Episteme, 14, 1 (2017) 39–48 © Cambridge University Press
doi:10.1017/epi.2016.51

ACCURACY FOR BELIEVERS

JULIA STAFFEL
staffel.julia@gmail.com

ABSTRACT

In *Accuracy and the Laws of Credence* Richard Pettigrew assumes a particular view of belief, which states that people don't have any other doxastic states besides credences. This is in tension with the popular position that people have both credences and outright beliefs. Pettigrew claims that such a dual view of belief is incompatible with the accuracy-first approach. I argue in this paper that it is not. This is good news for Pettigrew, since it broadens the appeal of his framework.

INTRODUCTION

In his book *Accuracy and the Laws of Credence*, Richard Pettigrew offers a comprehensive defense of what is often called "accuracy-first epistemology." The basic idea is appealingly simple: Assume that having accurate credences is fundamentally and exclusively epistemically valuable. Avail yourself of standard decision theoretic principles to determine which doxastic states are epistemically best. Derive as many commonly endorsed principles about rational credences as possible. The more you can get, the better for the view that accuracy is the fundamental epistemic good. The list of principles Pettigrew derives is quite impressive: probabilism, a chance-credence principle, an indifference principle, and plan conditionalization.

In making his arguments, Pettigrew assumes a particular view of belief, which states that people don't have any other doxastic states besides credences. This is in tension with the popular position that people have both credences and outright beliefs. Pettigrew claims that such a dual view of belief is incompatible with the accuracy-first approach. I argue in this paper that it is not. This is good news for Pettigrew, since it broadens the appeal of his framework.

In section 1, I explain the problem that the dual view of belief causes for the accuracy argument for probabilism. In section 2, I briefly state some of the motivations for holding a dual view of belief. In section 3, I explain how a dual view of belief can be reconciled with accuracy-first epistemology in a way that preserves the standard accuracy arguments. In section 4, I point out some open questions and avenues for further research.

1. ACCURACY AND THE PROBLEM OF OUTRIGHT BELIEF

In the current literature on the nature of belief, people distinguish between degrees of belief and outright beliefs. Degrees of belief, or credences, are fine-grained states. An agent can have any degree of belief in a proposition ranging from certainty that it is

https://doi.org/10.1017/epi.2016.51 Published online by Cambridge University Press

false to certainty that it is true. By contrast, outright beliefs are coarse-grained states. An agent can believe something, disbelieve it, or suspend judgment about it. It is commonly assumed that agents can outright believe claims they are not completely certain of. For example, I might have an outright belief that I will make turkey for Thanksgiving, even if I am not completely certain that I will do so. Philosophers disagree about which of these kinds of belief are genuine mental states. Some authors claim that our doxastic attitudes are exhausted by our credences, and that belief-talk is just a shorthand way of describing people's credences (Jeffrey 1970; Foley 2009; Lyon 2014; Pettigrew 2016). Others deny that people have credences, and claim that our minds contain only outright beliefs (Harman 1986). A third group claims that both credences and outright beliefs are genuine and distinct mental states (e.g. Lin and Kelly 2012; Wedgwood 2012; Clarke 2013; Lin 2013; Weisberg 2013, Forthcoming; Buchak 2014; Ross and Schroeder 2014; Easwaran and Fitelson 2015; Greco 2015; Sturgeon 2015; Tang 2015; Leitgeb 2016; Staffel 2016; Weatherson 2016). This third view is the most popular view in the current literature, and it presents a challenge for accuracy-first epistemology.

Accuracy-first epistemology is formulated as a theory of the epistemic value of and rational requirements on credences. This is appropriate if we assume that only credences are genuine doxastic attitudes, but not if we adopt a dual view of belief, according to which people have credences and outright beliefs. For accuracy-first epistemology to have far-reaching appeal, it should be able to accommodate this popular view of belief. Unfortunately, as Pettigrew argues in his book and in more detail in a separate paper, the accuracy-first view doesn't seem to combine well with the dual view of belief (Pettigrew 2015, 2016). The accuracy-dominance argument for probabilism seems to fail if we adopt the dual view of belief.

