the second involves the difficulties involved in giving a non-sensitivist account of learning from experience, and the third involves revisiting the analogy between belief dynamics on the one hand, and conversational dynamics on the other.

4.1

A number of recent papers have defended versions of the idea that belief (whether weak or strong) cannot be understood as involving a two-place relation between a subject and a proposition. Instead, they argue, we need at least one more argument place—people can believe propositions *for certain purposes* (tasks, situations, contexts), but not others. Moreover, the motivations they provide for these views have nothing in particular to do with the relationship between belief and credence.[22] Consider the following example, due to Jack Marley-Payne:

> Take the example of the inarticulate tennis player—let's call her Serena. Over the course of a rally she can execute a complex plan which involves hitting repetitive shots to first ground her opponent in one position and then wrong foot him in order to win the point. Moreover, she can calibrate her play in response to the court conditions, the abilities of her opponent, whether she desires to humiliate him etc. However, she is unable to explain what she was doing—indeed she may even say things about her play that turn out to be false. Her non-verbal behaviour suggests a belief that a cross-court shot will wrong foot her opponent while her verbal behaviour does not.

[22]See, e.g., **?** (Ms.), **?** (Ms.), **?**, **?**. Each of these papers draws in various ways on earlier ideas in **?** and **?**.

Marley-Payne's thought is that in cases like these, the question of what Serena believes *simpliciter* is misguided:

> [W]hereas belief is traditionally thought of as a two-place relation between a subject and a proposition, on my view it is a three-place relation between subject, proposition and task. No simple yes/no answer to the question 'does S believe that p' is acceptable in cases of conflicting behaviour.

The general idea here is that the very same agent might be guided by one body of information when pursuing certain tasks, but by other bodies of information when pursuing other tasks. While this might seem like straightforward inconsistency in belief,[23] and perhaps it is, the example of Serena shows that it needn't obviously involve irrationality, and needn't be eliminable by careful reasoning—I assume that she may reason as carefully as possible, but still fail to coordinate what she says with how she plays. Some knowledge may be inarticulable, and hard to integrate with the rest of one's information, while being genuine knowledge nonetheless.[24]

The sort of dependence of an agent's information (whether belief or knowledge) on her practical situation in the example of Serena is not exactly the same as the sort claimed to exist by the sensitivist. It's not that whether Serena believes that a crosscourt shot will wrongfoot her opponent depends on stakes, or what's salient, or anything quite like that. Nevertheless, the framework appealed to by Marley-Payne to handle the example applies equally naturally to cases significant for the sensitivist. That is, we can say that agents who believe that P for most purposes, may nevertheless fail to believe that P for

[23]Though see **?** for an argument that a closely related phenomenon does *not* amount to inconsistency in belief.

[24]See, e.g., **?**.

the purpose of betting their lives against a penny on whether P.[25,26]

Put more generally, the idea is that we don't have fully domain-general representations that guide our actions no matter what our situation. Rather, whatever task we're engaged in—with "task" understood broadly, so that "tasks" include conversation, and even just internal deliberation—we'll only treat certain possibilities as live, and which possibilities we do treat as live will depend on a wide range of factors, including broadly "practical" ones.[27]

4.2

Another independent motivation for sensitivism is that the sensitivist can offer a much more straightforward account of belief and credence updating than her rivals. Suppose you and I are planning to go to a party hosted by a friend of ours, and I've forgotten what time the party starts. I ask you, and you reply that it starts at 8. I have no reason to doubt what you say, so I accept your reply. Later, though, I remember that it's a surprise party, and so it's especially important to be there on time, before the guest of

[25]If we put significant weight on this motivation, we may have to slightly modify the way we characterized the sensitivist's position. Earlier, I said that the sensitivist's position is that one may abandon a belief once one has been offered a high-stakes bet on its truth. In the present framework, the point is put a bit differently—one does not *change* one's beliefs upon being offered a high stakes bet; rather, all along one believed P for many purposes, but not for the purpose of accepting a high-stakes bet. I suspect that this difference isn't a particularly significant one, however.

