

Belief without credence

J. Adam Carter¹ · Benjamin W. Jarvis² ·
Katherine Rubin³

Received: 14 May 2014 / Accepted: 31 July 2015 / Published online: 20 January 2016
© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract One of the deepest ideological divides in contemporary epistemology concerns the relative importance of belief versus credence. A prominent consideration in favor of credence-based epistemology is the ease with which it appears to account for rational action. In contrast, cases with risky payoff structures threaten to break the link between rational belief and rational action. This threat poses a challenge to traditional epistemology, which maintains the theoretical prominence of belief. The core problem, we suggest, is that belief may not be enough to register all aspects of a subject's epistemic position with respect to any given proposition. We claim this problem can be solved by introducing other doxastic attitudes—genuine *representations*—that differ in strength from belief. The resulting alternative picture, a kind of doxastic states *pluralism*, retains the central features of traditional epistemology—most saliently, an emphasis on truth as a kind of objective accuracy—while adequately accounting for rational action.

Keywords Belief · Rational action · Cognitive achievement · Contextualism · Truth · Virtue epistemology

✉ J. Adam Carter
j.adam.carter@ed.ac.uk
Benjamin W. Jarvis
benjamin_jarvis@alumni.brown.edu
Katherine Rubin
katherine.rubin@gmail.com

¹ University of Edinburgh, Edinburgh, UK

² New York University, Stern School of Business, New York, NY, USA

³ Weill Cornell Medical College, New York, NY, USA

Introduction

Consider the following case:

CAR INSURANCE: Your car insurance company offers you a deal. The offer is to insure your car for the next five years against theft for the nominal price of ten dollars. You do not live in circumstances where car theft is particularly prevalent. It is rational for you to believe that your car will not be stolen during the next five years. But, of course, car theft has been known to happen on rare occasion nearby. And, you do not consider yourself to be especially different from the owners of those stolen cars.

Cases like CAR INSURANCE appear to diminish the importance of rational belief, and this is because it seems you are rational to act in opposition to what you rationally believe is the all-things-considered best thing to do. After all, you believe that your car will not be stolen and also that, if your car will not be stolen, the insurance won't be needed; and further, you believe that it would be best not to spend money on insurance if it isn't needed. Accordingly, you believe that it would be best not to spend money on car insurance. Nevertheless, you rationally purchase the insurance.¹ This kind of case appears to show that the link between rational belief and rational action is broken; one cannot straightforwardly decide what to do on the basis of what one (even rationally) believes to be best.²

In contrast, rational credence appears to do much better here. Of course, you may rationally believe that your car will not be stolen—and thus believe that if you purchase the insurance, the end result will simply be less money for other things. However, if you are rational, you have at least some minimal credence in the proposition that your car will be stolen. The rationality of this minimal credence (along with the rationality of your preferences) can explain why it's rational to purchase the insurance.³

Cases such as CAR INSURANCE might seem to suggest that what is fundamental in epistemology is credence (alternatively known as “degrees of belief”) rather than belief. Perhaps the epistemology of belief might be important for reasons that aren't directly related to action, but a divorce between rational belief and rational action at least greatly diminishes the importance of having the capacity to believe rationally. The project of this paper is to explain why this threat to traditional epistemology—epistemology that focuses primarily on belief—may not be as severe as it first appears. The advantage that a credence framework has in accounting for rational action might

¹ Thanks to an anonymous referee at *Synthese* for helpful suggestions regarding the presentation of this case.

² We ignore here the possibility that one has beliefs about probabilities, e.g. a rational belief that there is a very small chance the car will be stolen. Cf. Stanley (2005, p. 10). This way of resolving the problem seems to succumb to an animals and small-children-style objection, but a more fundamental problem for this line of thought, we think, is the wrong content problem we raise at the end of Sect. 4. We thank an anonymous referee at *Synthese* for raising this line of discussion.

³ Cf. Jeffrey (1990).

well be secured by dispensing with credence in favor of other doxastic attitudes that are more similar to belief by being *representational*.⁴

Moreover, we suggest that this alternative framework can still allocate a special place to belief. Ultimately, we will characterize belief in terms of acceptance. Acceptance is the mental state of taking some propositional content $\langle p \rangle$ for granted—whether consciously or not—in practical reasoning and rational decision-making; acceptance ordinarily causes acting as if $\langle p \rangle$ is true.⁵ Acceptance of $\langle p \rangle$ involves using $\langle p \rangle$ as a working hypothesis. Acceptance of a propositional content is sometimes rational and sometimes not, depending on the strength of one’s epistemic position and the practical stakes. (As an important aside, we note that “epistemic position” is intended throughout to be a neutral way of talking about something like the subject’s evidential condition as determined by her total body of evidence without raising any question about what evidence is. Thus, in particular, “epistemic position” is not a way of referring to the subject’s doxastic states either taken collectively or individually.) Our proposal will be that belief is the weakest doxastic attitude that normally suffices for rational acceptance. So, it can make sense for a cognizer to traffic largely in beliefs, considering whether to take up stronger or weaker doxastic attitudes only when practical stakes are sufficiently high or low.⁶ Moreover, we will suggest that whether to take up stronger or weaker doxastic attitudes can, in a wide range of cases, be determined by what one believes. So, agents can regularly make decisions that are rational across a wide range of circumstances on the basis of rationally held beliefs alone.

Our project is significant because one of the deepest ideological divides within epistemology concerns the relative importance of belief versus credence. The moral of this paper is that one may have to look beyond decision-making and action to settle that debate. For instance, it may be more productive to consider Harman’s (1986) contention that understanding human reasoning requires a belief-based framework rather than a credence-based one.⁷ Alternatively, one might consider whether or not the basic doxastic elements should be representations in the sense highlighted in Sect. 4.

Before continuing, however, it’s worth clarifying what *isn’t* at issue in this paper. The credence-based framework is closely associated with probabilism, the view that, insofar as a subject’s assignment of credence is rational, it behaves like an assignment of probability. However, something very much like probabilism is also accepted among some traditional epistemologists: it is sometimes accepted that beliefs can be more or less rational, and moreover, that levels of rationality function like levels of likelihood, which should be understood as a kind of probability.⁸ Clearly, then, it cannot be a requirement that defenders of belief-based epistemology do without anything like

⁴ To be clear, we mean here ‘representational’ in the sense of definitively rendering the world in a particular way rather than having propositional content so as to be about entities in the world.

⁵ The distinction between acceptance and acting as if is clearly drawn in Ross and Schroeder (2012). On acceptance, see also Bratman (1992) and Weatherson (2005).

⁶ See Sect. 5 for more development of this element of the proposal.

⁷ See Broome (2013) and especially Staffel (2012) for some discussion.

⁸ See Smith (2010).

epistemic probability. We will discuss this theoretical possibility in Sect. 6, but strictly speaking, accepting it is not part of the belief-based picture. What's primarily at issue is not the viability of epistemic probability, but the extent to which epistemic statuses of credence are more fundamental than epistemic statuses of belief.

§1

Both belief and credal states qualify as doxastic attitudes, broadly construed. But, belief is an “on-or-off” attitude; believing a proposition stands in opposition to withholding belief with respect to that proposition. By contrast, credence comes in degrees; arguably, each of a continuum of opposing credal states is possible in principle (even if not in practice for limited human beings). Perhaps more importantly, the correctness of beliefs—their accuracy—is an all or nothing affair. A belief is correct if it is true, and incorrect if false. In contrast, to the extent that accuracy makes sense for credal states, it is a matter of degree.⁹ In taking up belief towards a proposition, the aim is to possess the truth (by believing it),¹⁰ yet all but the most extreme credal states do not take a definitive position per se on what the truth of a given matter is. For any proposition $\langle p \rangle$, credences seem to split the difference (with a particular weighting) between opposing positions on whether $\langle p \rangle$. For instance, if one holds .3 credence towards the proposition that it will rain today, one's credal state is not incorrect in any obvious sense if the proposition turns out to be false, i.e. it does not, in fact, rain. However, there is no obvious reason why this would change if one's credence is, instead, .9 (unless .9 credence is taken to additionally constitute belief or some similar “on-or-off” attitude)¹¹ since there is no *principled* difference between .3 and .9, only a *difference of degree*.

It is fairly obvious that there must be certain relationships between the respective *epistemologies* of beliefs and credal states.¹² For instance, it is clear that it can't simultaneously be rational to believe $\langle p \rangle$ while it is also rational to hold the minimum credence in $\langle p \rangle$. Nevertheless, the relationship between these epistemologies is not altogether clear. Echoing Weatherson (2005),¹³ we might ask:

⁹ See Joyce (1998). For further discussion of Bayesian attempts to locate what would play the role of correctness conditions for credal states, see also Hajek (unpublished). Unlike van Fraassen (1983, 1984) and Lange (1999) who appeal here to calibration, Hajek takes it that belief is to truth as credal states are to agreement with objective chance.

¹⁰ For some discussion, see Dummett (1959), Wright (1992), Engel (1998), Velleman (2000), Wedgwood (2002, 2007), Shah (2003), Shah and Velleman (2005), Gibbard (2005), Steglich-Peterson (2006), Lynch (2009), Whiting (2010), and Jarvis (2012).

¹¹ Moreover, it seems difficult to understand how credal states could straightforwardly constitute on-off belief given that the latter seems to be incorrect when the propositional content is false while the former does not.

¹² Cf. Hajek and Eriksson (2007) for a survey of different positions on the nature of credal states. Cf. Chalmers (2011) for a discussion of the nature of the objects of credal states.

¹³ In a parallel passage, Weatherson (2005), in fact, asks these questions about degrees of belief and belief *tout court* rather than their respective epistemologies. Thanks to an anonymous referee for asking us to clarify.

Do we really have two subject matters here (epistemology of degrees of belief and epistemology of belief *tout court*) or two descriptions of the one subject matter? If just one subject matter, what relationship is there between the two modes of description of this subject matter?

