

narrative of upward mobility. The distinctness of these ethical costs on the scales of social injustice begins to seem beside the point with respect to that most crucial take-away from this remarkable book.

Strivers' orientation towards interdependence plausibly makes them rich in relational goods. Material goods are clearly necessary for enjoying all the richness of interdependency, and strivers are clearly disadvantaged on that count. But their greater susceptibility to ethical *costs* does suggest that they are initially rich in something that matters deeply and incommensurably. So, although the most important beneficiaries of the ethical narrative are the strivers it portrays, it's worth making explicit one final upshot of this discussion: the ethical narrative might just help strivers *and* help their more materially advantaged peers by spotlighting the goods that an interdependent orientation would enable all of us more fully to appreciate.

*Harvard University
Cambridge, MA, USA
gschouten@fas.harvard.edu*

Settling the Unsettled: Roles for Belief

ELIZABETH JACKSON

1. Introduction

Uncertainty is a fact of life. We often form attitudes and make decisions from a position of uncertainty. Epistemologists capture this by appealing to *credences*, a measurement of our (un)certainly in propositions. On the standard way of modelling credences, 1 represents maximal certainty a proposition is true, and 0 represents certainty it is false. For example, while my credence in modus ponens is close to 1, my credence that a fair coin will land heads is 0.5 and my credence that it will snow in Tampa tomorrow is quite low.

While descriptive questions about credence are significant and interesting (see [Eriksson and Hájek 2007](#)), epistemologists have primarily focused on *normative* questions about credence: what makes a credence rational? According to an influential theory, Bayesianism, a credence function is rational only if it obeys two constraints: it is probabilistically coherent and updated by conditionalization.

While, at first blush, these Bayesian requirements on credences seem plausible, problems arise when they are applied to actual human thinkers. There is quite a bit of evidence from cognitive science that we are not rational Bayesian thinkers. Humans frequently make probabilistic errors, such as the base rate fallacy. And not only do we fail to meet the Bayesian requirements; these requirements seem implausibly strong. Consider: an agent with probabilistically coherent credences will assign credence 1 to all logical truths. But when it comes to an unproven mathematical theorem, most

human thinkers will (understandably) assign that a credence of 0.5. Is this really irrational?

Bayesians respond that their theory of rational credence is a theory of *ideal* rationality, and it may not straightforwardly apply to actual human thinkers. But then the obvious question is: what is it for an *actual human thinker* to have rational credences? This is one of the significant and underexplored questions that Julia Staffel addresses in *Unsettled Thoughts: A Theory of Degrees of Rationality* (2020).¹

More specifically, Staffel seeks to defend and precisify two main claims about Bayesian rationality. The first is that rationality comes in degrees, so one can be more or less rational. The second is that non-ideal thinkers should seek to approximate ideal rationality. To illustrate this, Staffel provides a helpful distinction between two kinds of goods: a winner-take-all good, like a dream job (there are no benefits to being the number-two candidate), and goods for which the closer you are to the ideal, the more benefits you glean, such as learning a language. The idea is that Bayesian rationality is not a winner-take-all case. Instead, the more rational you are, the more of the benefits of rationality you enjoy (2020: 4–5).

My goal in this critical notice is to summarize some of Staffel's main claims and arguments, and then comment further on some of the arguments in her later chapters, specifically chapters 8 and 9. In §2, I provide an overall summary of the book, and some general comments and observations. In §3, I explain and raise some questions about chapter 8, a chapter on different types of rationality. In §4, I explain and discuss the arguments from chapter 9, a chapter on the role that belief plays for non-ideal thinkers. I conclude in §5.

