

Article Type: Symposium

Corresponding author mail id: [jjoyce@umich.edu](mailto:jjoyce@umich.edu)

## Remarks on Richard Pettigrew's *Accuracy and the Laws of Credence*

James M. Joyce

Department of Philosophy, University of Michigan

Richard Pettigrew's *Accuracy and the Laws of Credence* is a marvelous book. I agree with most of it, but I will raise a few worries about some of its claims. In particular, will take issue with: (1) Pettigrew's handling of Bronfman's objection; (2) his rejection of dominance reasoning when the dominant option fails to satisfy certain requirements of rationality, and (3) his commitment to an *accuracy-only* epistemology on which, "the only constraints evidence can place on credence functions come from considerations of accuracy, together with decision-theoretic principles" (p. 29) I will argue that Pettigrew's accuracy-only approach should be replaced with what I call an *accuracy-centered* epistemology.

### 1. The Bronfman Objection (Chapter 5)

In 2003 Aaron Bronfman devised a clever objection to the accuracy dominance argument for probabilism. This argument shows that each inaccuracy score  $I$  associates each incoherent credence function  $c$  with a non-empty set  $D_I(c)$  of probability functions that dominate it. As Bronfman observes, it is easy to find pairs of scores for which  $D_I(c)$  and  $D_{I^*}(c)$  are *disjoint*. This seems worrying because, while  $I$  apparently encourages a shift from  $c$  to some state in  $D_I(c)$ ,  $I^*$  seems to encourage a shift into  $D_{I^*}(c)$ . But, if both scores are legitimate, then no *single* coherent credal state is univocally recommended as an improvement on  $c$ . Even worse, shifting from  $c$  into  $D_I(c)$  can make credences *less* accurate according to  $I^*$ . In light of this, Bronfman argues that the dominance argument depends on the false assumption that credences which are defective according to each accuracy score are thereby defective, full stop. If the various scores do not speak with one voice, why listen to any of them?

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the [Version of Record](#). Please cite this article as [doi: 10.1111/phpr.12502](https://doi.org/10.1111/phpr.12502)

This article is protected by copyright. All rights reserved

Pettigrew takes Bronfman's objection seriously (pp. 75-76), and seeks to avoid it by imposing a strong symmetry condition that eliminates every candidate accuracy score but one (the Brier score). I once suggested something similar, noting that Bronfman's problem evaporates if "only one [accuracy score] functions as the correct measure of epistemic disutility in any context" (2009, p. 290) I have misgivings about Pettigrew's symmetry requirement, though I will not discuss them here, and I no longer believe that any single inaccuracy score will be optimal either across or within contexts of epistemic evaluation. The class of legitimate scores is *easily* wide enough to sustain Bronfman's objection. Yet, the objection no longer troubles me. I have come to appreciate that its key premise is false:  $I$  and  $I^*$  do not make incompatible recommendations when  $D_I(c)$  and  $D_{I^*}(c)$  are disjoint. The various inaccuracy scores *do* speak univocally, though they don't say quite as much as we might have thought.

Bronfman's objection is based on the following premise:

**Dominators.** A person with dominated credences should always move to some dominating alternative, if she can.

This is false. Dominance principles are *knockout rules* that prohibit dominated alternatives, but (absent further information) say nothing about *which* undominated alternatives are optimal. An argument which shows that each option in set  $Y$  is dominated by some option in set  $X$  gives us a decisive reason to avoid  $Y$ s, but no reason to embrace any  $X$  (unless we independently know that the best option is either in  $X$  or  $Y$ ). The same holds for credal states. When I show you that your credences  $c$  are dominated by everything in  $D_I(c)$ , I am *not* recommending that you adopt one of this set's members as your credal state. Rather, I am telling you to consult your total evidence to see which credences it best supports, in  $D_I(c)$  or not. Contra Dominators, it can be a mistake for someone with dominated credences to adopt a dominating alternative, even when she can.

To illustrate, consider an argument from Easwaran and Fitelson (2012) which purports to show that the requirement of accuracy non-dominance conflicts with the *Principal Principle*. Ignoring complexities about undermining information,<sup>1</sup> the Principle reads as follows:

**PP.** If a believer with credences  $c$  sets  $c(T_p) > 0$  where  $T_p$  says that the probability function  $p$  gives the actual current chances, then  $c$  should satisfy  $c(\bullet|T_p) = p(\bullet)$ . So, if she learns (or knows) that  $p$  is the actual chance function then she should align her credences with  $p$ 's values.