These questions concern the extent to which the epistemic statuses of belief *settle* the epistemic statuses of credence—and vice-versa. The concern is whether the complete account of the epistemic statuses of beliefs *leaves a remainder* for the complete account of the epistemic statuses of credences—and vice-versa. There are, of course, metaphysical and epistemic interpretations of settling and remainder—corresponding, we think, to each of the two questions. The epistemic reading (the second question) might be stated (roughly) in the terms of Chalmers (2012): to what extent are the epistemic statuses of credence *scrutable* from those of belief—and vice-versa? The metaphysical reading (the first question) would be: to what extent do the epistemic statuses of credence metaphysically supervene on those of belief—and vice-versa? (The orthodox position is that while scrutability entails metaphysical supervenience, the converse may not be true, so these questions are genuinely distinct.)¹⁴

In addition, however, there is a *further* question about whether there is any explanatory order to the epistemic statuses of belief and credence, respectively. Could it be that when it isn't rational to believe $\langle p \rangle$, but rational to hold the minimum credence in $\langle p \rangle$, that the former explains the latter? There are three available positions on the fundamentality of belief-based versus credence-based epistemology:

Belief Fundamentalist Epistemology (BFE): The epistemic statuses of belief are more fundamental.

Credence Fundamentalist Epistemology (CFE): The epistemic statuses of credence are more fundamental.

Anti-Fundamentalist Epistemology (AFE): Neither kind of epistemic status is more fundamental.

AFE implies that, when it comes to the epistemic statuses of belief and credence, there is no (asymmetrical) constitutive dependence of one on the other. Consider, for instance, the status of (epistemic) rationality. The idea is that it is not the case that a belief is rational because a certain kind of credal state is rational, nor is a credal state rational because some kind of belief is rational. Matters of rationality for belief and credence might be interrelated, but one does not participate in (metaphysically) grounding the other. The degree of credence one is rationally committed to having in some proposition $\langle p \rangle$ is not, for instance, constitutively a matter of how easy it would be to come to have a rational belief with propositional content $\langle p \rangle$ even if the former has consequences for the latter and vice-versa.¹⁵

¹⁴ For some relevant discussion, see not only Chalmers (2012), but Ichikawa & Jarvis (2013).

¹⁵ Cf. Harman (1986, p. 22) on degrees of belief: “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....”

BFE and CFE both deny AFE. Consider first BFE. *Strong* BFE insists that the epistemic statuses of credence are *wholly* determined and constitutively explained by the epistemic statuses of belief (so there is no remainder). A non-trivial example might be the Williamsonian suggestion that rational credence is probability conditional on (the certainty) of those beliefs that qualify as knowledge.¹⁶ However, BFE itself insists only that the epistemic statuses of credence are *at least partly* grounded in those of belief while admitting that there may be other factors as well. (Of course, this “partial grounding” could not be reciprocal, but rather must be asymmetrical.)

Analogously, *strong* CFE insists that the epistemic statuses of belief are *wholly* determined and constitutively explained by the epistemic statuses of credence (so, again, there is no remainder). Consider this bi-conditional: a belief $\langle p \rangle$ is rational if and only if credence in $\langle p \rangle$ above some given threshold is rational.¹⁷ If the left-hand side is understood as the analysandum and the right-hand side is understood as the analysans, then what we have is an example of strong CFE. However, as with BFE, CFE need not be strong. A non-strong version of CFE might suggest that rational credence above a given threshold is a necessary precondition for rational belief, but that other factors unrelated to credence contribute to constituting the rationality of belief as well.

All three positions—BFE, CFE, and AFE—have an air of plausibility to them. Wedgwood (2012) notes that a spectrum of positions have been taken on the relationship between belief and credence.¹⁸ Presumably, the spectrum looks very similar when it comes to the relationship between the *epistemologies* of belief and credence. Our interest in this paper is this relationship between the epistemologies (which isn’t necessarily settled by the relationship between the psychologies of belief and credence). Nevertheless, we want to be very clear that our intention in this paper is not to rule out any of BFE, CFE, and AFE. Rather, we wish to show that a certain prominent consideration in favor of CFE—the one we will explore in the next section—can be resisted. This will open the door to epistemological frameworks that emphasize the importance of doxastic representations, i.e. BFE and AFE. But, it is well beyond the scope of this paper to give a positive argument for any such framework.

§2

As was noted in the introduction, CFE appears to gain support from cases such as CAR INSURANCE. Because in that case it is rational for you to believe that your car will not be stolen during the next five years, it will ordinarily be rational for you to believe that you will be financially worse off by accepting the offer to purchase insurance. Stipulate that CAR INSURANCE is an ordinary case. Stipulate also that you will gain no emotional comfort from having the insurance. A theft of the car would be bad, but it would not spell financial demise for you. Indeed, as far as this decision is concerned,

¹⁶ Williamson (1998, 2002).

¹⁷ This thesis is analogous to the Lockean theory of belief whereby believing is having credence above a certain threshold. For some relevant discussion here, see Christensen (2004), Sturgeon (2008), Foley (2009), and Fantl and McGrath (2009, Chap. 5).

¹⁸ Wedgwood (2012, p. 310).

the only considerations that are relevant for you are financial. Still, it could easily be rational for you to accept the offer to purchase this car insurance.

A credence-based epistemology can easily accommodate this result. It may be rational for you to believe that you will be financially worse off by accepting the offer. But, that need not imply that it would be rational to hold the maximum credence in this proposition. Indeed, if the former does imply the latter, then there may be bigger problems. It becomes more difficult to envision how rational belief is even possible—so skepticism threatens. Or if the former does not imply the latter, the relationship between rational action and credence becomes more obscure.¹⁹ If fairly ordinary beliefs are rational and rational belief entails rational maximum credence, then it could not plausibly be that rational maximum credence rationalizes staking one's life on the truth of the belief (because the rationality of ordinary beliefs certainly doesn't). In other words, we've moved quite far from even the spirit of Ramsey's suggestion that credences are closely associated with betting patterns.²⁰ Putting aside the theoretical option of an entailment between rational belief and rational maximum credence, it is plausible that although it is rational for you to believe that you would be financially worse off, it also could easily be rational to hold some non-minimal credence in the proposition that you would be financially better off by accepting the offer.²¹ And, this non-minimal credence could be above the threshold required to make it rational to accept the offer to purchase the car theft insurance for ten dollars.²²

Without the resources of a credence-based epistemology, it becomes more difficult to understand the rationality of accepting the offer. Ex hypothesi, what you believe is that you would be financially worse off by taking the offer, and no other non-financial considerations are relevant. So, how could it be rational to accept the offer?

The question is not merely how it could be rational for you to do what is, by your own beliefs, worse. Perhaps beliefs are simply not the kind of cognitive state to rely on in this situation. Using the terminology from the introduction, we might say, *pace Fantl and McGrath (2009)* and others, that the rationality of belief need not imply the rationality of acceptance.²³ Perhaps this is counterintuitive, but we can bite that bullet. Even so, a problem remains. How could it be rational for you to do what is, by your own beliefs, worse *unless* it is rational for you to have some further doxastic state that puts you in a position to appreciate that this action would be, in fact, rational?

¹⁹ See Kaplan's (2009) discussion of Williamson (2002) for some helpful discussion here.

²⁰ See Ramsey (1931), Skyrms (1980), and Christensen (2004, Chap. 5) for some discussion about the relationship between betting patterns and credence. For some discussion of how betting patterns and credence come apart, see Bradley and Leitgeb (2006). Relatedly, note that even though many deny a *metaphysical* connection between credences and betting patterns, it is much harder to deny a normative connection. See Christensen (2004). Thanks to an anonymous referee for encouraging this latter point.

²¹ It is noteworthy that Wedgwood (2012) responds to the dilemma introduced in this paragraph by postulating two kinds of credence: rational belief entails rational maximum practical credence, but not rational maximum theoretical credence.

²² Dutch-book arguments are used to show that making decisions in this kind of way makes it rational for credence to function as a kind of subjective probability—for discussion see, for example, Christensen (2004, Chap. 5) and Hajek (2008). But, as we pointed out in the introduction, probabilism is not really at issue in this paper.

²³ See also Engel (1998).

Presumably, it couldn't. So what are these further doxastic states that put you in a position to appreciate that this action would be, in fact, rational? Call this “the Further Doxastic State Question.”

One option for circumventing the Further Doxastic State Question is to deny that the description of CAR INSURANCE leaves much room for it to be rational to accept the offer. The description stipulates that it is rational for you believe that your car will not be stolen during the next five years. But then, one might say, it could not but be rational for you to reject the offer to insure for car theft given that financial considerations are all that matters. So, if it is ordinarily rational to accept this kind of offer (i.e. not in CAR INSURANCE, but in less unusual cases with a similar description), then it is only because, in these same ordinary circumstances, it is not rational to believe that your car will not be stolen during the next five years. This might be because, as a general rule, it is very difficult to be in a strong enough epistemic position to rationally have this belief. But then we're trending towards skepticism again—an unpalatable result.²⁴

Alternatively, an option would be to say the following: that even though it wouldn't ordinarily be difficult to be in a strong enough epistemic position rationally to have this belief, being in a “practical environment”²⁵ where this offer is available makes it difficult. The latter approach involves thinking that rationality for belief is subject to pragmatic encroachment: pragmatic factors about what courses of action play a role in determining the rationality of belief.²⁶

However, pragmatic encroachment creates problems for the rational governance of belief.²⁷ Rational governance of belief implies the agent believes, by and large, because it is rational to believe. Pragmatic encroachment implies that whether it is rational to believe depends on strength of epistemic position and the relevant pragmatic factors. So to believe, by and large, when it is rational and not, by and large, otherwise, one will have to be able to track (fallibly) the strength of one's epistemic position and the relevant pragmatic factors. How does one track (even fallibly) the strength of one's epistemic position and the relevant pragmatic factors? Call this “the Tracking Question.”

Again, credence-based epistemology provides a straightforward answer. One can track strength of one's epistemic position with regard to any proposition $\langle p \rangle$ with more fine-grained credal states. And, one can track pragmatic factors using one's credal states and preferences. The relative practical merits of acceptance versus non-acceptance of a proposition $\langle p \rangle$ can be measured by something like their relative expected values (as determined by credences and preferences). If belief rationally requires acceptance, then we can measure the practical merits of belief versus non-belief in the same kind of way.²⁸

²⁴ For some interesting and relevant discussion on this point, see [Greco \(2012\)](#).

²⁵ [Hawthorne \(2004\)](#).

²⁶ What we've rehearsed here is, more or less, Fantl and McGrath's (2009) argument from fallibilism to pragmatic encroachment. On pragmatic encroachment, see also [Hawthorne \(2004\)](#), [Stanley \(2005\)](#), [Ganson \(2008\)](#), [Hawthorne and Stanley \(2008\)](#), and [Grimm \(2011\)](#).