2. Summary and preliminary remarks

I will begin with an overview of the book. After introducing and explaining the aims of the book in chapter 1, chapter 2 provides an overview of Bayesian rationality. Staffel discusses probabilism and conditionalization, and other Bayesian norms such as the Principle Principle (one's knowledge of objective chances should constrain their credences), Indifference Principles (one should distribute their credence equally among symmetrical options) and Reflection Principles (one should defer to expert credences, including one's own future credences). Staffel mentions other normative concepts such as the Uniqueness Principle, higher-order evidence, Jeffrey conditionalization and more. The summary of these norms is very helpful, and I highly recommend chapter 2 for anyone interested in an overview of the Bayesian project.

Chapter 2 closes by explaining a crucial missing piece of the Bayesian framework: how does Bayesian rationality apply to non-ideal thinkers? Chapter 3 begins to answer this question by focusing on a central tenant of Bayesianism : coherence. Staffel considers several ways we might approximate coherence, and argues (in my view, convincingly) that we should prefer quantitative measures of coherence approximation to qualitative measures. She discusses four quantitative distance measures we might use: absolute distance, Euclidean distance, Chebyshev distance, and Kullback–Leibler divergence.

1 *Unsettled Thoughts: A Theory of Degrees of Rationality*. By Julia Staffel. Oxford University Press, 2020. 240pp

Staffel then turns to the question: why is it good to be coherent? The basic answer is that the more coherent one's credences are, the better they can perform their function. Staffel focuses on two functions of credence: guiding action and representing the world. In chapter 4, Staffel argues that reducing credal incoherence leads to less Dutch-book ability, meaning that more coherent credences will better guide action because they lead to fewer losses from Dutch books. In chapter 5, Staffel argues that there are certain ways of becoming coherent that guarantee improvements in accuracy. Thus, at least in these cases, more coherent credences better perform the function of representing the world than less coherent ones.

Chapter 6 turns to how we might approximate ideal rationality if thinkers are required to comply with norms in addition to coherence. Staffel discusses two main strategies: the *bundle strategy*, on which we take the principles of rationality together to determine the permissible credences, and measure incoherence as the distance from the nearest permissible credence function; and the *piecemeal strategy*, on which we measure the approximation to each rational norm separately, and then aggregate the results. Staffel discusses how these strategies fit with the view that there is a single epistemic good (e.g. accuracy) or multiple epistemic goods (e.g. accuracy, evidential fit, etc.). While this chapter does not commit to a single strategy or single view of epistemic goodness, Staffel ultimately suggests that the bundle strategy has significant advantages over the piecemeal strategy.

Chapter 7 considers thinkers who reason from irrational starting points. It considers three types of credal changes: changes without learning new information or adding new attitudes, changes without learning new information that involve adding new attitudes, and changes as a response to new evidence. Chapters 8 and 9 apply the previous results to other theories of rationality and the role of outright beliefs, respectively. We will discuss these chapters in more detail in the next two sections.

Overall, there is much to be said in praise of this book – there is not much more you can ask from a monograph. It is clear, well written and tightly argued; each chapter has a distinct thesis and a concise but informative overview of each section. Not only this, but the book is timely and needed; while Bayesian epistemology is popular and important, little work has been done on whether and how human thinkers should approximate Bayesian rationality. While, at points, the reader might feel that the book could be more committal, as much time is spent laying out various options (e.g. in chapters 3 and 6), I think this is appropriate, given the state of the debate – this book is about a vital but underexplored area of epistemic rationality, and much of the logical space has not been explored yet.

One of the things I most appreciate about the book is Staffel's use of formalism. Bayesian epistemology can be quite technical. While formal tools are extremely helpful and a key element of philosophical progress, they can also be overused, making simple concepts completely inaccessible, especially to those unfamiliar with the relevant formalisms. Staffel does an excellent job of using formalism without being overly technical. Thus I recommend Staffel's book to a wide audience, including traditional (or, as my friend calls them, 'casual') epistemologists.

I also enjoyed that this book touched on the permissivism/uniqueness debate in several chapters (e.g. chapters 6 and 7). Permissivism is the view that there are evidential situations that rationally permit more than one attitude toward a proposition. The relationship between permissivism and Bayesianism is interesting for several reasons.