[26]There are other ways one might challenge the idea that the resources needed to describe cases like Serena's can also be adopted by the sensitivist. E.g., if we hold it's only in cases of *inconsistent* behavior that the motivation to relativize belief to tasks exists, then one might argue that agents don't exhibit any inconsistency in acting as if P is true for most purposes, but not for the purpose of accepting a high stakes bet; perhaps *all* of their behavior could be consistently rationalized by the view that P is highly, but not maximally, probable. I'm sympathetic to the spirit of this response, but skeptical that agents' behavior with respect to the propositions they strongly believe *can* be consistently characterized as involving assigning them high but non-maximal probability. If it could, then there would be certain sorts of exploitation to which we would be immune—see §4.2 of **?**—but which I suspect we are, in fact, vulnerable to.

[27]See **?** for a somewhat similar position, including citations of the relevant psychological literature. On his view, we don't have stored credences. Rather, when we're deliberating, various psychological mechanisms will determine which possibilities we treat as live, and then credences will be assigned to those possibilities "on the spot", as it were. That is, rather than thinking of us as having stable, domain-general credences in a very wide range of possibilities, instead credences are only assigned to possibilities after some "filter" is applied (so that we end up with a small set of possibilities, tractable for working memory), and which possibilities get through the filter is a highly context-sensitive matter.

honor. Having remembered this, I decide to double check the invitation. What should we say about how my beliefs and credences evolve over the course of this episode?

The sensitivist can offer the following, natural account. When you first tell me that the party starts at 8, I come to (strongly) believe the content of your assertion, and to thereby have a credence of 1 that the party starts at 8.[28] This change can be modeled as conditionalization—the rest of my credences in live possibilities will equal my earlier credences conditional on the party's starting at 8. Later, when I remember that it's a surprise party, I retreat from that belief, and no longer have a credence of 1 that the party starts at 8. Given these new credences, and the costs associated with arriving late, it makes sense to double check the invitation. More generally, the sensitivist can allow that much belief updating occurs by simple conditionalization. Of course, she cannot allow that *all* updating occurs by conditionalization, since she needs to allow that we can "undo" the results of conditionalization. But this is not an unfamiliar idea.[29]

What about the non-sensitivist? Everybody will accept that at the last stage of the story, my credence that the party starts at 8 is less than 1—otherwise, there's no way to make sense of the decision to double check the invitation. But the non-sensitivist cannot accept that my credence is bumped down from 1 upon recognizing the practical significance of arriving late. So she needs an account on which it was never 1 in the first place. Which means she needs a different story from the sensitivist's about how I reacted to hearing your testimony—she cannot allow that I reacted by simply updating on the content of what you said. Some natural alternatives won't work. E.g., she might suggest that, rather than updating on the *content* of what you said, I instead update on *the fact that you said it*. Given a background belief that you're highly (but not perfectly) reliable, this would lead me to have a high (but not maximal) credence that the party starts at 8. But if it's a surprsie party, it still might make sense to double check the

[28]This is the picture in **?**. See also **?** on the attractions of the "package delivery model" of communication.

[29]See, e.g., the discussion of "contraction" in **?**.

invitation given the low cost of doing so, and the high cost of being late.

This only postpones the problem—just as practical considerations might lead me to retreat from a belief in the *content* of what you said, they also might lead me to retreat from a belief concerning which words you uttered, or even, in more unusual cases, which words I thought you just uttered.[30] Since the potential for practical considerations to lead us to retreat from our beliefs is ubiquitous, the non-sensitivist probably can't allow that *any* credences are ever updated by simple conditionalization.

Why is this a problem—why can't the non-sensitivist simply hold that our credences are updated via some other procedure that doesn't require any propositions to be assigned maximal probability? The reason is closely related to the broadly anti-idealizing approach to epistemology assumed in this paper. On the non-sensitivist view, our credences must always be distributed over a very large space of possibilities—one much larger than the space of "live" doxastic possibilities. But the same considerations that make it plausible that cognitively limited creatures like us need beliefs *also* make it plausible that cognitively limited creatures like us couldn't keep track of and appropriately update credences defined over such large a space.