²⁷ See [Ichikawa et al. \(2012\)](#). [Wedgwood \(2012, §6\)](#) addresses a similar issue.

²⁸ Cf. [Weatherson \(2005\)](#) and [Ross and Schroeder \(2012\)](#).

The emerging picture is one where, in one way or another, credence looks to be more theoretically interesting than belief. First consider the purely intellectualist approach where pragmatic factors don't matter for epistemic evaluation. On this approach, we run into the Further Doxastic State Question (so long, at least, as we steer clear of skepticism). We could easily answer this question by pointing to credence as the further doxastic state, but only by acknowledging a more fundamental relationship between rational credence and rational action than there is between rational belief and rational action. This may well obviate the need for belief in the cognitive system—at least in principle. At least as far as rational action is concerned, belief would turn out to be essential only if it is simply constituted by credence in some sort of way, e.g. belief is a matter of having credence above a certain threshold. When it comes to rational action, only a modest role for belief is apparently available, if any. In practice, it may be useful to have beliefs, because dealing with credence is cognitively taxing and we are rationally limited.²⁹ But, ideally, it would be better to make decisions using credences. And, one has to move from belief-based practical reasoning to credence-based practical reasoning whenever one perceives sufficient risk.

On a pragmatic encroachment approach, we run into the Tracking Question (so long as we continue to steer clear of skepticism). We could easily answer this question by pointing to credence as the tracker, but then one is rationally governing one's believing by having credal states. But why bother to manage one's beliefs by having credal states?³⁰ Why not just manage one's actions directly instead? Again, credence seems to obviate the need for belief in the cognitive system in principle (unless credence simply constitutes belief); and, it may obviate the need for belief in practice as well.

On both approaches, rational credence gives us a way of answering a pressing question. But, if rational credence is the way that we *must* answer these pressing questions, then BFE looks dubious. If we have to turn to rational credence to solve problems that arise from casting our epistemology solely in terms of belief, then it just looks like a credence-based epistemology has more resources than a belief-based epistemology in the sense that a belief-based epistemology simply cannot account for the same range of phenomena—in particular, either rational action (if we keep the relationship between rational belief and rational action loose)³¹ or rational belief governance (if we keep it tight).³² This would be a good reason (even if not an especially conclusive reason) to think that BFE is false. Moreover, CFE (but not necessarily strong CFE!) would be hard to avoid unless there is some other range of phenomena that a belief-based epistemology can account for but a credence-based epistemology cannot. (*Remember*: CFE merely entails the explanatory priority of credence-based epistemology; only strong CFE entails that credence-based epistemology settles belief-based epistemology without remainder.)

Of course, it is not our contention that either the Further Doxastic State Question or the Tracking Question have to be answered by pointing to credence; our aim is

²⁹ Hill and Schechter (2007).

³⁰ Cf. Christensen's (2004) discussion in §4.4 of Kaplan's (1996) theory of belief.

³¹ Cf. Hill and Schechter (2007) and Reed (2010).

³² See references in footnote 28.

to pursue a way for resisting this idea. Because we are particularly concerned by the Tracking Question,³³ we prefer to take an intellectualist approach (rather than opt for pragmatic encroachment). So, our interest here will be in resisting CFE by finding an answer to the Further Doxastic State Question. It is not our intention, however, to take a stand here about other ways that an advocate for the theoretical prominence of belief might respond to this dilemma. For all we say, there may be other promising avenues to pursue.

By way of reminder, the Further Doxastic State Question is “What are those further doxastic states that put you in a position to appreciate that some possible action would be, in fact, rational in cases like CAR INSURANCE in which, by your own beliefs, this action is worse?” One possible answer to this question is these further doxastic states are, in fact, beliefs.³⁴ So, while by your own beliefs, this action is worse, still by some other route, your beliefs let you appreciate that this action would be, in fact, rational. If it could be made to work, this kind of answer would definitely help the advocate of BFE.

This is not the route that we intend to pursue most immediately. In Sect. 7, we do propose that, in practice, rational action may be possible, by and large, for an agent trafficking only in beliefs precisely because an agent’s whole network of beliefs may register a lot about how secure any particular belief is in light of the agent’s experiences. Nevertheless, we think that a straightforward “further beliefs” answer to the Further Doxastic State Question confronts a sort of *wrong content problem*. In CAR INSURANCE, for instance, the relevant question seems to be whether your car will be stolen in the near future—something you believe to be false. Presumably, you have beliefs that are inferentially connected to this belief, e.g. that you live in a safe neighborhood. You may also have beliefs about how strong or weak your epistemic position is with respect to the proposition that your car will be stolen, e.g. the epistemic probability of theft. But, whether you will receive a benefit or encounter a loss from the course of action in question—purchasing insurance—doesn’t precisely turn on the truth or falsehood of most of these further beliefs. It turns precisely on whether or not your car will be stolen. (Assuming there is no emotional security gained from the purchase, purchasing insurance will benefit you if your car is stolen and otherwise be a loss.) So, many of these further beliefs aren’t directly related enough to the potential reason-giving facts for or against purchasing insurance. We conclude that many of these further beliefs cannot ultimately be the rationalizers for action in this case; they have the *wrong content* to do so. (Of course, the belief that your car will not be stolen has the right content, but the rational belief isn’t “strong enough” to rationalize forgoing insurance; in this instance, the problem appears to be not with the content, but the *attitude*.) In particular, beliefs about one’s epistemic position have the wrong content because although one’s epistemic position is relevant to whether one acts rationally, the aim of acting is not to act rationally per se, but to act to secure benefits, diminish losses, right wrongs, uphold justice, etc.—that is to say, to make appropriate changes in the world. Beliefs with the right content to rationalize action

³³ See Ichikawa et al. (2012).

³⁴ See, for example, Ichikawa (2012).

are precisely those directed outwardly at the sources of benefit, loss, right, and wrong in the world, not directed in a self-fascinated way at factors relevant to the subject's own rationality or evidence. To rationalize action, the beliefs must *be* rational and supported by evidence, not *be about* rationality or evidential support.

In addition, we suspect that while one's network of beliefs may register *a lot* about how secure any particular belief is in light of the agent's experiences, it doesn't register *everything*. Accounting for *maximally* rational action really does cause problems for a strong version of BFE that insists that the epistemic statuses of credence are wholly determined and constitutively explained by those of belief precisely because even an entire network of beliefs won't generally reflect all aspects of one's evidence that are relevant to selecting a course of action that is maximally rational. To adequately account for how one might, in principle, act maximally rationally in all possible cases, one needs to move away from a purely belief-based epistemology to an epistemology of a more general kind. This will be our course.

§3

Let's begin looking for an answer to the Further Doxastic State Question by thinking carefully about why a credence-based epistemology seems to handle CAR INSURANCE with such ease. Why doesn't the Further Doxastic State Question seem to arise for credences like it does for beliefs? For instance, why aren't there obvious cases in which questions about whether it is rational to have some credence are settled, but these answers conflict with answers to questions about which potential course of action is rational? We contend that the difference here is the fact that there is a continuum of credal states, whereas with belief, it seems like there's only belief and withholding belief.

It is not hard to imagine variants of CAR INSURANCE in which your epistemic position vis-à-vis the proposition that your car will not be stolen is much stronger. For instance, you might know that the area that you live and work in is renowned for its exceedingly low crime rate and very effective law enforcement. And, we can imagine variants in which your epistemic position is much weaker, e.g. in contrast to the original case, you might know that car theft is prevalent in your area, particularly theft of the brand of car you drive. But, these differences in epistemic strength can't be captured by a (first-order)³⁵ doxastic attitude towards the proposition if the only attitudes to choose from are believing and withholding. There are two attitudes, but at least three different epistemic positions. So, one attitude will have to double as the response to two different epistemic positions. Information about which of the two epistemic positions you are in will not be registered by this attitude. Merely adding another doxastic attitude won't help. Consider any two variants of our original case, CAR INSURANCE^X and CAR INSURANCE^Y that are strictly ordered by epistemic position with respect to the proposition in question. Mightn't we always find some third case,

³⁵ We contrast first-order doxastic attitudes with second-order doxastic attitudes that are about differences in epistemic strength. While differences in second-order doxastic attitudes can register differences in epistemic strength, a theory that relies heavily on them suffers from the *wrong content* problem discussed previously.

CAR INSURANCE^Z, that falls strictly between them in this regard? If so, this shows that the set of doxastic attitudes must be (at least) a dense topological space akin to the set of rational numbers in order to register all these different possible epistemic positions.

The reason that registering epistemic positions matters for rational action is that many differences in epistemic position with respect to some relevant proposition $\langle p \rangle$ can make for corresponding differences in whether a particular action is (maximally) rational. To better improve rationality with respect to action, it pays to be able to register more differences in one's epistemic position with respect to a proposition. That way, no matter what the threshold for rational action turns out to be vis-à-vis strength of epistemic position, one will be able to determine whether it is met. This discussion suggests a remedy to the advocate of belief-based epistemology who wants to answer the Further Doxastic State Question. Just introduce further doxastic attitudes to go alongside belief—one for every possible epistemic position one could have with respect to a proposition.

§4

A CFE advocate might complain that this strategy concedes too much. Isn't acknowledging an infinite number of possible doxastic attitudes towards a proposition just letting credences in through the back door? Here, it is important to pay close attention to which aspects of credence-based epistemology do and which don't help with the determinations of rational action. What helps is the large number of credal states. But there are other features that don't seem to matter.

Consider a disparity between belief and credence noted earlier: the correctness of beliefs—their accuracy—is an all or nothing affair; in contrast, to the extent that accuracy makes sense for credal states, it is a matter of degree. This disparity seems to account for the fact that although both beliefs and credal states have propositional content—and are, in some sense, about the world—only beliefs are a genuine kind of mental *representation*, i.e. mental states that definitively render the world as being a particular way.

This feature of credences played no role in our discussion last section. The explanation given there was that because the number of credal states is infinite, an infinite number of different epistemic positions can be registered with respect to some proposition. And, this turns out to be significant since any difference of epistemic position could be a relevant difference for rational action; having credences puts one in a position to select the rational action across these differences in epistemic position. Nothing in this reasoning turns on whether credences are a genuine kind of mental representation or not. The existence of an injective mapping from epistemic positions to credal states is all that matters.

