

Lockeans Maximize Expected Accuracy

KEVIN DORST

Massachusetts Institute of Technology

kmdorst@mit.edu

Forthcoming in *Mind*

The **Lockean Thesis** says that you must believe p iff you're sufficiently confident of it. On some versions, the 'must' asserts a metaphysical connection; on others, it asserts a normative one. On some versions, 'sufficiently confident' refers to a fixed threshold of credence; on others, it varies with proposition and context. Claim: the Lockean Thesis follows from **epistemic utility theory**—the view that rational requirements are constrained by the norm to promote accuracy. Different versions of this theory generate different versions of Lockeanism; moreover, a plausible version of epistemic utility theory meshes with natural language considerations, yielding a new Lockean picture that helps to model and explain the role of beliefs in inquiry and conversation. Your beliefs are your *best guesses* in response to the epistemic priorities of your context. Upshot: we have a new approach to the epistemology and semantics of belief. And it has teeth. It implies that the role of beliefs is fundamentally different than many have thought, and in fact supports a metaphysical reduction of belief to credence.

You ask me if I think Bob's in his office, and I reply that I'm confident he is. Have I answered your question? Have I told you what I think? Some theorists say 'No': there's a separate attitude—belief—and telling you that I have high credence doesn't settle the question of whether I believe.¹ In contrast, the **Lockean Thesis** says that I must believe Bob's in his office iff I'm sufficiently confident he is.² This thesis can be interpreted in different ways. (1) The 'must' could assert either a *metaphysical* or a *normative* connection: perhaps belief reduces to high credence; or perhaps they are distinct attitudes which simply ought to be in harmony. (2) The 'sufficiently confident' could refer to either a *fixed* or *variable* threshold of credence: perhaps there is a single threshold such that I must believe iff my credence is above that threshold; or perhaps the threshold can vary with proposition and context. Mix and match these interpretations as you like—they are all versions of the Lockean Thesis as I will understand it.

¹E.g. Levi (1967), Stalnaker (1984, 2006), Maher (1993), Kaplan (1996), Williamson (2000), Ganson (2008), Frankish (2009), Buchak (2014), Ross and Schroeder (2014), Smith (2016), and Staffel (2016).

²Sturgeon (2008), Foley (2009), Demey (2013), and Leitgeb (2013, 2014) defend versions of this view.

You ask me if you should be more confident that Bob’s in his office than that Alex is, and I reply that that’ll make you most accurate. Have I answered your question? Have I told you whether to be more confident? Some theorists say ‘No’: there’s a separate issue—rationality—and telling you that an attitude best promotes accuracy doesn’t settle the question of whether it’s rationally required.³ In contrast, **epistemic utility theory** says that rationality requires you to have an attitude that best promotes accuracy.⁴

Claim: Epistemic utility theory entails the Lockean Thesis. Different versions of the theory generate different versions of Lockeanism; moreover, a plausible version of epistemic utility theory meshes with natural language considerations, yielding a new Lockean picture that helps to model and explain the role of beliefs in inquiry and conversation. Your beliefs are your *best guesses* in response to the epistemic priorities of your context. Upshot: we have a new approach to the epistemology and semantics of belief. And it has teeth. It implies that the role of beliefs is fundamentally different than many have thought, and in fact supports a metaphysical reduction of belief to credence.

Plan: After taking a fresh look at Lockeanism through natural language (§1), we’ll survey the machinery and standard assumptions of epistemic utility theory (§2). Applying our machinery to beliefs, we’ll get a Lockean thesis with a fixed threshold (§3). Next, it’s time to adjust the knobs and dials. Dropping one of the standard assumptions (Extensionality) is independently motivated and leads to a plausible form of variable-threshold Lockeanism; and we can drop a second assumption (Separability) without loss (§4). At this point we’ll have our new Lockean picture. In closing, we’ll explore its epistemological and metaphysical significance (§5).