On the sensitivist picture, we assign non-zero credences to propositions only when they become "live", and we have various heuristics that keep the space of live possibilities tractable.[31] This is a unified story about belief *and* credence that respects the broadly anti-idealizing considerations mentioned earlier in this essay. We can put the worry to the non-sensitivist who is also a pluralist as follows. On her view, while the point of belief is to ignore certain possibilities, we nevertheless assign credences to such ignored possibilities, and update those credences over time, not via conditionalization—which would make things simpler, since it would cut down the space of relevant possibilities—but via some other method that never rules out possibilities. But then what's the

[30]See **?** (Forthcoming) for discussion.

[31]A closely related idea is stressed in **?**; instead of having stable, stored credences, on his view we assign credences to possibilities "on the fly", once we are treating the possibilities as relevant to some task.

point of ignoring possibilities? If we need to ignore them in order for deliberation and decision-making to be tractable, why don't we *also* need to ignore them for credence updating to be tractable? These points are especially pressing if credences play a central role in deliberation and decision-making under uncertainty, as they plausibly do.[32] I don't pretend to have shown that these questions are impossible to answer for the non-sensitivist, though I do hope to have identified serious tensions that they must address, at least if they take the anti-idealizing picture of belief seriously.

4.3

At the beginning of §4 I raised a familiar objection to the view that we have maximal credence in what we (strongly) believe—the challenge that there are few if any propositions that we'd be willing to bet on at arbitrarily unfavorable odds. I then considered the response that, in some sense, what we believe is sensitive to what bets we've been offered. I've since been trying to show that this isn't an ad hoc maneuver designed to save the connection between belief and maximal credence, but is instead independently motivated, both by considerations in the philosophy of mind, and by constraints that the anti-idealizing motivation for pluralism places on accounts of how credences are updated. The last such motivation I'll point to concerns ground we've already briefly covered: the close analogy between what the sensitivist says about belief, and what many philosophers of language say about conversational context and common ground. In particular, I'll argue that if we transpose the betting-based objection to sensitivism to a situation

[32]A notable dissenter is **?**, who grants that we need to treat many propositions as having probability 1 for *practical* purposes, but not theoretical ones. He holds that theoretical credences—which do not play a role in planning or deliberation, and whose sole function is to track how much justification we have for hypotheses—are defined over a vast space in which few or no possibilities are ruled out, while practical credences are defined over a much smaller space of "live" possibilities. I'm not sure of exactly the extent of my disagreement with Wedgwood, since I'm not sure I fully understand the practical/theoretical distinction, as he draws it. In particular, I'm not sure I understand the point of keeping track of degrees of justification, when those degrees are understood as systematically irrelevant to planning and deliberation. On the metaepistemological approach I favor, the ultimate raison d'etre of epistemological concepts like justification is the way they aid us in planning and deliberation. See, e.g., **?**, **?** (Forthcoming).

where some very unfavorable bet is offered not to an individual, but to a group engaged in conversation, then the mainstream account of how the conversational common ground will evolve closely parallels the sensitivist's account of how an individual's belief state evolves when offered such a bet on her own.

Suppose that Huey, Dewey, and Louie are playing a board game that uses dice, and Scrooge, seeing the potential for profit, joins the conversation.

SCROOGE: How is this game played?

HUEY: You roll the die, and then depending on how many pips are showing, you go that many spaces forward on the board, and then follow the instructions on the space you land.

DEWEY: So if it's showing 1, you go forward 1. If it's showing 2, you go forward 2, etc.

SCROOGE: Are those the only possibilities?

LOUIE: Sure—it's a regular, 6-sided die. No funny business.

SCROOGE: OK, if you're so sure, how about we play a different game. We roll the die. If it comes up 1-6, then you each get a shiny nickel. But if it balances on an edge, then you give me all the money in your piggybanks, and you agree to sign over 50% of your earnings to me, for the rest of your life.

LOUIE: Let's play. Dice never land on their edges. It's a free nickel!