The argument works as follows on the credence-only view: Assume that the most accurate credence to have in a truth is 1, and the most accurate credence to have in a falsehood is 0. The more your credences approximate the truth, the less inaccurate they are. The distance between your credences and the corresponding truth-values at a given world is measured by a continuous, strictly proper scoring rule, such as the Brier score.[1] If an agent has non-probabilistic credences, then there is an alternative probabilistic credence function she could adopt that is more accurate in every possible world. By contrast, if an agent has probabilistic credences, then there is no alternative credence function that is at least as accurate as the agent's credences in every world, and more accurate in at least one world. From these results it is concluded that having probabilistic credences is necessary for being rational.

An example will help illustrate this.[2] Suppose Craig has the following credences: Cr(p) = 0.7 and Cr(∼p) = 0.6. Cr is incoherent, because Craig's credences in p and ∼p don't sum to 1. We can compute the inaccuracy of Craig's credences for the p-world and the ∼p-world, and discover that there are alternative coherent credence functions

1 Suppose c is a credence function defined over a set of propositions F, and the function I_w indicates the truth values of the propositions in F at world w by mapping them onto {0,1}. Then the following function gives us the Brier score of c at w:

$$Brier(c, w) = \sum_{A \in F} (c(A) - I_w(A))^2$$

2 This example is the same as Pettigrew's (2015).

https://doi.org/10.1017/epi.2016.51 Published online by Cambridge University Press

he could adopt, such as $Cr^{alt}(p) = 0.55$ and $Cr^{alt}(\sim p) = 0.45$, that are more accurate in both the p-world and the ~p-world. Moreover, if Craig adopted a probabilistic credence function such as Cr^{alt} that dominates his existing credences, his new credences would no longer be accuracy-dominated.

Pettigrew invites us to consider cases in which agents can have outright beliefs in addition to credences. He assumes that for an agent to have a rational outright belief in a claim p, her credence in p must be at least as high as some specified threshold t. In order to measure the accuracy of an entire doxastic state, we have to add scores for outright beliefs. Pettigrew suggests the following general scoring method for doxastic states containing both credences and beliefs (2016: 54):

(i) True beliefs and false disbeliefs get a reward of R, where $R > 0$.
(ii) False beliefs and true disbeliefs get a penalty of W, where $W > 0$.
(iii) The total inaccuracy of an agent's doxastic state is computed by taking the inaccuracy score of the agent's credence function, adding a penalty of W for every false belief and true disbelief, and subtracting a reward R for every true belief and false disbelief.

For the dominance argument to still work once we include outright beliefs, it must be true that for any doxastic state d that contains incoherent credences or irrational beliefs, there is a doxastic state d* that contains only coherent credences and rational beliefs, and that at least weakly accuracy-dominates d.[3] But, as Pettigrew demonstrates, this is false. There are cases in which an irrational agent has outright beliefs that a rational agent cannot have, and the accuracy bonus the irrational agent gets from having these beliefs in worlds in which they are true outweighs the accuracy bonus the rational agent gets in those worlds for being rational.

We can build on the example of Craig to sketch in a bit more detail how this can happen. Readers who are interested in the proof of the result are invited to consult Pettigrew (2015). Suppose as above that Craig's credences are $Cr(p) = 0.7$ and $Cr(\sim p) = 0.6$. Additionally, we'll assume that he has an outright belief in p. His current accuracy scores are determined as follows:

> At the p-world, Craig is penalized for his 0.7 credence in p, and his 0.6 credence in ~p, since these credences are not maximally accurate. He is rewarded for having a true belief in p. The total score is the sum of the reward and penalties.
>
> At the ~p-world, Craig is penalized for his 0.7 credence in p, and his 0.6 credence in ~p, since these credences are not maximally accurate. He is additionally penalized for having a false belief in p. The total score is the sum of the penalties.