1 What a Lockean thinks

I say epistemic utility theory entails that high credence suffices for belief. But some will reply that this is a reductio of that theory—an example of what happens when we allow a formal framework to outstrip good sense. Haven’t we learned from lottery cases that high credence—no matter *how* high—doesn’t suffice for belief? Here’s your ticket; there are 1000 others. Should you believe yours will lose? Many say ‘No’: you should simply believe it’s *likely* to lose. Why’s that? Well, here things get messy. Maybe it’s because believing p is ruling out $\neg p$ (Hintikka 1962); or because beliefs require more than statistical evidence (Buchak 2014; Staffel 2016); or because beliefs are what you’d assert in a context of inquiry (Kaplan 1996); or because beliefs aim at knowledge (Williamson 2000); or because you should believe p only if it’d be *abnormal* if $\neg p$ (Stalnaker 2006;

³E.g. Berker (2013), Greaves (2014), Littlejohn (2015), and Carr (2017).

⁴See Pettigrew (2013b) for a helpful overview. For examples of the research program in action, see Joyce (1998, 2009), Greaves and Wallace (2006), Predd et al. (2009), Leitgeb and Pettigrew (2010), Easwaran (2013), Pettigrew (2013a, 2016a,b), Carr (2015), Schoenfield (2016, 2017), De Bona and Staffel (2016), Horowitz (2017), and Staffel (2017).

Smith 2016; Goodman and Salow 2017); or because beliefs are what you rely on as premises in reasoning (Frankish 2009; Ross and Schroeder 2014); or because beliefs are what you could conditionalize on without affecting your behavior (Weatherson 2005; cf. Ganson 2008). Regardless, you just shouldn't believe you'll lose.

I must admit, I used to agree. But then Hawthorne et al. (2016) helped me see that natural language offers a very different picture of belief—one that fits well with Lockeanism. In this section we'll focus on the modest point that Lockeanism is not a nonstarter. In §4.1 we'll see further that a plausible form of epistemic utility theory meshes with natural language in a variety of more subtle and interesting ways.

(But what does our use of *language* have to do with epistemology? Everything—at least according to the T-schema. If I'm right, our use of natural language suggests that the sentence 'You believe your lottery ticket will lose' is true. It follows that you believe your lottery ticket will lose. Since you are reasonable, it follows that reasonable people can believe lottery propositions, after all.)

To begin, we'd be well-served to follow Hawthorne et al. (2016) in focusing on judgments using 'think' instead of 'believe'. The former is much more colloquial⁵, and philosophers often use the latter in a technical way—flagged by jargon like 'full belief' and 'outright belief.' Thus the judgment that 'You shouldn't believe your ticket will lose' may be theoretically loaded. What happens when we change the verb?

First, Hawthorne et al. (2016: 1398) point out that—contra a natural first thought—'think' does not express a weaker state than 'believe.' If it did, then it would make sense to affirm the former and deny the latter. But it doesn't⁶:

(3) ? Tim thinks the dog is outside, but $\left\{ \begin{array}{l} \text{he doesn't believe} \\ \text{it's not as if he believes} \end{array} \right\}$ she is.

(Contrast: 'Tim thinks the dog is outside, but he doesn't know she is'.) This suggests that you think p only if you believe it. The converse holds as well:

(4) ? Tim believes the dog is outside, but $\left\{ \begin{array}{l} \text{he doesn't think} \\ \text{it's not as if he thinks} \end{array} \right\}$ she is.

The natural hypothesis is that you think p iff you believe p . I'll adopt it in what follows.⁷

First, note that it makes no sense to express high credence and deny belief:

(5) ? The dog is probably outside, but $\left\{ \begin{array}{l} \text{I don't} \\ \text{it's not as if I} \end{array} \right\}$ think she is.

This suggests that having high credence in rain requires thinking (hence believing) it. Similarly:

(6) ? I think the dog is outside, but $\left\{ \begin{array}{l} \text{it's not} \\ \text{it's not as if it's} \end{array} \right\}$ likely that she is.