HUEY: I don't know Louie. You're probably right. But maybe not. I mean, it *might* land on an edge. I've never seen it happen, but that doesn't mean it's impossible.

DEWEY: Huey's right, we shouldn't play the game. We'd almost certainly win, but if we lost, that would be *terrible*!

How should we model the evolution of the common ground in the conversation between Huey, Dewey, and Louie? Given Huey's and Dewey's initial answers, and Louie's

response to Scrooge, it's natural to model them as presupposing that a rolled die will land with one of its six faces showing up. After Scrooge's offer, and in particular, after Huey's hesitation, it's natural to interpret them as no longer presupposing this, and instead as allowing for the possibility, however improbable, that a rolled die might balance on its edge. While the details differ, essentially all accounts of the function of epistemic modals—like the "might" in Huey's response to Louie—agree that one of their main roles is to expand the space of possibilities taken seriously in a conversation. When Huey says the die might land on its edge, and the others accept his claim, they stop taking for granted that dice don't land on their edges. Moreover, if the above conversation is a natural one, this doesn't need to happen in response to anything one might recognize as *evidence* that sometimes land on their edges.[33] Rather, collectively considering such a possibility is a natural response to being offered a bet in which, if the possibility turned out to be actual, the consequences would be dire.

The sensitivist's account of the evolution of doxastic possibilities for a single subject parallels the above remarks about the evolution of the context set in a conversation. And this is a welcome parallel, I suggest. While the example above was framed as a conversation, it could just as naturally have been framed as an internal monologue. At the very least, the non-sensitivist has a challenge here that the sensitivist does not. Assuming she accepts the standard account of how the context set evolves in conversations like the above one, she needs to explain why, even though it doesn't *seem* like much hangs on whether the example is framed as a collective discussion, or an internal monologue, in fact this distinction is of great theoretical significance. By contrast, the sensitivist has a natural explanation of *why* there isn't such a huge difference here. The same considerations that make it the case that, individually, we need to ignore possibilities in planning and deliberation (that's what belief is for), *also* make it the case that we need to do so

[33]Perhaps the fact that Scrooge is offering the bet at all could be taken as such evidence. But we could set up the case so that it's clear Scrooge has never seen a die before, and is willing to offer the bet merely because of how skewed the payoffs are, rather than because he has some inside information suggesting he'll win.

collectively, when engaged in collective planning and deliberation (that's why linguists and philosophers of language always appeal to non-trivial context sets—sets that rule out lots of possibilities—to explain linguistic phenomena). And just as we sometimes stop collectively ignoring possibilities when our practical circumstances change—a practice for which we have specialized language, like "might"—we do the same thing in private. So it's unsurprising that not much seems to hang on whether we think of the example as a conversation, or an internal monologue—whichever it is, we'll model it essentially the same way.

5

So far I've been arguing that if we start with a certain broadly anti-idealizing motivation for thinking that outright belief is important for creatures like us, we should also accept that believing a proposition outright involves having maximal credence in it. Someone sympathetic to the broadly fallibilist considerations mentioned at the beginning of this essay, however, might accept all that, and take it to show merely that outright belief is epistemically inappropriate. After all, even if it's true that I can't help but take for granted that I have hands,[34] doesn't reflection on skeptical arguments show that this is unreasonable, and that I should allow for the possibility that I don't? Perhaps doing without belief and the maximal credences it entails isn't a psychologically realistic option for limited creatures like us, but that doesn't excuse unjustifiable certainty; ought doesn't imply can in epistemology.