Let's now try to come up with a rational doxastic state that is at least as accurate as Craig's current state in every world, and more accurate in at least one world. In order to determine when an outright belief is rational, we need to specify a threshold, which

3 For a doxastic state d* to weakly accuracy-dominate Craig's current doxastic state d, d* must be at least as accurate as d in every world, and more accurate than d in at least one world. Strong dominance requires that d* is more accurate than d in every world.

https://doi.org/10.1017/epi.2016.51 Published online by Cambridge University Press

we will assume to be 0.7. Craig's credence function Cr is incoherent, and so we know that there are alternative coherent credence functions that are more accurate in every world. It turns out that all of these more accurate coherent credence functions assign a credence to p that lies below 0.7. Thus, they all assign a credence to p that is below the threshold for rational belief, which means that none of these coherent credence functions are compatible with a rational belief in p. Let Cr* be one of these coherent credence functions that is more accurate than Craig's current credences in both worlds. Does Cr* weakly accuracy-dominate Craig's current doxastic state? Cr* is definitely better than Craig's current doxastic state in the ~p-world, because Cr* is more accurate than Cr, and Cr* does not get an added penalty for a false belief in p. But Cr* is worse than Craig's current doxastic state in the p-world, because, as Pettigrew demonstrates, the gain in accuracy resulting from changing the credences in p and ~p can't make up for the loss of the reward that results from Craig's true belief in p.

What if we instead try a coherent credence function that permits a rational belief in p, such as Cr**(p) = 0.7 and Cr**(~p) = 0.3? While this doxastic state is more accurate than Craig's actual doxastic state in the p-world, it is worse in the ~p-world. This is because this state differs from Craig's actual state only in the credence in ~p, and in the ~p-world Cr**(~p) = 0.3 is less accurate than Craig's credence of 0.6 in ~p. Any other doxastic state we can try has the same problem: if it is rational, then it is worse than Craig's current doxastic state in at least one world.

As Pettigrew (2015) shows, this problem can't be solved. We can find examples of irrational doxastic states that are not dominated by any rational doxastic states even if we switch to a different strictly proper scoring rule, or change the threshold for belief, or change the values for R and W.

Pettigrew's own view of belief escapes this problem, because it doesn't acknowledge outright beliefs as genuine mental states. On his view, which I will call *linguistic reductivism* (following Lyon 2014), outright belief talk is just that – a way of talking. Saying that Robin believes that birdwatching is fun is just a shorthand way of describing her credences. Belief-talk is like talk about tallness: Calling someone tall doesn't ascribe an additional property to them over and above the fact that their height is x inches. Similarly, saying someone believes something doesn't ascribe an additional mental state to them over and above their degree of confidence. And if beliefs aren't in the head, we don't need to include them in evaluations of the rationality of an agent's attitudes. Hence, a proponent of linguistic reductivism need not worry about the problems a dual view of belief poses for accuracy arguments.

2. AGAINST LINGUISTIC REDUCTIVISM

Many epistemologists reject linguistic reductivism in favor of a dual view of belief. They argue that cognitively limited agents like us need outright beliefs, because they simplify our reasoning. In adopting outright beliefs, an agent takes the believed claims for granted in her reasoning, which frees her from having to pay attention to small error probabilities. But for beliefs to play this simplifying role, they need to be mental states that people can reason with, which is at odds with linguistic reductivism (versions of this argument are endorsed by, among others: Lin 2013; Lin and Kelly 2012; Wedgwood 2012; Ross and Schroeder 2014; Tang 2015; Weatherson 2016; Weisberg Forthcoming).

https://doi.org/10.1017/epi.2016.51 Published online by Cambridge University Press