Easwaran and Fitelson imagine a believer who knows that a particular coin toss is biased 0.7 for *Heads*, and who assigns incoherent credences of 0.7 to *Heads* and 0.2 to *Tails*. They claim that invoking accuracy dominance to rule out these incoherent  $\langle 0.7, 0.2 \rangle$  credal state violates PP because (with  $I = \text{Brier}$ ) no probability function that dominates its assigns *Heads* a value of 0.7. But, why is that relevant? Unless the incoherent believer is required to adopt a credal state that dominates  $\langle 0.7, 0.2 \rangle$  this is a *non sequitur*. But, while accuracy dominance prohibits her from holding  $\langle 0.7, 0.2 \rangle$ , it does *not* follow that she should hold some state in  $D_I(\langle 0.7, 0.2 \rangle)$ . Instead, she should figure out which coherent state, whether in  $D_I(\langle 0.7, 0.2 \rangle)$  or outside it, is best supported by her evidence and adopt it.

What state might that be? On my view, any theory of evidence consistent with an accuracy-based epistemology must satisfy:

**SUPPORT.** If laws of rational estimation require a believer with total evidence  $E$  to have a higher estimate for the inaccuracy of  $b$  than for the inaccuracy of  $b^*$ , then  $b^*$  is better supported by  $E$  than  $b$  is.

The challenge lies in explaining when believers are rationally required to have a higher estimate of inaccuracy for one credence function than another. Two principles seem non-negotiable:

**SUPPORT<sub>DOM</sub>.** If  $b^*$  accuracy dominates  $b$ , then  $b^*$  is better supported than  $b$  by every consistent body of evidence.

---

<sup>1</sup> Here each candidate chance is assumed to satisfy:  $p(T_p) = 1$ ;  $p(E) = 1$  whenever  $c(E) = 1$ ;  $p$  encodes no other “inadmissible” evidence in the sense of Lewis (1980). Pettigrew discussion of what happens when these restrictions are lifted (Chapters 8-11) is splendidly illuminating.

SUPPORT<sub>CH</sub>. If  $E$  contains enough information about objective chances to determine that the objective expected inaccuracy of  $b^*$  exceeds that of  $b$  (and if  $E$  contains no inadmissible data), then  $b^*$  is better supported by  $E$  than  $b$  is.

The first principle entails that if  $b_E$  is nominated as the best supported credal state given  $E$ , then  $b_E$  must be a probability. So, we should hold coherent credences *no matter what our evidence*. The second principle will help us fix on the right coherent credences when we know enough about chances to make inferences about objective expected accuracies. Applying these principles to the Easwaran/Fitelson example, we should say that (i) all the credences in  $D_I(\langle 0.7, 0.2 \rangle)$  are *better* supported than  $\langle 0.7, 0.2 \rangle$  by the evidence  $ch(\text{Heads}) = 0.7$ , but (ii) the credence function *best* supported by that evidence is  $\langle 0.7, 0.3 \rangle$ , just as PP says.

As this case illustrates, it is no problem when two inaccuracy scores  $I$  and  $I^*$  generate different dominance sets for some  $c$  since a believer who holds  $c$  is *not* required to adopt a credal state in either set. Whether inaccuracy is measured by  $I$  or  $I^*$ , she is advised to hold the credences best supported by her evidence  $E$ , i.e., those that have the lowest estimated inaccuracy in light of  $E$ . Since these credences might not be in *either*  $D_I(c)$  or  $D_{I^*}(c)$ , Bronfman's objection never gets off the ground.

There would be a residual worry if, for some  $E$ , the credences with the lowest estimated  $I$ -value differed from those with the lowest estimated  $I^*$ -value. But, this will never occur as long as our theory of evidence is comprised of norms of *probabilistic* form. Such norms will always require a believer with total evidence  $E$  to estimate credal accuracies using expectations calculated with a *probability* function  $p_E$ .<sup>2</sup> Since inaccuracy scores are strictly proper  $p_E$  will always uniquely minimize  $p_E$ -expected inaccuracy with respect to *both*  $I$  and  $I^*$ . So, the right way to "proportion one's belief to the evidence" will not depend on which score is used.