Such a view needn't entirely eschew epistemological distinctions among beliefs. Perhaps, starting from the position that ideal doxastic conduct involves having high but not maximal credences, we might develop a set of norms of the second-best for creatures

[34]Or better, as Hume famously suggests, I can stop myself from assuming this, but only for brief periods, when explicitly trying to. I can't stop myself from going back to assuming this once I leave the study.

like us. Even if the best policy for forming credences isn't on the table, we might still be able to make distinctions among the policies that are available to us, perhaps on the grounds that some of them leave our credences *closer* to the ideal than others.[35] It's a widely, though not universally held view in political philosophy that the correct theory of justice in an ideal society would shed a great deal of light on the appropriate norms for a non-ideal one. The idea is, roughly, that non-ideal societies are *more* just when they are *closer*, in some sense, to what would make for justice in an ideal society.[36] If an attractive, plausible set of norms for outright belief could be vindicated in a parallel way—as the best approximation to the ideal available to limited creatures like us—then we'd have an attractive unification of ideal and non-ideal theory in epistemology. Call this view—that ideal doxastic conduct involves eschewing belief, and that norms for outright belief are justified to the extent that they represent approximations to that ideal—BISB, pronounced "busy bee" for "**B**elief **I**s **S**econd **B**est".

BISB would constitute a significant retreat from the broadly fallibilist position I've been arguing against, according to which we never *in fact* are maximally confident in anything. The sorts of considerations stemming from high stakes bets mentioned earlier are typically taken to motivate both descriptive *and* normative conclusions—that is, they're taken to show both that one *shouldn't* be maximally confident in anything (because it would be irrational to bet on anything at arbitrarily unfavorable odds), *and* that almost nobody *is* maximally confident in anything (because nobody is willing to bet on anything at arbitrarily unfavorable odds). BISB amounts to abandoning the descriptive claim, while maintaining the normative one.

Nevertheless, even though it's a weaker claim than the one I've been arguing against, and not directly threatened by the considerations I've raised so far, in the remainder of this paper I'll try to provide some grounds for pessimism about BISB. It's not at all ob-

[35] **?**, §5 briefly considers this suggestion.

[36] E.g., **?** writes that "the reason for beginning with ideal theory is that it provides, I believe, the only basis for the systematic grasp of [pressing and urgent matters we are faced with in everyday life]".

vious that there *is* a coherent ideal of never assigning maximal credences to propositions, and if there is, we have very little sense of what it might look like.

I've already argued that non-ideal creatures like us sometimes update our bodies of belief by ruling out possibilities, and sometimes by allowing them in. This parallels what philosophers of language say about how context sets evolve—sometimes possibilities are ruled out (e.g., when simple assertions are made and accepted), and sometimes they are allowed in (e.g., when certain epistemic modal claims are made and accepted). If BISB is right, then ideal doxastic conduct involves neither sort of update. It doesn't involve ruling out possibilities for obvious reasons—that's what outright belief is. And it also can't involve allowing new possibilities in, because if it did, that would mean that such possibilities were previously ruled out. Both of these features of BISB creature serious difficulties.

If ideal doxastic conduct never involves taking account of new possibilities—expanding the space over which credences are defined—then whatever space ideal credences are defined over must have a certain sort of completeness property; it could never be (reasonably) expanded. Is it plausible that there *is* such a space? Much literature in Bayesian philosophy of science address the problem of how new theories can be confirmed by old evidence, and it's typically among the presuppositions of the discussion that sometimes there are new theories—i.e., theories that previously weren't even among the propositions over which scientists' credences were defined.[37] So one way of putting the present question is whether this problem wouldn't arise for epistemically ideal creatures; would they already have conceived of all possible theories? If so, then ideal doxastic conduct needn't ever involve introducing new doxastic possibilities. But if not, then it must, and BISB is in trouble. Arguing for a "no" answer to this question is well beyond the scope of this paper. For now, I just want to show that the question touches on some deep and thorny issues in the metaphysics and epistemology of modality, and so to the extent that

[37]See, e.g., **?**.

the defender of BISB is must give a "yes" answer, that is a substantial commitment that requires serious defense.