How might a linguistic reductivist like Pettigrew respond? One possible way of resisting this argument is to point out that it appeals to the cognitive limitations of non-ideal reasoners. The aim of accuracy-first epistemology is to give principles of ideal rationality. If ideal epistemic agents have no cognitive limitations, they need not simplify their reasoning, thus they need not have outright beliefs. They only treat claims as true that they are certain of, which is already captured in their credence functions. Hence, a theory of ideal rationality doesn't need to make room for mental states that are only attributable to non-ideal agents.[4]

Unfortunately, this response is problematic. Principles of rationality are intended for evaluating the doxastic states of agents like us. It isn't particularly interesting to have a theory of epistemic rationality that only applies to ideal reasoners. If one's epistemic theory is too idealized, it is in danger of losing contact with the agents and doxastic states it is intended to evaluate. If agents like us in fact have two kinds of belief that we rely on in reasoning, then both of these kinds of belief should be evaluable as being rational or irrational by a fitting theory of epistemic rationality.[5]

This point is appreciated in other areas of formal philosophy that aim to give rational evaluations, such as decision theory (Lance 1995; Joyce 1999: section 2.6; Titelbaum 2013: 74). When an agent makes a decision, she needs to settle on a way of framing the decision problem she faces. The same decision problem can be framed in more fine-grained or coarse-grained ways, depending on how finely the agent divides up the available actions and the possible states of the world she might find herself in. For example, suppose Katarina is deciding whether to go ice-skating on a frozen lake. She is considering the outcomes of skating depending on whether the ice is thick or thin. If the ice is thick, skating will be fun, and if it's thin, it won't be fun at all, because she'll fall into the cold water. For Katarina's decision to go skating to be rational, the expected value of going skating must outweigh the expected value of not skating.

In framing the problem this way, Katarina takes it for granted that the thickness of the ice is the only thing about the state of the world that matters for evaluating the benefits of skating or not skating. She's ignoring other factors, for example the possibility that she could fall on the ice and injure herself. Framing the decision-problem in this coarse-grained way makes it a so-called small-world decision problem. By contrast, a so-called grand-world decision problem is maximally fine-grained, and completely exhausts the possibilities and actions the agent can distinguish between.

Ideal agents with unlimited computational resources might be able to only rely on grand-world decision problems. But this is impossible for agents like us. We need to rely on small-world problems all the time. For decision theory to be a useful normative theory that can help us evaluate the decision-making of agents like us, it needs to be applicable to small-world decision problems. It would be seriously misguided to argue that decision theory need only be applicable to grand-world problems, since ideal reasoners never consider small-world problems. This would render decision theory essentially

4 Of course, one might instead respond by denying that agents like us need to rely on outright beliefs to simplify our cognitive lives. It would be interesting to see how one might argue for this, but I don't have space here to explore this line of reasoning.

5 For a detailed and insightful discussion of normative modeling, see Mike Titelbaum's work on the topic (2013, Forthcoming).

https://doi.org/10.1017/epi.2016.51 Published online by Cambridge University Press

useless for evaluating decisions made by humans. The same point applies to a theory of epistemic rationality.

3. ACCURACY-FIRST EPISTEMOLOGY WITHOUT LINGUISTIC REDUCTIVISM

We thus face the following problem: how can we reconcile the argumentative strategies and results from accuracy-first epistemology with a dual view of belief? We know from Pettigrew's argument that scoring an agent's beliefs and credences jointly invalidates the accuracy-dominance argument for probabilism (and presumably other accuracy-based arguments as well, although Pettigrew doesn't demonstrate this). I will now explore a solution to this problem that Pettigrew doesn't consider. The distinction between small-world and grand-world decision problems can help us here.

In any decision situation, we have to choose one consistent way of framing our decision problem. We cannot treat a factor as both relevant and not relevant in a single way of framing a decision. Keeping this in mind, let's consider how beliefs and credences enter into our deliberations. Suppose I am currently highly confident, say 95% confident, that I am in good health. I also outright believe that I am in good health, in the sense that when I make plans and decisions, I usually take it for granted that my health won't interfere. The possibility that I might not be healthy is not a live possibility to me in those situations. However, there are some decision scenarios in which I don't, or wouldn't take it for granted that I am in good health. For example, in signing up for health insurance, I explicitly consider the small probability of major health problems. The important thing to observe is that in any episode of reasoning, I can rely either on my credence or on my belief in some claim p, but not on both. I can either take p for granted, or treat p as uncertain, but not both (see e.g. Lance 1995; Ross and Schroeder 2014; Weatherson 2016; Weisberg Forthcoming).