Here is another feature of credence that doesn't seem to matter: each credal state rationally precludes other credal states.³⁶ Holding a higher degree of credence in

³⁶ Suppose that, instead of working with credal states as they are traditionally understood, we work with at-least and no-greater-than credal states. Obviously, any given at-least or no-greater-than credal state will be rationally compatible with—i.e. not rationally precluded by—any number of other at-least or no-greater-than credal states. Nevertheless, it should be evident that working in a framework in which these sorts of

a proposition rationally precludes holding a lower. This is not to say that it would be irrational to hold a kind of mushy credence³⁷ that is indeterminate between the two. It is merely to say that given that one has committed oneself by holding a more determinate credal state, one is rationally precluded from simultaneously holding some distinct determinate credal state.

If the two features just discussed don't matter when it comes to explaining why credences appear to do better than belief vis-à-vis rational action, then we should be able to omit them without losing the virtues of the credence framework vis-à-vis rational action. In other words, the door is open to the following view: there is a hierarchy of basic doxastic attitudes with a very large cardinality.³⁸ For each of these doxastic attitudes, *D*, there is also a corresponding basic attitude of withholding *D*. Belief is among this hierarchy. But, there are stronger attitudes like being-sure or being-absolutely-certain. They are stronger in the sense that the rationality of belief doesn't entail but is compatible with the rationality of taking these attitudes. And, there are weaker attitudes like (at least) suspecting. They are weaker in the sense that the rationality of taking up these attitudes doesn't entail but is compatible with the rationality of belief.

Moreover, all of these attitudes are genuine mental representations that are entirely beholden to the world for their correctness (*simpliciter*), depending only on whether they are true. Indeed, in taking up any of these doxastic attitudes, the aim is to possess the truth by having that attitude. As an anonymous referee astutely noted, this latter "representational" difference seems to bring the first difference of rational compatibility of stronger and weaker doxastic attitudes with it. One way to appreciate this is to note that, were credences genuine representations, they would be very bizarre ones—since a .65 credence towards the proposition that it will rain today would, while increasingly likely to be accurate as one's epistemic position strengthened, would simultaneously become increasingly less rational. To fix this problem, genuine representations must not rationally preclude strictly stronger doxastic attitudes.

One might wonder how all these distinct doxastic attitudes could have the same aim of possessing the truth. Isn't belief individuated from other attitudes by having this aim?³⁹ Actually, no. In other work, we suggest that having the aim of possessing the truth is equivalent to being committed to implement good methods for possessing the truth on the basis of one's experiences.⁴⁰ This characterization doesn't say how good the methods must be; using better or worse methods corresponds to aiming more or less stringently at possessing the truth. Our proposal now is that every way of specifying how good the methods must be corresponds to a distinct doxastic attitude. When it

Footnote 36 continued

credal states are taken as fundamental puts one at no disadvantage relative to the more traditional framework when it comes to accounting for how an agent might, in principle, select a rational course of action.

³⁷ On mushy credences, see Elga (2010), White (2009), and Joyce (2010).

³⁸ Cf. Turri's (2010) suggestion that there is a hierarchy of speech acts. This is also similar to Williamson's (2002) countable hierarchy built on the suggestion that believing (*I know (p)*) is taking a stronger position believing (*p*). Our proposal is more radical and does not rely on giving up the KK principle.

³⁹ See references in footnote 10.

⁴⁰ See Jarvis and Katherine (2015).

comes to believing, one is committed to using methods that are fairly effective at landing true representations rather than false representations. However, in considering whether to take up or retain the stronger attitude of being sure that, one is committed to using methods that are strictly better in this regard.

At this point, it may be helpful to consider the familiar Jamesian point that the aim of possessing the truth is best understood as a mixture of two competing aims: truly representing and not misrepresenting.⁴¹ This dual-characterization has been espoused in different ways⁴² by Descartes (xx),⁴³ Alston (1985),⁴⁴ Foley (1987),⁴⁵ David (2001), Fallis (2006)⁴⁶ and many others. Importantly, truly representing and not misrepresenting are *competing* aims because one can only fulfill the first aim by putting oneself at risk of not fulfilling the second. Weighting the second aim more would lead one to be more cautious in order to avoid possible misrepresentation. Weighting the first aim more would lead one to be bolder in order to possess more truths.

Our proposal develops this Jamesian thought in a novel direction: different ways of mixing these two aims correspond to doxastic attitudes at different places in the hierarchy.⁴⁷ There are attitudes like being-sure, where the second aim is appropriately weighted more (so that misrepresentation is worse and failing to truly represent isn't so bad), and there are attitudes like suspecting, where the first aim is appropriately weighted more (so that failing to truly represent is worse and misrepresentation isn't so bad). But, the aim of all of these doxastic attitudes is to possess the truth by having the attitude. It's just the mixture of the twin aims—how important they should be in different situations—that varies.

We end this section by noting that “different ways of mixing” these two aims *may* but *need not* be construed as assigning each aim a scalar corresponding to its relative importance. (This scalar mixing approach appears to line up each distinct doxastic representation in the hierarchy with, what is in effect, some kind of minimal credal threshold.) Another way to mix the aims is to decide which sorts of error possibilities are tolerable in order take a chance on possessing the truth. For instance, misrepresenting because one is deceived by an evil Cartesian demon might be tolerable for belief, but not for absolute certainty. By tolerating *all* error possibilities, one is giving all the weight to truly representing. By tolerating *no* error possibilities, one is giving all the weight to not misrepresenting. But, of course, there are many ways

⁴¹ James (1897).

⁴² See David (2001), especially §1, for a nice survey of several of these positions.

⁴³ In Meditation IV, Descartes describes his aim as to arrive at the truth whilst taking precautions to “avoid ever going wrong.”

⁴⁴ Alston goes so far as to identify the dual-aims as constitutive of epistemic evaluation itself (Alston 1985, pp. 83–84).

⁴⁵ Foley puts the point in terms of epistemic rationality, which, for him, involves the goal of “now believing those propositions that are true and now not believing those propositions that are false” (Foley 1987, p. 8).

⁴⁶ As Fallis (2006, p. 180) puts it, “The two epistemic values most often discussed by epistemologists are having true beliefs and avoiding error.”

⁴⁷ Cf. Carter et al. (2015). Note also this seems contrary to the way in which James originally developed this point. It appears to undermine Kelly's (2013) Jamesian argument against uniqueness as defended by White (2005).*****

of tolerating *some* error possibilities, but *not others*. Each of these corresponds to a different way of mixing the aims of truly representing and not misrepresenting.⁴⁸

Our proposal in Sect. 6 develops this picture in more detail. In Sect. 6, we suggest that each doxastic attitude in the hierarchy might be characterized by a division of the space of possibilities into those that are more or less abnormal. To a first approximation, abnormality is the same as tolerability in the sense that *abnormal* error possibilities are *usually tolerable*. More exactly, possibilities are *more abnormal* for a doxastic attitude, *D*, to the extent that one has to have special evidence that they obtain before they *become intolerable* error possibilities for *D*. On the view in question, an agent rationally holds doxastic attitude, *D*, towards $\langle p \rangle$ if and only if $\langle p \rangle$ is true in, what are for *D*, the least abnormal possibilities that are strictly compatible with the agent's experiences.

§5

The picture we have developed thus far may not satiate those who contend that belief is theoretically significant. What we have is an alternative to credence-based epistemology, but it is representation-based rather than belief-based, per se.

In order to show that we have not abandoned belief, we must demonstrate how belief occupies a special place in the hierarchy of doxastic attitudes. Our suggestion is (roughly) that belief is the weakest doxastic attitude such that rationally holding that attitude towards a proposition normally suffices for rational acceptance of that proposition.⁴⁹ (Recall: acceptance is the mental state of taking some propositional content $\langle p \rangle$ for granted—whether consciously or not—in practical reasoning and rational decision-making; acceptance causes acting as if $\langle p \rangle$ is true.) Belief's being the 'weak-

⁴⁸ It is appropriate to consider a connection between the proposal being advanced here, and a particular kind of case that has been discussed in recent work by Reed (2010, p. 230), where a subject faces two different sets of practical stakes relative to the same proposition. Reed's case holds fixed a question you are asked—a question about Roman History, where the answer (of which you are confident) is that Caesar was born in 100BC—with two different practical stakes relative to this proposition which must be considered. He writes:

'...[T]here are two punishment/reward scenarios. The first is as before: for a correct answer, you receive a jelly bean, and for an incorrect response, you get a very painful electric shock. In the second, a correct answer results in a \$1000 reward, while an incorrect one leads only to a very mild slap on the wrist. In both scenarios, there is neither reward nor punishment for failing to answer. Although you must consider both scenarios simultaneously, you are not bound to give the same answer in each' (Reed 2010, p. 230).

While the context of Reed's presentation of the case is that of raising an objection to subject sensitive invariantism (pp. 228–230), it's helpful to note how the view we are proposing can offer a very natural way to think about such a case. With respect to the first punishment/reward scenario, it is not rational to answer on the basis of belief; with respect to the second, it is rational to answer on the basis of either belief or perhaps also the weaker attitude of suspecting-that. Thanks to an anonymous referee at *Synthese* for drawing our attention to such cases.

⁴⁹ A contrast may be helpful here. Consider Bratman (1987, p. 40): "To believe something is not merely to assign a high probability to its occurrence. I might assign a high probability to my failing to move the log without believing that I will fail ... what seems distinctive about believing that I will fail is that it puts me in a position to plan on the assumption of failure." On our view, fairly strong epistemic positions—of the sort Bratman seems to be alluding to by introducing 'high probability'—do (at least typically) countenance belief, so the relationship between believing and planning is much weaker, holding only ordinarily.

est' doxastic attitude with feature F means that there is no other doxastic attitude, D , with feature F such that rationally believing entails rationally holding D , but not vice-versa. 'Normally suffices' means that special circumstances—and in particular, risky payoff structures—have to become apparent in order to break the link between rational belief and rational acceptance. Risky payoff structures may be apparent if, for example, the subject has reason to believe or even suspect that standing causal regularities allow for significant changes in utility that would be attributable to having taken one course of action rather than another.⁵⁰ (Notice that credence is not obviously necessary for tracking risky payoff structures.) When the link between rational belief and rational acceptance is broken, it must be possible in principle for the subject to explain the link away by citing the special circumstances in question. But, no explanation is owed for why the link holds in canonical cases. 'Normally suffices' is plausibly vague; it is not clear which are the normal possibilities in which the link to rational acceptance has to hold for a doxastic attitude to count as belief. But, we think that this fuzziness corresponds to a genuine indeterminacy in which among the doxastic attitudes in the hierarchy is belief.