One reason to be skeptical is that it's not clear that there even is such a set as the set of all possible theories. The notion of a theory or hypothesis might be indefinitely extensible, in a sense analogous to what some philosophers of mathematics have claimed on behalf of the notion of number, or set. Maybe, for any set of theories, there is a larger one. Or for any set of distinctions among ways the world might be, there is a more fine-grained set of distinctions. There has been a great deal of debate in the philosophy of logic and mathematics over whether the notion of absolutely unrestricted quantification—quantification over absolutely everything—makes sense.[38] And at least some of the reasons for thinking it doesn't seem like they would carry over to provide reasons for doubt about the existence of a set not of all things, but of all theories. E.g., Agustin Rayo writes that absolutely unrestricted quantification "would require a final answer to what counts as a possible system of compositional representation. And I see no *prima facie* reason to think that our notion of representation (and our notion of linguistic representation, in particular) are constrained enough for this question to have a definite answer." (**?**, p. 29) If he's right, then it's similarly plausible that there's no definite answer to what counts as a theory—for any system of representation that allows one to generate a space of theories over which credences might be defined, there would be a richer system of representation that would allow for distinctions and possibilities not recognized by the previous one.

I haven't said anything to motivate the idea that the notion of a theory *is* indefinitely extensible in this way. My aim in raising these points is just to show that the defender of BISB takes on a serious commitment—if there is to be a coherent ideal of doxastic practice in which the space of possibilities over which credences are defined is never expanded, then the relevant notion of a possibility had better *not* be indefinitely exten-

[38]See **?** for a volume on the topic.

sible. There had better be a fixed set, never to be expanded, of doxastic possibilities over which ideal subjects' credences are defined. But this is a substantial commitment, and the broadly fallibilist considerations that can be mustered in favor of BISB don't seem to do much to motivate it.

Of course, if the commitments here were symmetrical—if the opponent of BISB were just as committed to *affirming* that the notion of a theory is indefinitely extensible as the defender of BISB is to denying it—then the preceding considerations might not seem to do much to motivate rejecting BISB. At best, they would show that the plausibility of BISB, somewhat surprisingly, is closely tied to debates in the philosophy of logic. But I think that's the wrong conclusion to draw. Consider analogous debates over the relationship between ideal and non-ideal theory in political philosophy. Those who hold that ideal theory has been overemphasized don't think that ideally just societies are *impossible*—only that we needn't get clear about what a perfectly just society would look like in order to ask questions about what would make for more justice in actual, flawed societies.[39] Likewise, the opponent of BISB needn't be committed to thinking that the ideal, belief-less creatures imagined by the proponent of BISB are impossible, only that we don't need to get clear on whether such creatures are possible (and if so, what their doxastic conduct is like) before asking what would make for better epistemic practice among believers like us. And to the extent that it is a difficult question whether such creatures are possible—perhaps because it turns on contested issues in the philosophy of logic—all the more reason not to hold our epistemological principles hostage to the answer.

For BISB to be tenable, we must be able to draw a relatively sharp distinction between epistemological questions about the shape of modal space—how we can know what is in principle possible, and what is not—and epistemological questions about how we can locate our own world in that space. This is because, for BISB to work, questions

[39]See, e.g., **?**, **?**.

about the shape of modal space must be capable of being settled through the exercise of ideal cognition, in advance of doing any serious empirical inquiry. That way, ideal doxastic practice never need involve learning things—updating our beliefs—in ways that lead to our newly ruling possibilities out, or in. Rather, ideal doxastic practice can involve merely changing one's probability distribution over a fixed space of possibilities. But if questions about the space of possibilities over which credences should be defined are themselves the sorts of questions that can be asked and answered in the course of ideal inquiry—if, e.g., ideal inquiry can involve learning that some scenario is possible, and should be ruled in, or impossible, and should be ruled out—then BISB looks to be in trouble.[40,41]

Again, I point to this debate not as an example of an area where BISB has intrinsically implausible commitments. I myself am quite ambivalent about the epistemology of modality. But I do think it's implausible that epistemology of modality is "first" epistemology, in the sense that we can only get on to the rest of the epistemological questions once the epistemology of modality is settled. BISB, though, has something very close to that commitment; if the proper methodology for discovering norms of belief is to try to see how well they approximate a certain sort of ideal, then we need a reasonably detailed picture of what that ideal is like in order to make progress in non-ideal epistemology. Otherwise, we won't be able to evaluate norms of belief, since we won't know what to compare them to. To continue the parallel with political philosophy, an analogous methodological assumption has motivated Rawlsian, "ideal theory first" approaches—the assumption that we need a reasonably detailed theory of what

[40]See, e.g., **?**, **?**.