⁵Skeptical? Record some conversations and count. My tally: 'think': 95; 'believe': 3.

⁶In hearing these sentences be sure not to put stress/emphasis on them, for doing so affects the semantic values of their parts—witness the acceptability of 'I believe you love me, but I don't *believe* it.'

⁷This hypothesis is helpful but inessential to my argument. First, the core argument from epistemic utility theory for Lockeanism stands on its own. Second, although 'think'-judgments are less theoretically loaded than 'believe'-judgments, all the natural language data holds up using either verb—feel free to replace occurrences of 'think' with 'believe' in what follows.

This suggests the converse: thinking suffices for having high credence.

More generally, we are happy to make conclusions about what to ‘think’ based on merely statistical evidence—as predicted by Lockeanism. You’re wondering whether it will rain? Seeing that there’s a 70% chance, I reply, ‘I think it will.’ Most people like starfruit, though few have tried it. You ask me what I’m eating and I reply, ‘It’s a starfruit—try it, I think you’ll like it.’ Student Steve asks Professor Peterson about a class. Knowing only that 90% of students get good grades—she replies, ‘I think you’ll find it manageable.’ Hopeful Holly clutches her lottery ticket, waiting for the numbers to fall. Alas, I think her ticket will lose. Don’t you? Examples could be multiplied. We are constantly having to figure out what to think based on merely statistical, inconclusive evidence—and outside the philosophy room we do so without pause or compunction.

...Or at least that’s what I want to suggest. The interested reader is referred to Hawthorne et al. (2016) (and below, in §4.1) for further discussion of these linguistic points. At this stage my claim is the modest one that Lockeanism is not absurd—the formal argument I’m about to give is not divorced from common sense.

(Maybe so. But aren’t Hintikka et al. right that there *is* a theoretically important attitude toward p that amounts to ruling out $\neg p$; that aims at knowledge; that you rely on in reasoning; etc? Yup!—It’s just not belief, says the Lockean. See §5.1.)

2 Epistemic utility theory

So our conclusion isn’t absurd—how do we reach it? Epistemic utility theory: a research program started by Rosenkrantz (1981), Oddie (1997), and Joyce (1998) that attempts to justify norms of rationality using the constraint that they must promote accuracy. Two core assumptions.

(1) First, there is a distinctive kind of epistemic value—a degree to which a given doxastic state is epistemically good or bad (for a given agent, in a given context) at a given possible world. We’ll make the traditional assumption that it’s a metric of *accuracy*. Model it using an accuracy function A that takes a doxastic state D , a set of propositions \mathcal{P} , a world w , and outputs D ’s Accuracy with respect to \mathcal{P} at w : $A(D, \mathcal{P}, w)$. Our understanding of accuracy can then be used to justify various constraints that A must meet (Joyce 2013, §6)—for instance, A should increase as you become more confident in truths.

(2) Our second core assumption is that the norms of rationality are constrained by the norm to try to promote accuracy. Two parts to this. (2.i) We say agents should ‘try’ to promote accuracy because this is a theory of subjective obligation—of what you should do, given the information you have. Though it remains implicit in many discussions, this point is highlighted by the fact that A —the metric of objective value—measures accuracy. Yet no one ever thought that rationality requires accurate beliefs. Instead, it requires adopting methods that are *expected* to be conducive to this end—that promote accuracy by your own lights.

(2.ii) To say rational norms are ‘constrained’ by promoting accuracy means that *if* a norm can be derived purely from considerations of accuracy, then it’s a rational norm. The basic idea is that rationality—whatever else it is—must be an optimal guide to truth; thus we should *start* with accuracy in theorizing about it. This has a methodological upshot: in setting up our framework we should not impose external constraints on the class of doxastic alternatives. For instance we do not want to stipulate that the class of credence functions includes only probabilistic ones, for we want to see whether accuracy alone will get us there. Generalizing, call this an **Accuracy First** methodology: we should leave open a wide class of doxastic alternatives, since we want to derive norms from accuracy *alone*. Again, Accuracy First is rarely stated, but it’s implicit in the way epistemic utility theorists set up their frameworks and paint the big picture of what they are trying to do.