Our theory of belief allows us to make belief the starting point of rational decision-making. When one is deliberating about what to do, it makes sense to start by considering whether one's beliefs support the conclusion that one of the possible actions—say ψ -ing—would be best.⁵¹ If so, ψ -ing is ordinarily rational. If not, it may not be immediately clear what to do. One may have to consider what other doxastic attitudes one has as well as the potential costs and benefits of both correctly and erroneously acting as if some particular action would be best.

It is beyond the scope of this paper to delve further into the mechanics of (bounded or unbounded) rational decision-making.⁵² It suffices for our purposes to emphasize that the hierarchy of doxastic attitudes introduced previously is compatible with putting belief at the center of the rational decision-making process. The suggestion is that our first resort when considering what to do is simply to consider whether there is any possible action that we believe to be best vis-à-vis our ends. In other words, belief has the first word on what means to our ends we should take, even if it doesn't have the last. That is a significant enough role for belief to have in determining action.

§6

At the end of Sect. 3, we suggested that a representation-based epistemology can imitate credence-based epistemology in a respect that's important for accounting for rational action.⁵³ Let us pause to defend this suggestion briefly. Just as there are many

⁵⁰ By 'apparent' we mean something like: 'psychologically salient'.

⁵¹ This is, in effect, Lin's (2013) Cliché Rule.

⁵² Notably, Lin (2013) gives a qualitative, belief-based decision theory that is straightforwardly generalizable to other doxastic attitudes that are representational like belief. One might try adding a procedure for determining which of these doxastic attitudes to use to make the decision.

⁵³ In effect, we have turned the tables on credence sympathizers such as Joyce (1998) and Moss (2013) who attempt to show that features of traditional epistemology—accuracy and Gettierization, respectively—can be incorporated into a credence framework.

credal states, so there can be many doxastic representational states. Moreover, just as credal states are assigned real numbers from the unit interval, so we could assign doxastic representational states real numbers from the unit interval in accordance with their strength. As a consequence, there would be enough doxastic representational states to register all the epistemic positions registered by credences. And, in principle, that should allow for equally fine-grained determinations of action by both. This should be obvious; in principle, one could assign credence level r (where $0 \leq r \leq 1$) to a proposition $\langle p \rangle$ if and only if level r corresponds with the strongest doxastic representational state it would be rational to hold towards $\langle p \rangle$, and calculate expected utilities in the usual way.

Imitation is an available theoretical option that puts representation-based epistemology on a par with credence-based epistemology. However, in the remainder of this section, we consider a potential reason not to favor imitation and consider what an alternative to imitation might look like.

Consider SAN DIEGO.

SAN DIEGO: John and Jane are wondering whether there are any mules in the San Diego Wild Animal Park that are cleverly disguised as zebras. John knows that students from UC Sunnydale and Caltech had planned to put a mule cleverly disguised as a zebra in the Park, but that the Caltech students cancelled their plans to participate. John doesn't know anything about how difficult this prank would be to pull off. Jane has heard a rumor that somebody might be interested in the prank, but doesn't know who is interested or what their plans might be. Jane works for the San Diego Wild Animal Park and knows about the security at the Park. She knows it would be very difficult for UC Sunnydale students to get past the security, but, because of differences in the engineering curriculum, the Caltech students could figure out how.

How should we compare the strength of the epistemic positions that John and Jane occupy with regard to the proposition (*There is no mule in the San Diego Wild Animal Park that is cleverly disguised as a zebra*)? John is in a better position to rule out possibilities in which cleverly disguised mules are placed in the Park by Caltech students, but Jane is in a better position to rule out possibilities in which cleverly disguised mules are placed in the Park by UC Sunnydale students. Which matters more? It is a theoretical possibility that this question doesn't have an answer.

We can put this point more abstractly. In the story, John's experiential condition is strictly incompatible with some of the possibilities in which the proposition highlighted last paragraph is false.⁵⁴ Many of these are possibilities that are not strictly incompatible with Jane's experiential condition. Let the set of these possibilities be S . Jane's experiential condition is also strictly incompatible with some of the possibilities in which this proposition is false. Many of these are possibilities that are not strictly incompatible with John's experiential condition. Let the set of those possibilities be S^* . Whether John's or Jane's epistemic position is stronger rests at least partly on whether, by having experiences strictly incompatible with S and S^* respectively, John's experi-

⁵⁴ "Strictly incompatible" in roughly the sense of Lewis (1996, p. 553).

ential condition is strictly incompatible with “more” or “less” possibilities than Jane’s is. But, the cardinality of these two sets may be the same. And, neither is a subset of the other. Maybe the possibilities of one of these sets are, in some suitable epistemic sense, collectively more probable than the possibilities of the other are. However, it strikes us that comparisons of epistemic probability of this sort need not make sense. After all, it is hardly clear that *⟨There is no mule in the San Diego Wild Animal Park that is cleverly disguised as a zebra⟩* should have a definitive epistemic probability for anyone. As a result, the epistemic positions of John and Jane may be *incommensurable*. John’s epistemic position may be better in one respect. Jane’s epistemic position may be better in another. But, there may be no fact of the matter about which is better *simpliciter*.

In effect, we are considering the theoretical possibility that epistemic positions need not have the structure of the real numbers between zero and one inclusive.⁵⁵ Epistemic positions may not be *well-ordered* or even *total ordered* by strength; the ordering may only be *partial*. For any two epistemic positions x and y , it may be that neither is x stronger than y , nor is y stronger than x , nor are they equal in strength. (The suggestion is that some epistemic positions might be incommensurable, not that all are.) If so, then degrees of beliefs—credences—are exactly the *wrong* kind of doxastic attitude to register epistemic positions because degrees of beliefs are well-ordered.⁵⁶ Persons in epistemic positions of incommensurable strength with respect to some proposition would have to assign less, more, or the same credence than one another even though their respective epistemic positions are neither stronger, weaker, nor exactly the same.

One might wonder how there could be epistemic positions of incommensurable strength. Couldn’t we discover the exact strength of an epistemic position for some proposition $\langle p \rangle$ by seeing what kinds of odds a rational agent in that epistemic position would need to bet on the truth of $\langle p \rangle$ —ignoring, of course, practical or moral considerations that might distort betting patterns? However, the theoretical possibility under consideration is precisely one in which the epistemic position itself does not warrant any particular betting pattern. Of course, if offered a series of bets with improving odds, a rational agent will be forced into adopting a betting pattern of some sort. However, in the theoretical possibility under consideration, the choice of betting pattern will not be adequately constrained by the agent’s epistemic position. She will be forced into a betting pattern that is, to some degree, arbitrary at least as far as her epistemic position with respect to $\langle p \rangle$ is concerned.

Moreover, we think that the betting pattern of this rational agent need not be directly revelatory of her total doxastic state either. By way of analogy, consider Buridan’s ass. Even if the ass opts for the stack of hay on the left rather than right, this need not indicate a preference on the part of the ass for the stack of hay on the left rather than

⁵⁵ It is worth noting that even credence sympathizers worry that epistemic positions need not have the structure of the unit interval. Cf. Sturgeon (2008, p. 159). Consider also that Joyce’s (2010) plea for imprecise credences seems to be based on the thought that, for any proposition $\langle p \rangle$, there are epistemic positions that are ambiguous so that there is no determinate fact as to how they compare with epistemic positions in which facts about statistical frequencies relevant to the truth of $\langle p \rangle$ are known.

⁵⁶ Unless, of course, degrees of belief come in more than one dimension as they do in the Dempster–Shafer theory. On this model, a proposition is assigned a degree of belief and a degree of ignorance. See Huber (2012) and Pryor (unpublished).

right. Ex hypothesi, the ass has no reason to prefer the stack of hay on the left, so, arguably, no such preference is rational. Of course, the ass should implement some plan or other in order to eat hay, but this plan might well be settled upon arbitrarily rather than wholly on the basis of an ungrounded preference for left. Similarly, we think, a particular betting pattern with respect to $\langle p \rangle$ need not reveal some *doxastic* attitude of definitive strength towards $\langle p \rangle$, particularly if that doxastic attitude would not be warranted by a current epistemic position that is not only fallible, but genuinely equivocal (i.e. not of definitive strength) with respect to $\langle p \rangle$. Instead, the betting pattern may reveal only a strategic coping plan of the agent—e.g. a *pragmatic* credal function rather than one constituting a genuine opinion with regards to $\langle p \rangle$. The rational agent may chose this strategic coping plan somewhat arbitrarily albeit partly by reference to a total doxastic state that is also equivocal (i.e. not of definitive strength) with regards to $\langle p \rangle$ so as to precisely mirror her equivocal epistemic position.

The hierarchy of doxastic representations introduced in Sect. 4 is well equipped to handle the theoretical possibility under consideration. The hierarchy *may* but *need not* be well-ordered; in fact, it may only be partial ordered. So, it might do better at registering sometimes equivocal and therefore only partially ordered epistemic positions than the credal scale. Of course, a hierarchy of doxastic attitudes wouldn't have to be a hierarchy of doxastic *representations* in order not to be total ordered by strength. Still, we can see that representation-based epistemology may have some flexibility that credence-based epistemology does not obviously have—at least insofar as credence is construed as a kind of subjective probability.

How much this flexibility matters depends partly on whether we can find a case where difference of epistemic position matters for whether an action is rational even though this difference is between epistemic positions that are of incommensurable strength. We won't pursue further whether there are such cases. We think it may be interesting enough that, in principle, a representation-based psychology might be better equipped to register differences in the strength of epistemic positions than a credence-based psychology is regardless of whether this makes any difference for rational action.