## 2. Does Dominance Need Strengthening? (Chapter2)

This reading of dominance arguments raises questions about another aspect of Pettigrew's view. He claims that the mere fact that  $c$  is accuracy dominated is insufficient, by itself, to undermine

---

<sup>2</sup> This assumes that there will be a *unique* probability  $p_E$  for every consistent body of evidence  $E$ , but nothing important would change if norms of evidence recommended *sets* of probabilities.

its credentials as a rational credence function:  $b$ 's dominance of  $c$  only counts against  $c$  when  $b$  meets some minimal conditions of rationality (specifically, being both undominated and not overly "modest").

This is not how dominance principles are normally understood. Ordinarily, if  $A$  dominates  $B$ , then  $B$  is deemed unchoiceworthy even when  $A$  is unchoiceworthy. As far as I can see, the only reason to be tempted toward Pettigrew's view is the mistaken idea that, in addition to knocking out dominated alternatives, dominance arguments provide a kind of *endorsement* for dominating alternatives. The suggestion would be that in invoking  $A$ 's dominance to torpedo  $B$  we implicitly propose  $A$  as a serious contender for choiceworthiness. If this were right, then we surely would want  $A$  to meet certain minimum standards of rationality before using it in dominance arguments. However, when we point out that  $A$  dominates  $B$ , we imply nothing about  $A$ 's choiceworthiness (unless we already happen to know that either  $A$  or  $B$  is sure to be our best option). Perhaps every option that dominates  $B$  is atrocious. It doesn't matter!  $B$  fails to be choiceworthy simply because it is dominated by *something*! This something need not itself be choiceworthy.

This reflects the general situation. If we keep in mind that using  $A$ 's dominance of  $B$  to eliminate  $B$  does not commit us to recommending  $A$  in any way (except to say that it is better than  $B$ ), then we will not be tempted to restrict dominance arguments by requiring the dominant alternative to clear some bar or have some property that all serious contenders for choiceworthiness must have. In particular, to prove that  $c$  is impermissible we need *only* show that  $D_I(c)$  is non-empty. This tells us that  $c$  is not the best justified state given our evidence, and that's all we need to knock it out of contention.<sup>3</sup>

### 3. Accuracy-Only Epistemology (Chapter 10)

Pettigrew might object to the forgoing because, by embracing SUPPORT, I am straying from the accuracy-*only* path. My error, he contends, lies in thinking that norms of evidence "have their source both in the value of respecting evidence *and* in the value of accuracy," (p. 29) while the

---

<sup>3</sup> Pettigrew offers a version of the old "name your fortune" game to show that being dominated does not prevent an option from being choiceworthy when *every* option is dominated. (p. 21) He assumes the case cannot be a "rational dilemma" with *no* choiceworthy options. I have the opposite view: decisions in which every option is dominated are rational dilemmas *par excellence*. See my (20xx).

correct approach would *derive* them from accuracy considerations via decision-theory. The situation is a bit more nuanced than Pettigrew makes out. While do I deny that evidential norms like PP, or any norm beyond Dominance, can be *derived* solely from the requirement to have accurate beliefs, I also deny that “respecting evidence” is a *separate* epistemic good that stands apart from accuracy. Let me start with the derivation point.

Pettigrew aims to derive all legitimate norms of evidence from pure accuracy norms by means of purely “decision-theoretic” principles that are free of all evidential entanglements (to preserve the accuracy-only character of the derivations). Unfortunately, the decision-theoretic principles that Pettigrew invokes are far from epistemically innocent. His justification of PP provides an example. Ignoring subtleties, Pettigrew seeks to derive PP from the following allegedly decision-theoretic norm:

**Chance Dominance (CD):** It is impermissible for an rational agent with credences  $\mathbf{c}$  to choose  $o$  over  $o^*$  when, for every probability function  $\mathbf{p}$  with  $\mathbf{c}(T_p) > 0$ , the  $\mathbf{p}$ -expected utility of  $o^*$  exceeds that of  $o$ .

In reality, CD is a conjunction of a decision-theoretic principle and a norm of evidence. To see why we need some terminology. Credal states can be modeled either: *precisely*, as single, sharp credence functions; *imprecisely*, as sets of such functions; or *vaguely*, as “fuzzy” sets of such functions. On any of these models we can make sense of the idea that a believer’s credal state is *confined to a family of probability functions*  $\mathbf{C}$ , i.e., her credal state is either an element (precise), a subset (imprecise), or a fuzzy subset (vague) of  $\mathbf{C}$ . Lastly, say that  $\mathbf{p}$  is a *candidate chance* for a believer just when her credal state is confined to the set  $\mathbf{Ch} = \{\mathbf{c} : \mathbf{c}(T_p) > 0\}$ . In these terms, we can write CD as a conjunction of two *independent* requirements:

**EU.** A agent whose credal state is confined to  $\mathbf{C}$  should prefer  $o^*$  to  $o$  when  $Exp_{\mathbf{c}}(\mathbf{u}(o)) > Exp_{\mathbf{c}}(\mathbf{u}(o^*))$  for every  $\mathbf{c} \in \mathbf{C}$ .