(Accuracy First is crucial, and explains why I do not say Hempel (1962), Levi (1967), and Maher (1993) are founders of epistemic utility theory. They too ask what to believe given your epistemic values. But—crucially—they reject Accuracy First, instead imposing external bounds on permissible belief states. Example: deductive closure. This constraint is not motivated by considerations of accuracy, and is the reason—in fact, the *only* reason—why they do not end up with Lockeanism.)

So we have our machinery in place. To get a grip on how it works and the assumptions we’ll use, let’s begin with the classic example. Consider Rational Rachael, who assigns credences C to a set of propositions \mathcal{P} modeled as subsets of a (finite) set of doxastic possibilities W . Suppose we want to explain why her credences are probabilistic. Accuracy First requires us to start with a wide range of credence functions as doxastic alternatives—say, every function from propositions to real numbers. The goal is to argue that only the probabilistic ones respect the norm to promote accuracy.

To do this we have to motivate some constraints on A ; the following three are standard. First, it should be *proper*:

Propriety: Any probabilistic credence function C expects any other particular credence function C^* to be less A -accurate than itself.

If A were *improper* then Rachael would be in a bizarrely unstable state, for in trying to maximize accuracy she’d have to change her credences based on no evidence (Gibbard 2008). (Although it’s needed for probabilism, we won’t need Propriety for Lockeanism.)

Second, A should be *extensional*: the accuracy of Rachael’s credences with respect to p depends only on (1) her credence in p and (2) p ’s truth-value. This is because, intuitively, A is a distance metric—a measure of how *close* Rachael’s credences are to the truth. Consider: the distance between her left foot and the nearest spider depends only on (1) where she puts her left foot and (2) where the world puts the spider. Likewise: the distance between her credence and the truth-value depends on only (1) where she puts her credence and (2) where the world puts the truth-value. Formally:

Extensionality_C: For any p, q : if $C(p) = C(q)$ and p and q have the same truth-value, then $A(C, \{p\}, w) = A(C, \{q\}, w)$.

Rachael has .7 credence in both *Nearby Spider* and *Nearby Snake*, and in fact both are true—there are spiders and snakes nearby. Suppose her accuracy for *Nearby Spider* is 8.4. Can we guess her accuracy for *Nearby Snake*? 8.4, according to Extensionality_C.

Finally, *A* should be *separable*: the accuracy of Rachael’s credences with respect to a set of propositions $\{p_1, \dots, p_n\}$ is the sum of its accuracy with respect to each member of the set. Consider: the distance she’d have to move her feet to stand on the spider and snake is the sum of the distances between (1) her left foot and the spider and (2) her right foot and the snake. Likewise: the distance she’d have to move her credences to hit on the truth of *p* and of *q* is the sum of the distances between (1) her credence in *p* and its truth-value and (2) her credence in *q* and its truth-value. Formally:

Separability_C: $A(C, \{p_1, \dots, p_n\}, w) = \sum_i A(C, \{p_i\}, w)$

Rachael is wondering how accurate her views on *Nearby Spider* and *Nearby Snake* are. I tell her how accurate they are for *Nearby Spider* and how accurate they are for *Nearby Snake*. Have I answered her question? Yes, according to Separability_C.

Propriety, Separability_C, and Extensionality_C have become the standard operating procedure assumptions about *A*, so let’s suppose it obeys them.⁸ Our final step is to clarify the sense in which Rachael can be thought of as ‘trying to promote accuracy.’ At a bare minimum she must never adopt a *dominated* state: a state *D* that she knows is less accurate than a different state *D**. Theorem: if *A* satisfies Propriety, Separability_C, and Extensionality_C, then Rachael’s credal state is undominated iff it is probabilistic (Predd et al. 2009).