[41]I don't mean that there aren't ways that a defender of BISB could allow for our learning that scenarios are possible or impossible in *some* sense, while still maintaining her view. E.g., The two-dimensional framework defended by **?** and **?** allows that we can learn which propositions are metaphysically possible or necessary in a certain sense (a "secondary") sense, while still holding that ideal cognition involves shifting one's credences over a fixed space of *ideally epistemically possible* worlds. For some evidence that this *is* Chalmers' view—not just that there is a fixed space of ideally epistemically possible worlds, which he defends in numerous places, but also that ideal cognition never rules out any such worlds as impossible, see **?**, and in particular the discussion of "insulated idealization".

the ideally just society looks like in order to make progress on questions of justice in actual societies. So if BISB is right, then questions about the epistemology of modality are utterly central to epistemology more broadly. In fact, only given certain answers to such questions—the answers on which the belief-less ideal is coherent—can the rest of epistemology get off the ground. And that commitment is hard to swallow—it seems to me that we can ask epistemological questions about how creatures like us should go about our doxastic business while remaining agnostic about whether various of our limitations could, in principle, be overcome by creatures radically different from us. Even if we think such creatures could exist, I doubt we need a detailed picture of what their doxastic conduct would be like before we can ask epistemological questions about ourselves.

6

In this paper I've defended the simple view that outright belief is an endpoint of the scale of degreed belief—that we are maximally confident in what we (outright) believe. I've also argued that the "sensitivist" commitmenst necessary to defend this position from otherwise devastating counterexamples can be independently motivated in a variety of ways. The view that outright belief is of epistemological significance, in large part because of the role it plays for cognitively limited creatures like us, is increasingly popular. If I'm right, then the same motivations for taking belief seriously at all are *also*, ultimately, motivations for accepting the simple view about its relationship to credence.

Bibliography

Buchak, Lara. 2013. "Belief, Credence, and Norms." *Philosophical Studies* pages 1–27.

Chalmers, David, 2012. *Constructing the World*. Oxford University Press.

Chalmers, David J., 1996. *The Conscious Mind*. Oxford: Oxford University Press.

Christensen, David, 2004. *Putting Logic in its Place*. Oxford: Oxford University Press.

Clarke, Roger. 2013. "Belief Is Credence One (In Context)." *Philosophers' Imprint* 13:1–18.

———. 2015. "Preface Writers Are Consistent." *Pacific Philosophical Quarterly* 96.

Craig, Edward, 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.

Earman, John. 1989. "Old Evidence, New Theories: Two Unresolved Problems in Bayesian Confirmation Theory." *Pacific Philosophical Quarterly* 70:323–340.

Egan, Andy. 2008. "Seeing and Believing: Perception, Belief Formation and the Divided Mind." *Philosophical Studies* 140:47 – 63.

Elga, Adam, and Agustín Rayo. Unpublished Manuscript. "Fragmentation and Information Access." .

Gibbons, John, 2013. *The Norm of Belief*. OUP Oxford.

Greco, Daniel, 2013. "Probability and Prodigality." In Gendler, Tamar Szabo, and John Hawthorne, editors, "Oxford Studies in Epistemology, Volume IV," Oxford University Press.

———. 2014. "Iteration and Fragmentation." *Philosophy and Phenomenological Research* 88.

———. Forthcoming. "Cognitive Mobile Homes." *Mind* .

Greco, Daniel, and Brian Hedden. Forthcoming. "Uniqueness and Metaepistemoloy." *Journal of Philosophy* .

Hansson, Sven Ove, 2011. "Logic of Belief Revision." In "Stanford Enyclopedia of Philosophy," http://plato.stanford.edu/entries/ogic-belief-revision/.

Harman, Gilbert, 1986. *Change in View: Principles of Reasoning.* Cambridge, MA: The MIT Press.