This suggests a particular way of thinking about how to model an agent's doxastic state. Pettigrew proposes to include all of the agent's beliefs and credences when calculating their accuracy score, presumably because these are all the doxastic states the agent has in some sense available to her. But there's another way of thinking about which of the agent's beliefs should be included when evaluating whether her doxastic state is rational. We can model the agent's doxastic state as the combination of attitudes she relies on in a particular context or situation. Since agents only ever rely on either their credence or their belief in a particular claim in a given context of reasoning, only one or the other will be included in a model of her doxastic state. This way of modeling an agent's doxastic state strikes me as no less natural than the one Pettigrew accepts. It mirrors the idea that we can represent a decision problem as a grand-world problem or a small-world problem, but not both at once. Moreover, it captures the idea that outright beliefs have value for us, because we need them to reason efficiently. But we only need them – and thus their value only becomes relevant – when we can't or don't use our credences. Hence, since we never simultaneously rely on a credence and a belief in the same claim, we don't need to score them together, since only the value of one or the other is relevant in that context.

Notice that this proposal does not claim that agents constantly change their attitudes, for example by shifting whether or not they believe something between contexts. That would not be a very attractive proposal. Instead, the proposal suggests that agents have both credences and outright beliefs at their disposal, and one or the other can be relied on in any given context.

https://doi.org/10.1017/epi.2016.51 Published online by Cambridge University Press

Here's how we might implement this proposal formally: Given that an agent who is relying on a belief in p in a context treats p as true in that context, it is a natural suggestion to represent this belief as Cr(p) = 1 in that context. For example, in a context in which I rely on my belief that I am healthy, my attitude would be represented as Cr(Julia is healthy) = 1, whereas in a context in which I rely on my credence, my attitude would be represented as Cr(Julia is healthy) = 0.95. Several authors in the literature have suggested views along these lines. Clarke (2013) and Greco (2015) argue that we should identify a full belief with a credence of 1 in a proposition, but allow that whether a proposition is assigned credence 1 can vary between contexts. Similarly, Wedgwood (2012) argues that reasoners like us have theoretical credences (the credences that we adopt purely in light of our evidence), and practical credences. The latter are simplified versions of our theoretical credences that allow us to reduce the complexity of reasoning problems. An outright belief is identified with a practical credence of 1. Tang (2015) also defends the view that outright beliefs can be identified with high probability estimates that get rounded up to 1. Of course this is not the only way in which one might formally capture the proposal that we should only include an agent's credence *or* her outright belief in a given claim p in a model of her doxastic state in a context, but this proposal relies on a framework that is already familiar.

This way of modeling doxastic states gets us at least some way towards reconciling accuracy-first epistemology with a dual view of belief. Since agents' attitudes are modeled relative to contexts, and each model only includes either a belief or a credence towards a proposition, but not both, it is not possible to devise a counterexample to the accuracy-argument for probabilism like the case of Craig. Moreover, since a belief is represented as a credence of 1 in a particular context, and we are thus using a standard modeling framework to capture the dual view of belief, we can easily run all of the familiar accuracy-based arguments to show that an agent's doxastic state in a context should obey rational norms. This way of modeling doxastic states can be applied to both ideal and non-ideal agents. Relativizing our models of doxastic states to contexts doesn't hinder our ability to model ideal agents' doxastic states, and it enables us to model non-ideal agents' doxastic states who need to rely on outright beliefs in many situations.