But dominance is a special case. It arises from requiring Rachael to promote *A* by her own lights—to *try* to promote *A*. What does this come to generally? Presumably that she must select an option she *estimates* will best promote *A*. Note that such estimates are judged on a ‘closeness counts’ metric: if there are 63 jellybeans in the jar, then my estimate of 58 is worse than your estimate of 59—even though neither of us were exactly right. Using such a (proper, separable, extensional) ‘closeness counts’ metric, we can prove—on pain of dominance—that Rachael’s *estimate* of *A* must equal her mathematical *expectation* of *A* (de Finetti 1974; Pedersen and Glymour 2012; Pettigrew 2012). This is a weighted average of the accuracy of *D* across worlds, with weights determined by her credence that those worlds are actual. Formally, her *Expected Accuracy* of a doxastic state *D* with respect to the total set of propositions \mathcal{P} is:

$$EA(D, \mathcal{P}) = \sum_w (C(w) \cdot A(D, \mathcal{P}, w))$$

Example: if she’s $\frac{3}{4}$ confident that *D* has accuracy 4, and $\frac{1}{4}$ confident that it has accuracy 1, then its expected accuracy is 3.25.

To promote accuracy by her own lights, Rachael must select a state that maximizes expected accuracy according to her credences:

⁸They are all assumed by Joyce (1998), Predd et al. (2009), Leitgeb and Pettigrew (2010), Easwaran (2013), Pettigrew (2013a) and Briggs and Pettigrew (2016).

Maximize Expected Accuracy: Given credences C , a doxastic state D is rational only if $EA(D, \mathcal{P}) \geq EA(D^*, \mathcal{P})$ for every other D^* .

This norm is standard.⁹ But we needn't appeal to orthodoxy: the above-mentioned proof suggests that Maximize Expected Accuracy is part and parcel with our core requirement to promote accuracy by your own lights. If it wouldn't be rational for Rachael to maximize EA , then (by the above-mentioned proof) she shouldn't do what she *estimates* will best promote A —which just shows that A wasn't the correct metric of accuracy in the first place.

3 Fixed Lockeanism

The epistemic utility machinery has primarily dealt with credences, so our first step is to make it work for beliefs. The assumptions we'll start with are Accuracy First, Extensionality, Separability, and Maximize Expected Accuracy. In this section we show that translating them to apply to beliefs yields a normative, *fixed* Lockean Thesis: rational agents have a threshold t such that, for all p : they believe p iff their credence in p is above t . This is our base result. In the next section I'll argue that it can and should be generalized: dropping Extensionality and Separability yields a normative, *variable* Lockean thesis, wherein the threshold varies with proposition and context.

But first let's situate our argument, since some theorists have engaged in similar projects.¹⁰ In particular, Easwaran (2016) has some very similar formalism—my Theorem 1 and Corollary 2 are identical to two of his results (see his sections 3.2 and G.2; he proved them first). Three key differences. First, I provide a new picture for *why* we should drop Extensionality, and explore how the resulting view meshes with natural language to yield a new picture of belief. Second, I argue that we can drop the (seemingly) essential assumption of Separability (cf. Easwaran 2016, §D). Third, philosophically we are going in opposite directions. Easwaran is attempting to reduce talk of credences to talk of beliefs; I'm trying to do the opposite. So whereas my project can be seen as an extension of the epistemic utility research program, Easwaran's cannot—in rejecting the independent existence of credences he's rejecting the foundations of that program.

Our conclusion isn't absurd, and it's not old news—time to see how we get there. Consider Rational Rachael; we want to prove she's a Lockean. Our story begins—as these stories often do—by letting the idealizations rip. (But bear with me: in §5.1 I'll suggest we can drop them.) So suppose Rachael has a probabilistic credence function C defined over a set of propositions \mathcal{P} , the subsets of (finite) W . We're interested in her *belief-state* B —the set of propositions she believes. It encodes all her relevant attitudes: if she believes p , then $p \in B$; if she disbelieves p , then $\neg p \in B$; and if she suspends judgment

⁹Greaves and Wallace (2006), Leitgeb and Pettigrew (2010), Easwaran (2013), Schoenfield (2016, 2017), and Shear and Fitelson (2017).