First, traditional Bayesianism implies an interpersonally permissive view of rationality (see Meacham 2014); this is notable because permissivism has the reputation of setting a 'low bar' for rationality. However, Bayesian rationality is idealized rationality. The lesson is that there are many ways to be a permissivist, and not all set a low bar for rationality.

I also found Staffel's discussion of the relationship between permissivism and conditionalization noteworthy (ch. 7, sec. 2), as it bears on the question of whether Bayesianism is consistent with *intrapersonal* credal permissivism (permissivism involving a single agent's credences). The traditional interpretation of conditionalization is that thinkers should adjust their credences *when and only when* they get new evidence. This raises the question: does conditionalization rule out intrapersonal permissivism? Not necessarily; two possibilities are worth exploring. First, the intrapersonal permissivist could accept that there may be more than one way to rationally change one's credences in response to new evidence, depending on how one interprets or weighs the new evidence. Second, one could weaken conditionalization to the view that thinkers should adjust their credences *when* they get new evidence, but rational thinkers can sometimes adjust their credences in response to non-evidential changes (such as practical reasons, a change in stakes, a change in what hypotheses one is considering etc.). Generally, the prospects for credal intrapersonal permissivism merit further exploration – including the relationship between these credal changes and susceptibility to Dutch books.

3. *Varieties of rationality*

Now we will more closely examine two later chapters. In chapter 8, Staffel discusses varieties of rationality, focusing on five distinctions. Here, we will focus mainly on the relationship between propositional and doxastic rationality and the relationship between evaluative and ameliorative approaches to rationality.

An attitude is propositionally rational if and only if (iff) it is rational to hold it; an attitude is doxastically rational iff it is rationally held. Staffel uses this distinction to shed light on two cases. In case 1, a logic student considers an unknown tautology and forms a 0.5 credence in it. In case 2, a logic student considers an unknown tautology and forms a credence of 1 in it. This illustrates the problem of logical omniscience: according to Bayesian norms, you ought to have a credence of 1 in all tautologies, but this seems to require logical omniscience. And many will judge that the student with credence 0.5 is rational, and the student with credence 1 is irrational.

Staffel suggests that a potential solution to this problem involves the distinction between propositional and doxastic rationality. If we interpret the claim that it is rational to have a credence of 1 in tautologies as a claim about *propositional* rationality, then we can maintain that the student with credence 1 is propositionally, but not doxastically, rational. By contrast, the student with credence 0.5 is doxastically, but not propositionally, rational. They are both rational in some sense (154–55).

While Staffel's solution is attractive, it also has several noteworthy – and, for some, potentially unattractive – consequences. First, as Staffel notes (155) it is standardly thought that doxastic rationality entails propositional rationality. However, on this view, doxastic and propositional rationality crosscut each other.

Second, for this solution to help with the general problem of logical omniscience, it requires that each thinker's evidence supports every tautology to degree 1. And this seems to presuppose a particular view of evidential support. There are two general camps regarding evidential support: psychological and non-psychological views. On the psychological views, whether one's body of evidence supports a certain proposition p depends heavily on certain psychological facts, for example, whether one sees that their evidence supports p , whether one could derive p from their evidence, etc. On the non-psychological views, there is an objective evidential support relation that does not depend on a thinker's psychology.

To establish the claim that each thinker's evidence supports every tautology to degree 1, one must accept a view of evidential support that is fully non-psychological. On this view, the credences that are propositionally rational for you could be completely beyond your ken – things you could never reason to or figure out, even given unlimited time and resources. Of course, these camps form a continuum, and the fully psychological notions of evidential support might be problematic too – plausibly, there are cases where one is not directly aware of what their evidence supports. But the fully objective notion of evidential support seems to take things too far in the other direction, as it may be impossible for a thinker to determine what attitudes are propositionally rational for them. Thus, while I agree with Staffel that Bayesians should pay more attention to the distinction between propositional and doxastic rationality, the right view of evidential support may fall somewhere in between the extreme psychological view and the extreme non-psychological view. Therefore, in general, it is not clear to me that each thinker's evidence supports every tautology to degree 1.