4. OPEN QUESTIONS

We've thus found a way to combine accuracy-arguments with a dual view of belief that doesn't run into the kind of counterexamples Pettigrew describes, and that relies on a natural assumption about which doxastic attitudes to include in epistemic evaluations. However, this proposal is not without problems and it raises important new questions. First, we need to know which outright beliefs are rational for an agent to adopt. There are numerous views on this in the literature, some of which are based on accuracy considerations (Easwaran 2015; Easwaran & Fitelson 2015).

Once we know what an agent believes, we also need to know in which contexts the agent should rely on her beliefs, and in which contexts she should consult her credences. This is a question that likely can't be answered solely on accuracy-based grounds. We can find a discussion of a related issue in the literature on pragmatic encroachment (see e.g. Hawthorne 2004; Stanley 2005; Ross and Schroeder 2014). Pragmatic encroachment

https://doi.org/10.1017/epi.2016.51 Published online by Cambridge University Press

about knowledge is the view that practical considerations can make a difference to whether an agent's true belief is a case of knowledge or not, and thus to whether the belief can be asserted and acted upon. A related hypothesis for our purposes is that whether or not an agent should rely on her credences or on her beliefs can depend on practical matters. For example, in decision scenarios where the stakes are high, or errors can be very costly, it seems often better to rely on one's credences. When I have to decide whether to purchase health insurance or not, or whether to go on an expedition that would cut me off from civilization for a long period, it is appropriate to explicitly take into consideration the small probability that I might be or become sick. More generally, it seems plausible on the dual view of belief that different combinations of the agent's available beliefs and credences will be appropriate to rely on in different contexts. We need to know how agents can and should handle selecting the appropriate combination of attitudes.

The claim that agents can switch between relying on credences and relying on beliefs in different contexts also raises the question of how agents can ensure that the set of attitudes used in a given context is coherent. Answering this question seems especially tricky given that our response should be compatible with the claim that the function of outright beliefs is to simplify our reasoning. If anything, it would seem that switching between beliefs and credences introduces complications that make it harder to be coherent. Clarke (2013) and Weisberg (Forthcoming) address this problem to some extent.

5. CONCLUSION

We started by asking how the dual view of belief can be made compatible with the framework of accuracy-first epistemology. Pettigrew argues that this is not possible. He proves that, if we measure the accuracy of an agent's credences and outright beliefs at the same time, the argument for probabilism doesn't go through. I argued that it is equally, if not more plausible, to adopt a different way of measuring the accuracy of an agent's doxastic state. Since an agent only ever relies on her credence or her belief in a proposition in a given context, but not on both, only one of them should be included when measuring the accuracy of the agent's doxastic states in that context. If we adopt this way of modeling and scoring agents' doxastic states, all the standard arguments in accuracy-first epistemology still go through.

While the dual view of belief is very popular in the literature, there are many open questions the view must address, such as which beliefs are rational for an agent to adopt, and how she should select which attitudes to rely on in a given context of deliberation. Moreover, it needs to be explained how agents can switch between different combinations of beliefs and credences in a way that preserves the purported simplifying role of having beliefs in addition to credences. Yet, these questions arise independently of whether the dual view of belief is combined with the accuracy-framework or not. Hence, they don't cause problems for the view I am proposing here any more than they do for the dual view on its own. Accuracy-first epistemology, on this view, provides rational constraints on doxastic states in a context. Additionally, it could provide constraints on which outright beliefs to adopt, given a particular rational credence function. Accuracy-first epistemology wouldn't be tasked with answering the question of whether to rely on a belief or a credence in a given context, because this is plausibly not a purely epistemic question.

https://doi.org/10.1017/epi.2016.51 Published online by Cambridge University Press

Consequently, accuracy-first epistemology has a broader appeal than Pettigrew recognizes. Its arguments can be endorsed not only by linguistic reductivists about belief, but also by proponents of the far more popular dual view of belief.

ACKNOWLEDGEMENTS

I would like to thank Brian Talbot, Richard Pettigrew, and an anonymous referee for helpful comments.