¹⁰Easwaran and Fitelson (2015), Pettigrew (2015, 2016c), Easwaran (2016), and Shear and Fitelson (2017).

on p , then $p, \neg p \notin B$.¹¹ Accuracy First says there are no external requirements on B —it need not be consistent, or deductively closed, or nonempty.

We have Rachael’s beliefs; what do they aim at? Epistemic utility theory says: accuracy! Beliefs aim to be true, and to avoid being false. This motivates translations of Separability and Extensionality for the case of beliefs.

First, Separability: the accuracy of Rachael’s beliefs for a set of propositions is the sum of her accuracy for each member of the set. Consider: her score on the exam is the sum of her scores on each question. Likewise: the overall accuracy of her beliefs for a set $\{p_1, \dots, p_n\}$ is the sum of her accuracy for each p_i . Formally:

$$\textbf{Separability}_B: A(B, \{p_1, \dots, p_n\}, w) = \sum_i A(B, \{p_i\}, w)$$

This is a direct translation of Separability_C, above.

Second, Extensionality: the accuracy of Rachael’s beliefs about p depends only on (1) her attitude toward p and (2) p ’s truth-value. She could either believe (b), suspend (s), or disbelieve (d), and p could be either true (t) or false (f); so there are six options: bt, df, st, sf, bf, dt . But she disbelieves p iff she believes $\neg p$, and p is true iff $\neg p$ is false; so disbelieving a falsehood is believing a truth ($df = bt$) and disbelieving a truth is believing a falsehood ($dt = bf$). Moreover, it’s no better or worse to suspend judgment on truths than falsehoods ($st = sf$). The six become three: $bt=df, st=sf, bf=dt$. So Extensionality requires A to take one of three values for a proposition, depending on whether Rachael believes it truly, suspends, or believes it falsely. Call these values T , S , and F , respectively.

Suspending judgment is neither accurate nor inaccurate: $S = 0$.¹² What about T and F ? They correspond to what William James called our two ‘great commandments as would-be knowers’ (1897: §VII). His claim was that ‘Be accurate!’ actually breaks down into two separate injunctions: (1) ‘Seek truth!’, and (2) ‘Avoid error!’ Rachael could fully satisfy (1) by believing everything, and fully satisfy (2) by believing nothing; we only get an accuracy metric once we weigh these factors against each other. In particular, there’s no need to assume that $|T| = |F|$; so what constraints should these values obey? A true belief is more accurate than suspending judgment, which itself is less inaccurate (more accurate) than a false belief: $T > 0 > F$.

Moreover, it’s plausible that (contra James) Rachael will be doxastically *conservative*: she’ll avoid error more fervently than she’ll seek truth. Why? Well here’s a fair coin—does she believe it’ll land heads? Or tails? Or both? Or neither? Clearly neither.

¹¹My treatment of suspension of judgment as *non-belief* is a modeling choice, not a substantive one. Friedman (2013, 2017) has forcefully argued that suspension is distinct from mere non-belief, but the question for our purposes is whether they have different *accuracy* profiles. They cannot. For the accuracy of suspending/non-believing on p cannot depend on whether p is true or false. Thus suspension always yields some accuracy score S , and non-belief always yields some accuracy score N . If $S \neq N$, then one option accuracy-dominates the other. Thus at the level of accuracy there cannot be two distinct, potentially rational options; our model needn’t discriminate between suspension and non-belief.

¹²But this assumption is inessential; see (Easwaran 2016, §C).