**Chance Expert (CE).** Assuming that chances are probabilities, a believer’s credal state should be confined to the convex closure of the set of her candidate chances. This is the set  $\mathbf{Ch}^+$  of functions  $\mathbf{c}(\bullet) = \sum_p \lambda_p \mathbf{p}(\bullet)$  where  $\mathbf{p}$  ranges over  $\mathbf{Ch}$ , and where, for each  $\mathbf{p}$ , the  $\lambda_p$  are non-negative real numbers summing to 1.

EU is a decision-theoretic norm if ever there was one. CE is plainly epistemic: it tells believers how to proportion the strength of their beliefs to evidence about the objective chances. It asks a believer with sharp credences to satisfy  $c(\bullet) = \sum_p c(T_p) \cdot p(\bullet)$  — equivalently  $c(\bullet|T_p) = p(\bullet)$  for all  $p \in Ch$  — so that her credences are her expectations of the chances. (Those with imprecise/fuzzy beliefs will have credal states given by determinate/fuzzy subsets  $Ch^+$ .) Taken together EU and CE require an agent’s subjective expectations to agree with her estimates of the objective chance expectations, so that  $Exp_c(\bullet) = Exp_c(Exp_{ch}(\bullet))$ . CD follows directly.

Unfortunately, this nice connection between objective and subjective expectations only holds if we assume *both* EU and CE. Neither Chance Dominance nor the Principal Principle can be derived from EU alone. To see why, go back and reread Pettigrew’s discussion of CD on p. xx of the *Precis*: “The Principal Principle... is a variant of a law of rational choice that we might call Chance Dominance... So you should take it.” There is a lacuna in this reasoning. The only “law of rational choice” in view is EU, which has you buy the bet only if your *subjective* expectation of its payoff exceeds £1 (with utility = money). But, how do we get from “the bet’s objective expected utility exceeds £1” to “your subjective expected utility for it should exceed £1”? To bridge this objective-to-subjective gap we need CE, which confines  $c$  to  $\{p : p(\text{heads}) > 0.6\}$ . If we did not require this, then you would have no reason to assign an expected utility in  $[1.2, 2]$ . More generally, you have no reason to align your subjective expectations with the chance expectations unless you have reason to align your credences with the chances. So, CD presupposes CE. But, CE is merely PP generalized to cases where credences might not be sharp! We have no accuracy-only justification for the Principal Principle here.

The circular pattern recurs when options are credal states and (dis)utilities are *I-values*. Here Chance Dominance is again a conjunction of a decision-theoretic norm and an evidential one.

**Minimize Subjective Estimated Inaccuracy (EI).** A believer’s whose credal state is confined to a set of probabilities  $C$  is committed to regarding  $b^*$  as a better credal state to occupy than  $b$  whenever  $Exp_c(I(b)) > Exp_c(I(b^*))$  for every  $c \in C$ .

**Chance Expert (CE).** Assuming that chances are probabilities, an epistemically rational believer’s credal state should be confined to  $Ch^+$ , the convex closure of the set of her candidate chances.

As before, it takes *both* norms to justify PP. And, again as before, there is no significant difference between CE and PP. We have no non-circular, accuracy-*only* justification for the Principal Principle here.

Pettigrew briefly considers the objection I am raising (pp. 129-132), but dismisses it by arguing that

- (i) CD is no worse off than accuracy non-dominance. Even though EI-plus-probabilism entails accuracy non-dominance, using it to justify probabilism does not beg the question. This is because accuracy non-dominance is *more basic* than probabilism.
- (ii) Chance Dominance more basic than PP and CE because: (ii-a) it binds believers who lack precise credence functions, and whether or not they are coherent; and (ii-b) while PP and CE require credences to be confined to  $Ch^+$ , CD only has  $b^*$  preferred to  $b$  “on rare occasions” where all candidate chances assign  $b^*$  a lower expected inaccuracy.