Second, I want to address some of Staffel's remarks regarding the relationship between evaluative and ameliorative approaches to rationality. Evaluative approaches to rationality evaluate thinkers from a third-person point of view and do not purport to tell them how to be more rational. Ameliorative approaches take the first-person point of view and seek to guide thinkers towards rationality. Staffel is clear that she is interested only in the former: 'The theory of degrees of rationality I propose is an evaluative theory ... [it cannot] answer questions about how particular irrational thinkers should reason or change their credences' (160). She also remarks, 'There is thus no direct path from an evaluative theory of epistemic rationality to a theory that tells us how particular thinker[s] should change their credences in particular situations' (161).

While of course evaluative rationality is important in its own right, I was slightly disappointed that Staffel took her conclusions to have no implications for how a thinker should change their credences from the first-person point of view. Given that one of the goals of the book was to provide an account of rationality for non-ideal thinkers, it would be nice for these thinkers to be able to utilize this account to inform their reasoning and attitude formation.

Further, Staffel suggests that ameliorative questions about rationality are largely empirical, since they concern psychological questions about how people in fact reason and even pedagogical questions about how good reasoning should be taught (162). However, the question of what ameliorative norm applies given one's attitudes and situation seems different from the much more practical and empirical question of how we should teach and/or motivate people to be rational. Consider moral philosophy. Much of normative and applied ethics is about determining the true action-guiding

norms, but what ethicists are doing seems separable from practical questions about how we might, say, effectively teach children about morality in a classroom setting. Thus it seems possible to engage in a philosophical study of ameliorative norms (in both epistemology and ethics) that does not primarily involve empirical psychology. An application of Staffel's evaluative theory to ameliorative rationality strikes me, then, as a promising area of further philosophical research.

4. *Roles for belief*

Chapter 9 turns from credence to the role and rationality of outright belief. The chapter presents a puzzle about outright belief, then argues for a particular solution. Staffel assumes that humans have outright beliefs in addition to credences. She argues that this is because beliefs simplify reasoning by ruling out small error possibilities. In some contexts, it is appropriate to treat a claim in which one has a high credence as true, as this makes one's reasoning process much easier. Furthermore, ruling out a small error possibility does not result in a substantial loss of accuracy (see [Ross and Schroeder 2014](#), [Tang 2015](#), [Jackson 2019](#)).

On this picture, it is natural to think that whether we rely on our belief that p or our credence in p can vary with context. In some contexts (e.g. low stakes), we treat a claim as true, while in others (e.g. high stakes), we treat it as probable but uncertain. A strategy that has been proposed for how we might manage beliefs and credences across contexts is called pseudo-conditionalization. When a thinker uses pseudo-conditionalization, they conditionalize on p , but not because they learn p with certainty, but because the context allows them to treat p as true. However, pseudo-conditionalization is very computationally demanding. If the role of outright belief is to simplify reasoning, it is implausible that managing outright beliefs would require such a complex strategy. Pseudo-conditionalization, then, seems to undermine belief's primary functional role.

Staffel summarizes the puzzle by noting that the following four claims are all plausible, but we need to reject one (176):

- (1) Human thinkers have outright beliefs in addition to credences because outright beliefs help simplify reasoning processes.
- (2) Reasoning processes can involve mixtures of credences and outright beliefs, and it is flexible from context to context which outright beliefs are relied upon.
- (3) Outright beliefs and credences in a context are determined by pseudo-conditionalizing on a set of background credences.
- (4) Pseudo-conditionalizing is difficult to execute for human reasoners, because it is computationally expensive.