If the hierarchy of doxastic attitudes isn't a total ordering by strength, what kind of structure might it have? As alluded at the end of Sect. 4, perhaps, the strength of doxastic attitudes is ordered by what sorts of possibilities are “relevant” in considering whether to hold the attitude or, equivalently, by which sort of possibilities are “irrelevant” so that error in them is tolerable.⁵⁷ In considering whether to believe a proposition—e.g. *(There is a goldfinch in the garden)*—certain skeptical possibilities in which the proposition is false—e.g. possibilities involving fake goldfinches—might be irrelevant. They might be irrelevant in the sense that the strict compatibility of these possibilities with one's experiential condition need not count against belief. But, in considering whether to be absolutely certain that the proposition is true, these same possibilities might be relevant; they must be strictly incompatible with one's experiential condition in order for taking up this attitude of absolute certainty to be rational. The idea would be that for every set of possibilities, there is a unique doxastic attitude

⁵⁷ Cf. Austin (1946), Dretske (1970), Goldman (1976), and Pritchard (2010).

for which this is the set of relevant possibilities. Suppose we have doxastic attitudes D and D^* with sets X and X^* as their sets of relevant possibilities, respectively. Then, D is stronger than D^* if and only if X is a proper superset of X^* . If X is not a proper superset of X^* , then it will turn out that D is not stronger than D^* . And, if X^* is also not a proper superset of X , then it will turn out that D^* is not stronger than D . Then, D and D^* turn out to be incommensurate with respect to strength.

As attractive as this suggestion might be, it can't be quite right. The problem is that while certain possibilities might not be relevant now, they might become relevant in more bizarre circumstances. Imagine that you know that there is a prankster in neighborhood planting fake goldfinches in gardens. These bizarre circumstances might make possibilities involving fake goldfinches relevant even if they weren't relevant before. It might be that it wouldn't be rational to believe *(There is a goldfinch in the garden)* unless one's experiences are strictly incompatible with these possibilities. This suggests that we move to a framework instead where possibilities are more or less abnormal, where more abnormal possibilities become relevant to whether to believe on the condition that one discovers oneself to be in more bizarre circumstances. Stronger doxastic attitudes will be ones for which "crazier" possibilities are taken to be more normal, and thus more easily relevant to whether to hold those attitudes. Weaker doxastic attitudes will be ones for which only slightly surprising possibilities are taken to be more abnormal, and thus less easily relevant to whether to hold those attitudes.

In fact, the kind of framework that we have in mind is explored formally by those interested in non-monotonic consequence relations.⁵⁸ With classical monotonic consequence, if $\langle p \rangle$ is a consequence of $\langle q \rangle$, then it is also a consequence of $\langle q \rangle$ and $\langle r \rangle$. This is no longer valid for non-monotonic consequence relations. Non-monotonic consequence relations are, of course, an attempt to formalize genuinely "ampliative" reasoning where the conclusion may "go beyond" what is contained collectively in the premises. Because conclusions "go beyond" the premises, one may have to backtrack on the conclusion upon learning more even without giving up any of the previous premises. For this reason, the same conclusions don't necessarily follow from strengthened premise sets. As it turns out, non-monotonic consequence relations can be characterized or "represented" by truth-preservation in the set of "least abnormal" possible worlds.^{59,60} (Good ampliative reasoning is reasoning to a conclusion that

⁵⁸ For a solid introduction to non-monotonic logic, see Makinson (2005).

⁵⁹ This result is reported by Huber (2012). But, it should be intuitive enough to those familiar with Lewis (1973) semantics for counterfactuals. The counterfactual relation between antecedent and consequent is, in effect, a non-monotonic consequence relation since Strengthening the Antecedent does not hold. And, famously, whether this counterfactual relation holds depends on whether the consequent is true in the nearest (read: 'least abnormal') worlds in which the antecedent is true.

⁶⁰ An anonymous referee wondered whether the incommensurable doxastic attitudes we proposed leads to permissivism as discussed in White (2005). The characterization of rationality we give here, in fact, entails anti-permissivism. For any (determinate) doxastic attitude, either it will be true in what are for it all the least abnormal possible worlds compatible with the subject's experiential condition or it won't. Thus, any two subjects that differ with respect to the rationality of a doxastic attitude must *ipso facto* differ in their experiential conditions in a way that matters for rationality. So, in effect, they won't "share the same total evidence" with respect to the proposition in question. This proves where uniqueness fails, there is not the same total body of evidence—i.e. the contrapositive of anti-permissivism. Notably, this proof does not rely on a trivialization of anti-permissivism as the argument does not presuppose that *all* difference in

can be false even when one's premises are true; more abnormal possible worlds will be—by definition—those where this happens.) Different non-monotonic consequence relations can thus be characterized or “represented” by different orderings of possible worlds in terms of normality. When you impose a normality ranking on possibilities, you are, in effect, imposing a non-monotonic consequence relation.

Our (toy) suggestion, then, appears to amount to this: doxastic attitudes within our hierarchy are individuated by non-monotonic consequence relations—or more intuitively, by the differences in what counts as valid ampliative reasoning for that attitude—that generate, at a particular time, the doxastic attitudes from (the same full set of) premises offered up by a subject's experiences. By way of concrete example, reasoning (in the absence of any particular background information) from the starting point of a perceptual experience as of something with a barn façade to *⟨That thing is a barn⟩* may be valid non-monotonic reasoning for belief, but not for a much stronger doxastic attitude. The idea is this difference may participate in individuating belief from its stronger counterpart (in roughly the same way having a non-negative predecessor helps to individuate the number one from zero). More generally, weaker doxastic attitudes have very strong non-monotonic consequence relations—you can get a lot more out (in terms of “deductive” strength of the conclusion) for what you put in (in terms of the collective “deductive” strength of the premises). At the limit would be an absurd doxastic attitude that it would be rational to take towards any proposition. This corresponds to a limiting case of a “non-monotonic” consequence relation that allows you to get everything on the conclusion side out of anything on the premise side. (It is a limiting case because this “non-monotonic” consequence relation is monotonic; monotonic consequence relations are a special case of “non-monotonic” consequence relations on our understanding.) Stronger doxastic attitudes have very weak non-monotonic consequence relations. At the limit would be the doxastic attitude of absolute certainty that it would only be rational to take towards propositions that strictly follow from one's experiential condition. This corresponds to another limiting case of a “non-monotonic” consequence relation that allows you to only get out on the conclusion side what you put in on the premise side. (Again, it is a limiting case because this “non-monotonic” consequence relation is monotonic.)

Notice that non-monotonic consequence relations are not total ordered by strength. Of course, some non-monotonic consequence relations are stronger than others, in the sense that the consequences of the former (for any premise set) include the consequences of the latter (for that same premise set) and more besides. So, the doxastic attitude individuated by a stronger non-monotonic consequence relation will be (strictly) weaker than the doxastic attitude individuated by the weaker non-monotonic consequence relation. But, in many cases, we don't have “consequence inclusion” of

Footnote 60 continued

experiential condition leads to differences in same total evidence; as far as the argument is concerned, the sets of least abnormal possible worlds compatible with two genuinely distinct experiential conditions might be entirely the same in certain cases. (To put the point yet another way, the argument is, in fact, compatible with the enrichment discussed in the next footnote, whereby some genuinely distinct evidential conditions might support rationalizes every proposition to exactly the same extent; the argument merely relies on the idea that where there are differences in rationality, this is clearly not the case.)

this sort. So, correspondingly, we will have two doxastic attitudes that are potentially⁶¹ incommensurable in strength.

By way of illustration, consider again John and Jane from SAN DIEGO. Ex hypothesi, John's experiential condition is strictly incompatible with possibilities of S ; so, in principle, he doesn't need anything from a non-monotonic consequence relation to rule them out. But, he would need something from a non-monotonic consequence relation to rule out the possibilities of S^* ; those aren't strictly compatible with his experiential condition. Jane is in the reverse position. For John, whether it is rational for him to take up a doxastic attitude, D , towards *(There is a mule in the San Diego Wild Animal Park that is cleverly disguised as a zebra)* depends on whether the non-monotonic consequence relation individuating D lets him rule out the possibilities of S^* given his experiential condition. For Jane it depends on whether the non-monotonic consequence relation individuating D lets her rule out the possibilities of S given her experiential condition. Suppose that the latter is true, but the former isn't. In such a case, Jane is rational in taking up D , but John isn't. Still, there should be some other non-monotonic consequence relation that lets John rule out the possibilities of S^* given his experiential condition, but doesn't let Jane rule out the possibilities of S given hers. And, this non-monotonic consequence relation should individuate a further doxastic attitude, D^* that it will be rational for John to take up but not Jane. So, D and D^* are of incommensurable strength; the rationality of taking up one doesn't entail the rationality of taking up the other.⁶²

Some of the doxastic attitudes individuated by non-monotonic consequence relations are, of course, completely ridiculous.⁶³ There is a doxastic attitude within this hierarchy that it is rational to take up towards *(I am a recently disembodied spirit deceived by a Cartesian demon)* in the ordinary sort of case where one has a perceptual experience as of one's hand—at least assuming the latter experience doesn't provide conclusive reason rejecting this proposition. This is simply because there is a way of ranking possibilities by normality so that bizarre possibilities including Cartesian demons get ranked the most normal. Belief is obviously not an attitude of this sort. The non-monotonic consequence relation individuating belief would have to be characterized by a normality ranking on possibilities that is fairly intuitive; the possi-

⁶¹ The framework we are exploring can be enriched. For instance, for any two experiential conditions, e and e^* , it can be stipulated that e rationalizes a proposition $\langle p \rangle$ to at least the same degree as e^* in the sense that if a doxastic attitude, D , towards $\langle p \rangle$ is epistemically rational on the basis of e^* , it is also epistemically rational on the basis of e . Such stipulations would have the effect of imposing more order on the envisioned hierarchy. If such stipulations were made for every pair of experiential conditions and every proposition, the effect would be a total ordering on doxastic attitudes (so that, once again, a credal scale becomes entirely appropriate for registering differences in strength of epistemic position).

⁶² Here, for simplicity and illustrative purposes, we're ignoring the possibility of enrichment raised in the previous footnote. In fact, because it can be enriched, our framework does not preclude that the experiential conditions of Jane and John might make the same doxastic attitudes epistemically rational. The point is that it allows for the kind of theoretical possibility discussed in text, not that it mandates it.