I am not persuaded. (i) ignores a key difference between the cases. The power of dominance reasoning is that it does not depend on how credences are apportioned out among events. When I know  $o^*$  dominates  $o$  according to your desires, I do not need to know *anything* about your beliefs to conclude that  $o$  is not your best choice. I don’t even need to know whether you are coherent! Choosing dominated options is wrong for incoherent agents for the same reason it’s wrong for coherent ones: it commits one to incurring sure losses or passing up sure gains. Likewise, when  $b^*$  accuracy dominates  $b$  you are obliged to see  $b$  as suboptimal *whatever your beliefs*. So, accuracy non-dominance follows from EI *alone*, without any help from probabilism, which is why we can use the former to justify the latter. This is *not* true for EI and CE. So, there is no parity here.

As to (ii-a), my formulation of CE renders its first part moot. CE does not force believers to adopt precise credences unless they have precise beliefs about precise chances. It does ask them to occupy credal states (precise, imprecise, fuzzy) confined to  $Ch^+$ , but CD requires this too! If all candidate chances are *probabilities* it can be proven that, for any inaccuracy score  $I$ , the only way to avoid being chance dominated is by having a credal state wholly confined to  $Ch^+$ . (Theorem I.D.5, p. 92). This won’t be true if incoherent chances are allowed, but in that case CD



can sometimes recommend  $c^*$  over  $c$  even when  $c$  dominates  $c^*$ , something an accuracy-based epistemology will never permit.

Pettigrew defends (ii-b) by saying that CD “merely says that, on the rare occasions on which [the candidate chances] all agree in their ordering of two credence functions with respect to accuracy... you should adopt that ordering yourself.” (132) This makes it seem as if CD is a limited, local norm that applies much less frequently than CE and PP. In fact, it applies just as widely. For non-trivial cases, there will be infinitely many pairs  $(b^*, b)$  where  $b^*$  chance-dominates  $b$ . CD requires that, in any such pair, a believer’s credal state must not contain a chance-dominated  $b$ . As already noted, the only way to ensure this is by having a credal state confined within  $Ch^+$ , just as CE requires.

#### 4. Accuracy-Centered Epistemology

The lesson here is that we should replace the dream of an accuracy-*only* epistemology with that of an accuracy-*centered* epistemology that (a) concedes that evidential norms, like PP, cannot be derived from accuracy considerations alone, and yet (b) does not embrace “respecting evidence” as an *independent* goal that stands apart from that of having accurate beliefs.

To see how this can be achieved, note first that many evidential norms can be framed as *expert principles*. Let  $\Pi$  be a probabilistic information source, i.e., a random variable whose potential values are probability functions about which a believer might have or gain evidence.  $\Pi$  might be *the chances at t*, *a meteorologist’s forecast chance of rain*, *my credences this afternoon*, or even *the actual truth-values*. Let  $\Pi_p$  mean that the (rigidly designated) probability  $p$  gives  $\Pi$ ’s actual values. A believer with credences  $c$  defers to  $\Pi$  as an expert when  $c(\Pi_p) > 0$  and  $c(\bullet|\Pi_p) = c(\bullet)$ . It follows directly that her estimates will be her expectations of the expert’s expectations. An *expert principle* requires rational believers with credences  $c$  and evidence  $E$  to defer to some  $\Pi$  as an expert, thus telling believers how to update their beliefs in light of evidence about an expert’s values. PP is an expert principle with  $\Pi$  as the current chances.

We can relate the notion of an epistemic expert to considerations of accuracy using two closely related concepts (where  $p$  ranges over all probabilities with  $c(\Pi_p) > 0$  and  $b$  ranges over all credences).

- $c$  sees  $\Pi$  as *trustworthy* when  $Exp_{c(\bullet|\Pi_p)}(\mathbf{I}(\mathbf{b})) \geq Exp_{c(\bullet|\Pi_p)}(\mathbf{I}(\mathbf{p}))$ , with equality iff  $\mathbf{p} = \mathbf{b}$ .
- $c$  sees  $\Pi$  as *reliable* when  $Exp_c(\mathbf{I}(\mathbf{b})) \geq Exp_c(\mathbf{I}(\Pi)) = \sum_p c(\Pi_p) \cdot Exp_{c(\bullet|\Pi_p)}(\mathbf{I}(\mathbf{p}))$ .