Ultimately Staffel suggests we reject claim 3. Claim 3 undermines belief's functional role of simplifying reasoning. Staffel importantly points out that, when trying to determine the norms governing a mental state, we must consider that state's functional role (188). So if the functional role of belief is to simplify reasoning, then the norms for belief should complement that role – we should not posit super-computationally demanding norms for belief. (In contrast, the norms for credence can be quite demanding, since the functional role of credence is not simplification, but accuracy.) Instead of pseudo-conditionalizing, Staffel suggests we use a simpler strategy for managing beliefs

and credences. When we rely on our belief that p , we treat p as if it had probability 1 in that context. However, we need not adjust our credences in all other claims accordingly, as pseudo-conditionalizing requires. This can lead to incoherence – both within a context and across contexts. Nonetheless, the incoherence is only slight (190ff).

I want to explore the possibility of denying another claim in the puzzle – claim 1. Staffel motivates claim 1 by objecting to two other roles that belief might play; I will argue that her objections are not ultimately compelling. Then I will explain my main worry for the claim that the only role for belief is that of simplifying reasoning. Finally, I will suggest some other roles for belief that vindicate Staffel's assumption that we have both beliefs and credences.

Staffel motivates claim 1 by considering views on which the main function of belief is something other than simplifying reasoning. She considers two alternative functions. The first is that outright beliefs are necessary as a basis for moral judgement; the second is that outright beliefs are necessary for knowledge. We will take them in reverse order, and I will argue that Staffel's reasons for rejecting each alternative are not decisive.

Consider the idea that outright beliefs play a role in knowledge. It is widely thought that knowing p requires believing p . Then, if knowledge is valuable, outright belief is valuable too. Staffel rejects this role for belief based on the way she defines beliefs and credences. She defines a credence as an attitude that encodes uncertainty, and a belief as an attitude that does not encode uncertainty (184). Given this, credences can constitute knowledge because probability-beliefs *are* credences. There are two problems with this, however. The first is that this is a non-standard way of understanding beliefs and credences. When it comes to the simplifying role of belief, it does seem that probability-beliefs and credences function similarly (neither rules out the possibility of not- p). At the same time, most think that probability-beliefs and credences are distinct, except those who think credences *reduce to* probability-beliefs (see [Moon and Jackson 2020](#)), but Staffel explicitly assumes that belief and credence are irreducible. Then it is plausible that probability-beliefs are beliefs, not credences. Second, even if we accept Staffel's non-standard way of understanding beliefs and credences, it also does not seem that knowledge always encodes uncertainty. One can know that p is flat-out true. This seems difficult to explain if the attitude involved in knowledge always encodes uncertainty. For these reasons, even though probability-beliefs can constitute knowledge, this does not seem like a satisfactory reason to reject this role for belief.

Second, Staffel considers the idea, proposed by [Buchak \(2014\)](#), that beliefs are necessary as a basis for moral judgement. High credences can be based on statistical evidence, but we should not praise or blame someone based on statistical evidence. For example, Buchak provides a case in which you know either Barbara or Jake stole your cell phone, and you know men are ten times more likely to steal cell phones than women. You should not believe Jake stole it or blame Jake for theft, even though you should have a high credence he is a thief ([Buchak 2014](#): 292).

Staffel is sympathetic to Buchak's claim that mere statistical evidence is not a basis for rational belief, and that we ought not blame someone based on mere statistical evidence. However, Staffel considers a case where we have high credence based on non-statistical evidence that someone transgressed. In this case, Staffel says it is unclear that the credence cannot play a role in holding her responsible. If credence can underlie moral judgement, 'why is outright belief needed as a middleman?' (184). Staffel further notes that legally, in civil cases, the preponderance of evidence standard applies, on

which someone can be convinced if the evidence makes it more than 50% likely that they are guilty, and the evidence is not purely statistical. This further supports Staffel's contention that high credence *based on the right kind of evidence* can play a role in blame, even if high credence based on mere statistical evidence cannot (184).