Hawthorne, John, and Jason Stanley. 2008. "Knowledge and Action." *Journal of Philosophy* 105:571–90.

Holton, Richard, forthcoming. "Intention as a Model for Belief." In Vargas, Manuel, and Gideon Yaffe, editors, "Rational and Social Agency: Essays on the Philosophy of Michael Bratman," Oxford University Press.

Jackson, Frank, 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis.* Oxford University Press.

Jeffrey, Richard C., 1970. "Dracula Meets Wolfman: Acceptance Vs. Partial Belief." In Swain, Marshall, editor, "Induction, Acceptance, and Rational Belief," volume 26, pages 157–185. D. Reidel.

Leitgeb, Hannes. 2014. "The Stability Theory of Belief." *Philosophical Review* 123:131–171.

Levi, Isaac, 1974. *Gambling with Truth: An Essay on Induction and the Aims of Science.* The MIT Press.

———, 1980. *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance.* The Mit Press.

Lewis, David. 1982. "Logic for Equivocators." *Noûs* 16:431–441.

Marley-Payne, Jack. Unpublished Manuscript. "Task-Indexed Belief." .

Mills, Charles W. 2005. ""Ideal Theory" as Ideology." *Hypatia* 20:165–184.

Moss, Sarah. 2012. "Updating as Communication." *Philosophy and Phenomenological Research* 85:225–248.

———. 2013. "Epistemology Formalized." *Philosophical Review* 122:1–43.

———, Forthcoming. *Probabilistic Knowledge*. Oxford University Press.

Norby, Aaron. 2015. "Uncertainty Without All the Doubt." *Mind and Language* 30:70–94.

O'Neill, Onora. 1987. "Abstraction, Idealization and Ideology in Ethics." *Royal Institute of Philosophy Lectures* 22:55–69.

Rawls, John, 1971. *Theory of Justice*. Cambridge Massachussetts: Belknap Press.

Rayo, Agustín, 2013. *The Construction of Logical Space*. Oxford University Press.

Ross, Jacob, and Mark Schroeder. 2014. "Belief, Credence, and Pragmatic Encroachment." *Philosophy and Phenomenological Research* 88:259–288.

Staffel, Julia. Forthcoming. "Beliefs, Buses and Lotteries: Why rational belief can't be stably high credence." *Philosophical Studies* .

Stalnaker, Robert. 1978. "Assertion." *Syntax and Semantics (New York Academic Press)* 9:315–332.

———, 1984. *Inquiry*. Cambridge, MA: The MIT Press.

———. 2006. "On Logics of Knowledge and Belief." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 128:169–199.

———, 2011. *Mere Possibilities: Metaphysical Foundations of Modal Semantics*. Princeton University Press.

———, 2014. *Context*. OUP Oxford.

Stanley, Jason, 2011. *Know How*. Oxford University Press.

Sutton, Jonathan. 2005. "Stick To What You Know." *Nous* 39:359–396.

Swanson, Eric, 2009. "How Not to Theorize About the Language of Subjective Uncertainty." In Egan, Andy, and B. Weatherson, editors, "Epistemic Modality," Oxford University Press.

———. forthcoming. "The Application of Constraint Semantics to the Language of Subjective Uncertainty." *Journal of Philosophical Logic* pages 1–26.

Urmson, J. O. 1952. "Parenthetical Verbs." *Mind* 61:480–496.

Uzquiano, Gabriel, and Agustin Rayo, 2006. *Absolute Generality*. Oxford University Press.

Weatherson, Brian. 2014. "Games, Beliefs and Credences." *Philosophy and Phenomenological Research* 88.

Wedgwood, Ralph. 2012. "Outright Belief." *Dialectica* 66:309–329.

Williamson, Timothy, 2000. *Knowledge and its Limits*. Oxford University Press.

Yalcin, Seth. 2007. "Epistemic Modals." *Mind* 116:983–1026.

———. 2012. "Bayesian Expressivism." *Proceedings of the Aristotelian Society* 112:123–160.