REFERENCES

Buchak, L. 2014. 'Belief, Credence, and Norms.' *Philosophical Studies*, 169: 285–311.
Clarke, R. 2013. 'Belief Is Credence One (In Context).' *Philosophers' Imprint*, 13: 1–18.
Easwaran, K. 2015. 'Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities.' *Noûs*, Online First.
—— **and Fitelson, B.** 2015. 'Accuracy, Coherence, and Evidence.' *Oxford Studies in Epistemology*, 5: 61–96.
Foley, R. 2009. 'Beliefs, Degrees of Belief, and the Lockean Thesis.' In F. Huber and C. Schmidt-Petri (eds), *Degrees of Belief*, pp. 37–48. Synthese Library 342. New York, NY: Springer.
Greco, D. 2015. 'How I Learned to Stop Worrying and Love Probability 1.' *Philosophical Perspectives* 29: 179–201.
Harman, G. 1986. *Change in View*. Cambridge, MA: MIT Press.
Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
Jeffrey, R. 1970. 'Dracula meets Wolfman: Acceptance vs. Partial Belief.' In M. Swain (ed.), *Induction, Acceptance, and Rational Belief*, pp. 157–85. Dordrecht: D. Reidel.
Joyce, J. M. 1999. *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
Lance, M. N. 1995. 'Subjective Probability and Acceptance.' *Philosophical Studies*, 77: 147–79.
Leitgeb, H. 2016. *The Stability of Belief. How Rational Belief Coheres with Probability*. Oxford: Oxford University Press.
Lin, H. 2013. 'Foundations of Everyday Practical Reasoning,' *Journal of Philosophical Logic*, 42: 831–62.
—— **and Kelly, K.** 2012. 'Propositional Reasoning that Tracks Probabilistic Reasoning.' *Journal of Philosophical Logic*, 41: 957–81.
Lyon, A. 2014. 'Resisting Doxastic Pluralism: The Bayesian Challenge *Redux*.' Unpublished Manuscript.
Pettigrew, R. 2015. 'Accuracy and the Credence-Belief Connection.' *Philosophers' Imprint*, 15 (16).
——. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
Ross, J. and Schroeder, M. 2014. 'Belief, Credence, and Pragmatic Encroachment.' *Philosophy and Phenomenological Research*, 88: 259–88.
Staffel, J. 2016. 'Beliefs, Buses and Lotteries: Why Rational Belief Can't be Stably High Credence.' *Philosophical Studies*, 173: 1721–34.
Stanley, J. 2005. *Knowledge and Practical Interests*, Oxford: Oxford University Press.
Sturgeon, S. 2015. 'The Tale of Bella and Creda.' *Philosophers' Imprint*, 15 (31).
Tang, W. H. 2015. 'Belief and Cognitive Limitations.' *Philosophical Studies*, 172: 249–60.
Titelbaum, M. 2013. *Quitting Certainties*. Oxford: Oxford University Press.
—— Forthcoming. 'Normative Modeling.' In J. Horvath (ed.), *Methods in Analytic Philosophy: A Contemporary Reader*. New York, NY: Bloomsbury.
Weatherson, B. 2016. 'Games, Beliefs and Credences.' *Philosophy and Phenomenological Research*, 92: 209–36.
Wedgwood, R. 2012. 'Outright Belief.' *Dialectica*, 66: 309–29.

https://doi.org/10.1017/epi.2016.51 Published online by Cambridge University Press

Weisberg, J. 2013. 'Knowledge in Action.' *Philosophers' Imprint*, 13 (22).
—— Forthcoming. 'Belief in Psyontology.' *Philosophers' Imprint*.

JULIA STAFFEL is an Assistant Professor in the Department of Philosophy at Washington University in St. Louis. She received her Ph.D. in philosophy from the University of Southern California in 2013. She mainly works in the areas of traditional and formal epistemology, with a focus on questions about rationality and reasoning.

https://doi.org/10.1017/epi.2016.51 Published online by Cambridge University Press