⁶³ Perhaps in response to this possibility, an anonymous referee worries whether the doxastic states pluralism we develop will lead to licensing intractable disagreement with distinct rationally held doxastic attitudes towards contradictory propositions. However, rationally co-held doxastic states from the hierarchy don't "clash" in a way required for disagreement—in the way that believing $\langle p \rangle$ and rejecting $\langle \text{not } p \rangle$ don't clash despite having contradictory contents.

ble worlds ranked as more normal by the ranking function for belief really are more normal. For that reason, rational belief provides a *pro tanto* reason for acceptance: because rational beliefs are true in possible worlds that really are more normal, it is ordinarily reasonable to rely on them in deliberating about what to do.

It is not clear whether in the envisioned hierarchy there will be a unique weakest doxastic attitude such that rationally holding that attitude towards a proposition normally suffices for rational acceptance of that proposition. This is not just because ‘normally suffices’ is vague. For example, a number of doxastic attitudes incommensurate in strength may normally suffice for rational acceptance even though no strictly weaker attitudes would. We think that this would simply indicate some further indeterminacy in which among the hierarchy of doxastic attitudes is belief.

Before we move on, it may be helpful to say something more about the relation between the envisioned hierarchy of doxastic attitudes and Bayesianism. Although the envisioned hierarchy is only partially ordered by strength, an ideally rational agent should be able to superimpose a kind of pragmatic credence function for purposes of decision-making partly on the basis of her doxastic attitudes.⁶⁴ This may be the best way to deal with a kind of situation raised earlier: a well-order series of proposed bets on the truth of $\langle p \rangle$ with improving odds. Certainly, the model of doxastic attitudes under consideration doesn’t prohibit an agent from assigning pragmatic credences and doing so in accordance with the axioms of probability. Indeed, the model might naturally be supplemented so as to give some positive guidance as to how to assign pragmatic credences. For instance, perhaps relatively high pragmatic credence should be assigned to at least certain propositions that it is rational to believe.⁶⁵ Presumably, the model might put other constraints on the assignment of pragmatic credence as well (or instead). Once pragmatic credence is assigned in conformance with these constraints, rational decision-making could proceed by calculating expected utilities in the usual way.

Of course, as alluded to earlier, a pragmatic credence function generally won’t be *uniquely* determined from the doxastic representations that it is rational to hold, but we shouldn’t expect it to be. Some arbitrariness in assigning (any kind of) credence to a proposition is to be expected once it is conceded that the space of epistemic positions (as ordered by strength) doesn’t have the structure of the unit interval. If we can’t even make sense of the idea that some stronger epistemic position, e , is (literally) twice as strong as another epistemic position, e^* , then the choice of whether or not to register the difference in strength with a credence that’s twice as large seems arbitrary. Unless strengths of epistemic position line up with numbers from the unit interval, they likely won’t dictate a particular assignment of credence. The arbitrariness of the assignment of credence will obviously carry over to rational decision-making. This is an apparent

⁶⁴ Cf. Wedgwood (2012).

⁶⁵ Some restriction on which rationally believed propositions receive high credence is necessary, since, on the model under consideration, rational belief is closed under multi-premise deduction. A speculative suggestion: high credence rationally believed propositions might be ones such that the rationality of belief exhibits less in the way of asymmetrical dependence. Roughly speaking, a rational belief, B , *asymmetrically depends* on a rational belief, B^* , just in case if B^* weren’t rational (because the subject’s experiences were relevantly different), B wouldn’t be either, but if B weren’t rational (because the subject’s experiences were relevantly different), B^* still might be.

cost of acknowledging that epistemic positions can be of incommensurable strength (and hence aren't like real numbers).⁶⁶

Nevertheless, the assignment of pragmatic credences to propositions on the basis of doxastic attitudes from the entire hierarchy isn't nearly as arbitrary as an assignment merely on the basis of beliefs. The doxastic hierarchy is built to be able to register *all aspects* of the strength of one's epistemic position towards a proposition rather than merely whether the epistemic position is strong enough for belief to be rational. In effect, an ideally rational agent will "mirror" the epistemic position she stands in with respect to a proposition by registering all the doxastic attitudes and withholdings of doxastic attitudes from the hierarchy that are rational. If this total "mirroring" state doesn't uniquely constrain pragmatic credence, it's simply because the agent's *epistemic position* doesn't uniquely constrain pragmatic credence. The total "mirroring" state *will* constrain pragmatic credence insofar as the agent's epistemic position does. Thus, any pragmatic credence assignment properly based on this total mirroring state is as good as it gets for decision-making, even if the assignment isn't uniquely proper.

In any case, the principal point we want to make is that it is possible to embed the machinery for rational decision-making from the Bayesian framework within the model of doxastic attitudes that we have been considering. Consequently, *reasons to favor calculating expected utilities as the uniquely ideal way of making rational decisions are not obviously reasons to reject this model*. At same time, reasons to favor other methods are not necessarily reasons to reject the model either since other methods of decision-making could be appended to the model instead.⁶⁷

Furthermore, on this model, epistemic rationality has nothing to do with probability. Rather, for a doxastic attitude to be epistemically rational, it must be true in what are for it all the least abnormal possible worlds compatible with the subject's experiential condition.⁶⁸ This is a *safety* conception of epistemic rationality rather than a probabilistic one.⁶⁹ Given this significant departure from anything like probabilism, *it cannot be said that the model of doxastic attitudes under consideration is parasitic on the Bayesian framework*. Together, the two italicized statements in this and the preceding paragraph suggest that the model under consideration is neither obviously implausible (at least as an idealization), nor a mere imitation of Bayesianism.

Of course, it doesn't follow that we should immediately accept the model either. The model suggests that, at any point of time, an agent is delivered a set of premises—presumably by his or her experiences up to that point in time—that determine the rationality of various kinds of doxastic representations by way of their individuating

⁶⁶ Cf. Joyce (2010).

⁶⁷ For instance, a generalization of Lin's (2013) approach. See footnote 54.

⁶⁸ An anonymous referee pointed out that this theory of epistemic rationality has the alleged problem of entailing that it is epistemically rational to believe any necessary truth. Ichikawa & Jarvis (2013) argues at length that the entailed is not, in fact, a genuine problem, but exactly the right result. However, even putting that issue to the side, this entailment is shared with probabilism and most theories that one might be able to formalize in principle. Thus, one might take it to be an unintended consequence of an otherwise useful idealization.

⁶⁹ See, for instance, Pritchard (2005, 2007) and Sosa (1999) for some representative articulations of a safety condition. We're thinking of safety in a fairly formal way (as truth in nearby worlds meeting some criterion) that doesn't commit us to any particular doctrine of Pritchard or Sosa.

non-monotonic consequence relations. But, of course, the deliverance of a premise set by one's experiences may well be a point of concern for some.⁷⁰ And, there are probably other problems as well. Fortunately, we need not defend this model. We merely raise it as a theoretical alternative to emphasize that representation-based epistemology need not imitate credence-based epistemology, which may turn out to be an advantage rather than a deficit.

§7

In this last section, we will consider a problem for our proposal. The suggestion in Sect. 5 was that belief has the first word when it comes to rational decision-making. The implication is that cognizers can largely get along in their rational decision-making by trafficking in beliefs. But, belief does not have the last word. There are certain cases—as illustrated by CAR INSURANCE—where a rational agent needs to move beyond beliefs to consider other doxastic attitudes. However, this raises the question of how a rational agent is to keep track of whether it is rational to hold stronger or weaker doxastic attitudes in order to act rationally in these cases. Must a cognizer simultaneously be trafficking in these other doxastic attitudes in order to act rationally when these cases arise? If so, then little advantage is gained by pointing out that a rational agent need not consider these attitudes ordinarily. If a rational agent has to be managing these attitudes anyway alongside belief, then cognition is already very taxing.

We think that this problem is not fatal for our proposal. Depending on a rational agent's cognitive resources, it may make sense to keep track of some other doxastic attitudes—e.g. being-sure or suspecting—alongside of beliefs. However, we also think that a great deal of information about whether to hold other doxastic attitudes towards a proposition is captured by our system of beliefs. Perhaps this is easiest to see when the doxastic attitude in question is weaker than belief. It doesn't seem so far-fetched to suppose that we can evaluate the plausibility of more speculative theories on the basis of our beliefs. Presumably, though, this can help us to appreciate whether weaker doxastic attitudes towards these speculative theories are rational. Notice, though, that this should also help us to appreciate whether strong doxastic attitudes are rational: if it is rational to suspect $\langle p \rangle$, then this limits how strong a doxastic attitude it would be rational to take towards $\langle not p \rangle$.

Our system of beliefs may also be able to tell us quite a bit about whether to hold stronger doxastic attitudes because we have extensive knowledge about how our belief-management system works. For instance, we know that beliefs about features in our environment that are currently visible are of higher quality because visual perception is fairly dependable. So, stronger doxastic attitudes towards these contents are typically rational. On the other hand, we know that beliefs that are more removed from immediate perceptual experience tend to be of lower quality as are beliefs that are partly maintained by memory. So, stronger doxastic attitudes towards these contents may not be rational. Of course, even beliefs removed from immediate experience can be of

⁷⁰ Cf. McDowell (1994). We submit that this aspect of the model might be altered at the expense of making it more baroque; we won't attempt this alteration here.

fairly high quality if they have support from multiple sources. Often (even if not very often) we know something about how beliefs are supported, i.e. we have beliefs about their credentials. This may allow us to see that stronger doxastic attitudes are, in fact, rational. Indeed, even if we can't remember the exact credentials of some particular belief, we can usually tell quite a bit about which faculties have been involved in managing it by the belief's content. These discriminatory abilities again tell us something about whether a stronger doxastic attitude might be rational. Obviously, some information about whether to hold stronger doxastic attitudes towards propositions that we believe is lost, but a surprising amount seems to be stored implicitly by believers who are—in roughly the sense of Sosa (2011)—reflective. By remembering and otherwise forming rational beliefs about the sources of belief and the general quality of those sources, a bounded cognizer can do a fairly good job of determining whether to hold doxastic attitudes that are stronger than belief.

The upshot is, as we alluded at the end of Sect. 2, that there is a certain sense in which the answer to the Further Doxastic State Question might very well be “further beliefs.” By way of reminder (again), the Further Doxastic State Question is “What are those further doxastic states that put you in a position to appreciate that some possible action would be, in fact, rational in cases like CAR INSURANCE in which, by your own beliefs, this action is worse?” The answer implicit in Sects. 4–5 is “the withholding of a doxastic attitude that is stronger than belief (but necessary for rationally acting as if the belief is true given the special circumstances of the case).” However, what we are now suggesting is that whether to withhold this stronger doxastic attitude might be roughly determined on the basis of one's system of beliefs. Consequently, in a certain sense (that avoids the wrong content problem mentioned in Sect. 2), the further doxastic states might turn out to be beliefs after all. That obviously would work in favor of those that emphasize the theoretical importance of belief.