While Staffel's objection is thought-provoking, there are responses available to Buchak. Buchak might simply disagree with Staffel's contention that high credence *based on the right kind of evidence* can underlie moral judgement, as Buchak suggests that mere credence is always insufficient for blame. We take a stand on whether someone is guilty, then blame them in proportion to the severity of their wrongdoing. In our blaming behaviours, we do not (and should not) appeal to our credences that the offender is guilty, whether our credence is based on statistical evidence or not.

This blame-requires-belief view seems to offer a simpler explanation of the data, as opposed to the view on which blame requires a high credence, and that credence's being based on both the right kind and an adequate degree of evidence. On Buchak's view, every time we rightfully blame someone, we believe they are guilty. Staffel's alternative view, on which some, but not all, high credences play a role in blaming, begins to look complex and even potentially ad hoc. Therefore Buchak's view provides a clean, unified explanation of the data: beliefs are the states that enable us to rightfully blame.

Finally, concerning Staffel's point on the preponderance of evidence standard, it may be that sometimes moral and legal norms come apart. We might have a practical reason (e.g. deterrence) to use the preponderance of evidence standard in court, but that does not mean we ought to blame someone just because our credence in their guilt is slightly above 0.5 (and based on non-statistical evidence).

Generally, then, I am not convinced that we should rule out these alternative roles for belief. Furthermore, the claim that beliefs only simplify reasoning leads to its own 'Bayesian challenge' involving belief and acceptance. Namely: why posit beliefs in these cases, when it seems as though acceptances can do all the relevant work? One accepts p iff one acts as if p is true. Belief is a *mental state*, primarily sensitive to epistemic factors. Acceptance is a *policy*, motivated by practical factors. But in the cases that motivate the simplifying role of belief, one treats a proposition p as if it has probability 1 for practical, rather than epistemic, reasons. One's evidence does not establish p to probability 1, but one treats p as if it has probability 1 to make one's reasoning easier. This looks much more like acceptance than belief; it is not clear that we have a reason to posit an additional mental state.

Staffel addresses this worry (chapter 9, n.5) and suggests that acceptances are more voluntary than beliefs. However, there are some recent compelling arguments that beliefs are sometimes under our voluntary control (see [Roeber 2020](#)). Furthermore, it does seem as though we can control what possibilities we consider in decision-making, especially since reasoning is an active and deliberate process. A second potential response is that beliefs rule out small error possibilities, whereas acceptance can justify rounding up to 1 in cases where the error possibilities are much larger. I agree that acceptance can sometimes justify rounding p up to 1, even if the probability of p on your evidence is quite low (see the cases in [Jackson forthcoming](#)). However, if one uses this method to distinguish belief from acceptance, this raises the question: how much can we 'round up' for the attitude to still count as a belief? Can we round up to belief from 0.6 credence? Or from 0.5 credence? The answer seems to depend on what is at

stake, but this is exactly what we would expect if all of these roundings were acceptances (a practical policy) rather than beliefs (an epistemic attitude).

This may push us, especially those of us inclined to belief-credence dualism, to think that the main function of belief is not simplification. In addition to playing a role in knowledge and playing a role in our moral judgements, other dualists, including Ross and Schroeder (2014) and Buchak (forthcoming), have suggested that beliefs play a related cluster of additional roles: allowing us to take a stand, have a view of the world, and represent the world such that certain claims are true. It is not clear that credences can play this role, since sometimes we take a stand on uncertain claims. Buchak (forthcoming) argues that this provides a novel view in the epistemology of disagreement. If we conciliate with our credences, but remain steadfast in our beliefs, we can take a stand on the truth of certain claims – especially significant moral, political or religious matters – but lower our credences as we encounter counterevidence. Our credences, then, function to track our precise level of evidential support for a claim, and help determine when we should give up a belief. At the same time, beliefs allow us to take a stand on our central life commitments, even in the face of counterevidence.