If a coherent believer treats  $\Pi$  as trustworthy, then learning  $\Pi\mathbf{p}$  (and no more) leads her to have a lower expectation for  $\mathbf{p}$ 's inaccuracy than any other (rigidly specified) probability. By SUPPORT, she should then see  $\mathbf{p}$  as the credence function that is best justified by her posterior evidence (her prior evidence augmented with  $\Pi_p$ ). A coherent believer treats  $\Pi$  as *reliable* when she assigns it a lower expected inaccuracy than any (rigidly specified)  $\mathbf{b}$ . A simple dominance argument shows that trustworthiness implies reliability.

For our purposes, the key point about experts is this:

FACT.  $c$  defers to  $\Pi$  as an expert iff  $c$  sees  $\Pi$  as trustworthy, and only if  $c$  sees  $\Pi$  as reliable.<sup>4</sup>

Committing to an expert norm thus requires us treat its expert as an especially accurate source of information.

This might seem to offer hope for accuracy-only approaches since FACT seems to provide a kind of template for deriving evidential norms. If we can just prove that epistemic rationality requires  $\Pi$  to be trustworthy, so that

$$(\#) \quad Exp_{c(\bullet|\Pi_p)}(\mathbf{I}(\mathbf{b})) > Exp_{c(\bullet|\Pi_p)}(\mathbf{I}(\mathbf{p})) \text{ for all } \mathbf{p} \text{ and } \mathbf{b} \text{ with } \mathbf{p} \neq \mathbf{b},$$

then we will have established  $\Pi$  as an expert for all rational believers. But, this is a false hope. To prove (#) we must *already* know how to calculate expectations conditional on  $\Pi_p$ . But, since  $\mathbf{I}$  is strictly proper, we know  $Exp_{c(\bullet|\Pi_p)}(\mathbf{I}(\mathbf{b})) > Exp_{c(\bullet|\Pi_p)}(\mathbf{I}(c(\bullet|\Pi_p)))$  for all  $\mathbf{b} \neq c(\bullet|\Pi_p)$ . So, the only way to secure (#) is by having  $\mathbf{p}(\bullet) = c(\bullet|\Pi_p)$  when  $c(\Pi_p) > 0$ , which is just to say that  $\Pi$  must be an expert for  $c$ . Thus, we must assume that  $\Pi$  is an expert to prove its trustworthiness, and cannot appeal to its trustworthiness to justify giving it expert status.

---

<sup>4</sup> The converse of the last clause is true as long as  $c$  continues to regard  $\Pi$  as an expert after conditioning on  $\Pi_p$  for each candidate  $\mathbf{p}$ ,

The take-home point is that one cannot deduce a source's trustworthiness from pure accuracy considerations alone: one also needs to invoke norms of evidence. If we have *already* identified certain sources as experts, then accuracy-centered epistemology can help us find others. For example, it follows from Greaves and Wallace (2006) that if  $\Pi$  is trustworthy and  $0 < c(E) < 1$  then  $\Pi(\bullet | E)$  is also trustworthy. The hard part lies in identifying trustworthy sources to use as inputs in the first place. Considerations of accuracy, unsullied by evidential entanglements, will not do the job themselves. This is not to say that norms of evidence are more basic than norms of accuracy. The two are sides of the same coin. When we endorse an expert principle, like PP, we commit to viewing its expert as an especially accurate information source, and thus to seeing obeying the principle as integral to the rational pursuit credal accuracy. Of course, to make these commitments we must have good reasons for thinking that a policy of aligning our beliefs with the expert's values will promote accuracy. However, these reasons will not be found in an epistemology that relies only on accuracy scores and decision-theory. Evidence must be a part of the picture from the start.

## References

- Joyce, J. M. (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In F. Huber and C. Schmidt Petri, eds., *Degrees of Belief*. Springer.
- Joyce, J. M. (2017). Accuracy, Ratification and the Scope of Epistemic Consequentialism. In J. Dunn and K. Ahlstrom-Vij, eds., *Epistemic Consequentialism*. Oxford University Press.
- Joyce, J. M. (20xx). Are There Rational Dilemmas? (in preparation)
- Lewis, D. (1980). A Subjectivist's Guide to Objective Chance. In R. C. Jeffrey, ed., *Studies in Inductive Logic and Probability, Volume II*. University of California Press, Berkeley.
- Pettigrew, R. (2016) *Accuracy and the Laws of Credence*. Oxford University Press.