Conclusion

The principal project of this paper is to defuse a certain problem for those that maintain the theoretical prominence of belief. The problem comes to light when attempting to account for rational decision-making without making reference to rational credence (as a way of tracking either rational action or rational belief), particularly in cases with risky payoff structures (as illustrated by CAR INSURANCE).

However, at its core, the problem is really that a psychology with only belief and withholding belief struggles to have the capacity to register all the nuances of the subject's epistemic position with respect to any given proposition. We suggest the problem can be solved by adding to the psychology. Further *representational* doxastic attitudes can be used to reflect more about one's epistemic position. Because they are representational in the sense that their truth conditions are conditions for accuracy-entailing correctness, the epistemology of these doxastic attitudes should be very similar to the epistemology of belief (and unlike the epistemology of credence).

We proposed that *ideal* rational decision-making might draw on these other doxastic representations. However, whether to hold these different doxastic representations can often be determined on the basis of one's beliefs. So, *in practice*, belief without

credence might well be enough to account for rational decision-making across a wide range of cases. Either way, the resulting alternative picture retains the central features of traditional epistemology—most saliently, an emphasis on truth as a kind of objective accuracy—while adequately accounting for rational action.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Alston, W. (1985). Concepts of epistemic justification. *The Monist*, 68(1), 57–89.
- Austin, J. L. (1946). Other minds. *Proceedings from the Aristotelian Society Supplement*, 20, 148–187.
- Bradley, D., & Leitgeb, H. (2006). When betting odds and credences come apart: More worries for Dutch book arguments. *Analysis*, 66(290), 119–127.
- Bratman, M. (1987). *Intention, plans and practical reason*. Cambridge, MA: Harvard University Press.
- Bratman, M. (1992). Practical reasoning and acceptance in a context. *Mind*, 101(401), 1–16.
- Broome, J. (2013). *Rationality through reasoning*. Oxford: Blackwell.
- Carter, J. A., Jarvis, B., & Katherine, R. (2015). Varieties of Cognitive Achievement. *Philosophical Studies*, 172(6), 1603–1623.
- Chalmers, D. (2011). Frege’s puzzle and the objects of credence. *Mind*, 120(479), 587–635.
- Chalmers, D. (2012). *Constructing the world*. Oxford: Oxford University Press.
- Christensen, D. (2004). *Putting logic in its place: Formal constraints on rational belief*. Oxford: Oxford University Press.
- David, M. (2001). Truth as the epistemic goal. In M. Steup (Ed.), *Knowledge, truth and duty: Essays on epistemic justification, responsibility and virtue* (pp. 151–169). Oxford: Oxford University Press.
- Descartes, R. (1641/1985). *The philosophical writings of descartes* (Vol. 2). Cambridge: Cambridge University Press.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67(24), 1007–1023.
- Dummett, M. (1959). Truth. *Proceedings from the Aristotelian Society*, 59, 141–162.
- Elga, A. (2010). Subjective probabilities should be sharp. *Philosopher’s Imprint*, 10(5), 1–11.
- Engel, P. (1998). Believing, holding true and accepting. *Philosophical Explorations*, 1(2), 140–151.
- Fallis, D. (2006). Epistemic value theory and social epistemology. *Episteme*, 2(3), 177–188.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Foley, R. (1987). *The theory of epistemic rationality*. Cambridge, MA: Harvard University Press.
- Foley, R. (2009). Belief, degrees of belief and the Lockean thesis. In F. Huber & C. Schmidt-Petri (Eds.), *Degrees of belief* (pp. 37–47). Dordrecht: Springer.
- Ganson, D. (2008). Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139(3), 441–458.
- Gibbard, A. (2005). Truth and correct belief. *Philosophical Issues*, 15(1), 338–350.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(20), 771–791.
- Greco, D. (2012). The impossibility of skepticism. *Philosophical Review*, 121(3), 317–358.
- Grimm, S. (2011). On intellectualism in epistemology. *Mind*, 120(3), 705–733.
- Hajek, A. (2008). Dutch book arguments. In P. Anand, P. Pattanaik, & C. Puppe (Eds.), *The Oxford handbook of rational and social choice* (pp. 173–195). Oxford: Oxford University Press.
- Hajek, A. A puzzle about degrees of belief. Manuscript.
- Hajek, A., & Eriksson, L. (2007). What are degrees of belief? *Studia Logica*, 86(2), 183–213.
- Harman, G. (1986). *Change in view: Principles of reasoning*. Cambridge, MA: MIT Press.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.
- Hill, C., & Schechter, J. (2007). Hawthorne’s lottery puzzle and the nature of belief. *Philosophical Issues*, 17(1), 102–122.
- Huber, F. (2012). Formal representations of belief. In *Stanford encyclopedia of philosophy*, Fall.

- Ichikawa, J. J. (2012). Knowledge norms and acting well. *Thought*, 1(1), 49–55.
- Ichikawa, J. J., & Jarvis, B. (2013). *The Rules of Thought*. Oxford: Oxford University Press.
- Ichikawa, J. J., Jarvis, B., & Katherine, R. (2012). Pragmatic Encroachment and Belief-Desire Psychology. *Analytic Philosophy*, 53(4), 327–343.
- James, W. (1897). *The will to believe*. New York: Longmans, Green and Company.
- Jarvis, B. (2012). The Dual Aspects Theory of Truth. *Canadian Journal of Philosophy*, 42, 209–233.
- Jarvis, B., & Katherine, R. (2015). The Evaluative and Normative Roles of Knowledge (Manuscript).
- Jeffrey, R. (1990). *The logic of decision*. Chicago, IL: University of Chicago Press.
- Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4), 575–603.
- Joyce, J. M. (2010). A defense of imprecise credences in inference and decision-making. *Philosophical Perspectives*, 24, 281–323.
- Kaplan, M. (1996). *Decision theory as philosophy*. Cambridge: Cambridge University Press.
- Kaplan, M. (2009). Williamson's casual approach to probabilism. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge*. Oxford: Oxford University Press.
- Kelly, T. (2013). Evidence can be permissive. In M. Steup & J. Turri (Eds.), *Contemporary debates in epistemology* (pp. 298–312). Oxford: Blackwell.
- Lange, M. (1999). Calibration and the epistemological role of Bayesian conditionalization. *Journal of Philosophy*, 96(6), 294–324.
- Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Lin, H. (2013). Foundations of everyday practical reasoning. *Journal of Philosophical Logic*, 42(6), 831–862.
- Lynch, M. P. (2009). *Truth as one and many*. Oxford: Oxford University Press.
- Makinson, D. (2005). *Bridges from classical to nonmonotonic logic*. London: King's College Publications.
- McDowell, J. (1994). *Mind and world*. Cambridge, MA: Harvard University Press.
- Moss, S. (2013). Epistemology formalized. *Philosophical Review*, 122(1), 1–43.
- Pritchard, D. (2010). Relevant alternatives, perceptual knowledge and discrimination. *Noûs*, 44(2), 245–268.
- Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.
- Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, 158(3), 277–297.
- Pryor, J. Uncertainty and undermining. Manuscript.
- Ramsey, F. (1931). Truth and probability. In R. Braithwhite (Ed.), *The foundations of mathematics and other logical essays* (pp. 156–198). London: Routledge and Kegan Paul.
- Reed, B. (2010). *A defense of stable invariantism*. *Noûs*, 44(2), 224–244.
- Ross, J., & Schroeder, M. (2012). Belief, credence and pragmatic encroachment. *Philosophy and Phenomenological Research*. doi:10.1111/j.1933-1592.2011.00552.x.
- Shah, N. (2003). How truth governs belief. *Philosophical Review*, 112(4), 447–482.
- Shah, N., & Velleman, J. D. (2005). Doxastic deliberation. *Philosophical Review*, 114(4), 497–534.
- Skyrms, B. (1980). Higher order degrees of belief. In D. H. Mellor (Ed.), *Prospects for pragmatism* (pp. 109–137). Cambridge: Cambridge University Press.
- Smith, M. (2010). What else justification could be. *Noûs*, 44(1), 10–31.
- Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives*, 13, 137–149.
- Sosa, E. (2011). *Reflective knowledge: Apt belief and reflective knowledge* (Vol. 2). Oxford: Oxford University Press.
- Staffel, J. (2012). Can there be reasoning with degrees of belief? *Synthese*, 190(16), 3535–3551.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Steglich-Peterson, A. (2006). No norm is needed: On the aim of belief. *Philosophical Quarterly*, 56(225), 499–516.
- Sturgeon, S. (2008). Reason and the grain of belief. *Noûs*, 42(1), 139–165.
- Turri, J. (2010). Epistemic invariantism and speech act contextualism. *Philosophical Review*, 119(1), 77–95.
- van Fraassen, B. (1983). Calibration: A frequency justification for personal probability. In R. S. Cohen & L. Laudan (Eds.), *Physics, philosophy and psychoanalysis* (pp. 295–319). Dordrecht: D. Reidel.
- van Fraassen, B. (1984). Belief and the will. *Journal of Philosophy*, 81(5), 235–256.
- Velleman, J. D. (2000). *The possibility of practical reason*. Oxford: Oxford University Press.
- Weatherston, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19(1), 417–443.
- Wedgwood, R. (2002). The aim of belief. *Noûs*, 36(16), 267–297.
- Wedgwood, R. (2007). *The nature of normativity*. Oxford: Oxford University Press.

- Wedgwood, R. (2012). Outright belief. *Dialectica*, 66(3), 309–329.
- Weisberg, J. Belief: Full and partial. Manuscript.
- White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 19(1), 445–459.
- White, R. (2009). Evidential symmetry and mushy credence. In T. Szabo-Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 3, pp. 161–189). Oxford: Oxford University Press.
- Whiting, D. (2010). Should I believe the truth? *Dialectica*, 64(2), 213–224.
- Williamson, T. (1998). Conditionalizing on knowledge. *British Journal for the Philosophy of Science*, 49(1), 89–121.
- Williamson, T. (2002). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, C. (1992). *Truth and objectivity*. Cambridge, MA: Harvard University Press.