This discussion raises the question: how do these roles of belief fit together? If beliefs allow us to take a stand, underlie our moral judgement and play a role in knowledge, this suggests that we would have beliefs even if we did not have cognitive limitations. However, if beliefs only simplify reasoning, ideal reasoners might not have beliefs (at least it is unclear what role beliefs would play; they would appear superfluous). Thus there may be a tension between these roles of belief: on the simplifying view, it is very natural to think that beliefs are a product of our cognitive limitations, but, on the other views, the reason we have beliefs is not due to our limitations, so ideal reasoners would have beliefs.

One possible response to this tension is to embrace one role for belief but not another. Staffel may be sympathetic to the view that the only (or at least primary) role of belief is simplifying reasoning. A second option is to say that beliefs play one or several of the other roles – taking a stand, being necessary for knowledge, underlying moral judgements – but to deny that beliefs play a simplifying role in our reasoning. On this view, it is natural to think that acceptances, but not beliefs, play the simplifying role that Staffel describes. A third option is to embrace all the potential roles for belief. Maybe some of our beliefs play the role of simplifying reasoning, and other beliefs underlie moral judgements or enable us to take a stand. Maybe some beliefs even play both roles. On this third option, ideal reasoners might have some beliefs, but may have fewer beliefs than ordinary human thinkers. These possibilities all merit further exploration.

Chapter 9 nonetheless makes a lot of interesting moves, especially on the assumption that beliefs simplify reasoning. And if beliefs (or acceptances) simplify reasoning, I am convinced that they probably do not do so via pseudo-conditionalization.

5. Conclusion

This book covers a lot of ground and opens up many fruitful areas of further research. Chapter 10 is especially helpful in this regard, as it outlines the main themes of each chapter and where those themes might be further explored. I highly recommend this book to a large array of people, including those working on rationality, especially bounded rationality and degrees of rationality, those working on norms for credences,

those working on Dutch book and accuracy justifications for probabilism, those working on diachronic rationality (see especially ch. 7), and those working on belief and credence (see especially ch. 9). The book will also be of interest to those exploring questions about irrationality, especially questions about what attitudes irrational agents should have and how one should reason from irrational starting points. Finally, those interested in empirical questions regarding Bayesianism, for example how we succeed and fail to approximate Bayesian rationality according to cognitive science, should read this book.²

Ryerson University
Toronto, ON, Canada
lizjackson111@gmail.com

References

- Buchak, L. 2014. Belief, credence, and norms. *Philosophical Studies* 169: 285–311.
- Buchak, L. Forthcoming. A faithful response to disagreement. *Philosophical Review*.
- Eriksson, L. and A. Hájek. 2007. What are degrees of belief? *Studia Logica* 86: 183–213.
- Jackson, E. 2019. How belief-credence dualism explains away pragmatic encroachment. *Philosophical Quarterly* 69: 511–33.
- Jackson, E. Forthcoming. Belief, faith, and hope: on the rationality of long-term commitment. *Mind*.
- Meacham, C. 2014. Impermissive Bayesianism. *Erkenntnis* 79: 1185–217.
- Moon, A. and E. Jackson. 2020. Credence: a belief-first view. *Canadian Journal of Philosophy* 50: 652–69.
- Roeber, B. 2020. Permissive situations and direct doxastic control. *Philosophy and Phenomenological Research* 101: 415–31.
- Ross, J. and M. Schroeder. 2014. Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research* 88: 259–88.
- Staffel, J. 2020. *Unsettled Thoughts: A Theory of Degrees of Rationality*. Oxford: Oxford University Press.
- Tang, W.H. 2015. Belief and cognitive limitations. *Philosophical Studies* 172: 249–60.

2 Thanks to Julia Staffel and Lara Buchak for helpful discussion about the content of this review.