But if she cared more about seeking truth than avoiding error, why not believe both? She'd then be guaranteed to get one truth and one falsehood, and so be more accurate than believing neither! Yet believing both is *not* as accurate as believing neither—so belief is conservative. Upshot: we impose a **Conservativeness** constraint to capture the sense in which Rachael has 'more to lose' in forming a belief than she does to gain. That is, the jump in accuracy from suspending judgment (S) to believing a truth (T) is smaller than corresponding jump from believing a falsehood (F) to suspending judgment (S): $|T - S| < |F - S|$. Equivalently, when $S = 0$: $|T| < |F|$.

Upshot: Extensionality says that a belief-state gets one of three values for a given proposition: T for a true belief; F for a false belief; 0 for no belief. Formally:

$$\mathbf{Extensionality}_B: A(B, \{p\}, w) = \begin{cases} T & \text{if } p \text{ is true and believed} \\ F & \text{if } p \text{ is false and believed} \\ 0 & \text{if } p \text{ is not believed} \end{cases}$$

We can combine this with Separability_B to determine Rachael's accuracy for the question of whether p , i.e. $\{p, \neg p\}$. If she believes only the true proposition she gets accuracy T ; if she believes only the false one she gets F ; if she suspends judgment she gets 0; and if she believes both she gets $T + F$ (which, by Conservativeness, is less than 0).

Summing up: epistemic utility theory commits us to Accuracy First and Maximize Expected Accuracy; moreover, its accuracy-theoretic foundations seem—at a first pass—to motivate Separability_B and Extensionality_B. So—at a first pass—epistemic utility theory commits us to a normative, fixed Lockean Thesis: there is a single threshold of credence t (which turns out to be $t = \frac{-F}{T-F}$) such that Rachael must believe everything more likely than t , and must not believe anything less likely than t . Formally:

Theorem 1. (Easwaran 2016, §3.2) If A is separable and extensional, a belief-state¹³ B maximizes expected accuracy iff, for all p :

- (i) If $p \in B$, then $C(p) \geq \frac{-F}{T-F}$
- (ii) If $p \notin B$, then $C(p) \leq \frac{-F}{T-F}$

(See the Appendix for all proofs.) Why does this happen? Extensionality_B ensures that believing p is an epistemic *bet*, which pays out T if p is true and costs F if p is false. When should Rachael take this bet? Maximize Expected Accuracy says: whenever the

¹³In scoring only Rachael's beliefs, I sideline a worry raised by Pettigrew (2015) about scoring joint credal-belief states. He shows that if we *presuppose* Lockeanism, then there's a hitch in the accuracy-dominance argument for probabilism: if C is not probabilistic, the joint state $\langle C, B \rangle$ may not be dominated by a probabilistic and Lockean state $\langle C^*, B^* \rangle$. But this issue does not arise on my account, for the order of explanation is reversed: we only get a Lockean constraint *after* establishing probabilism, for that constraint comes from Maximize Expected Accuracy—a norm that only makes sense for probabilistic credences (Joyce 1998, 590). Precisely: we start with any joint state $\langle C, B \rangle$. The accuracy-dominance argument shows that if C is non-probabilistic, $\langle C, B \rangle$ is dominated by a state $\langle C^*, B \rangle$ with C^* probabilistic. Only then does it make sense to apply my expected accuracy argument to show that if $\langle C^*, B \rangle$ is not Lockean, it is exceeded in expectation by a state $\langle C^*, B^* \rangle$ that is. This is how we avoid Pettigrew's worry. (See Staffel (2017) for a more general response.)

possible benefit (T) is likely enough that it's worth the possible cost (F). (Example: if $T = 1$ and $F = -2$, the threshold is $\frac{2}{1+2} = \frac{2}{3}$.) Next, Accuracy First ensures that Rachael decides whether the bet is worth the risk for each proposition individually—deciding whether to bet on $p \wedge q$ is independent of whether to bet on p and on q . Finally, Separability_B ensures that this dynamic for individual propositions is maintained when all propositions are in view. Hence Lockeanism.

Conclusion: the standard assumptions of epistemic utility theory imply a fixed Lockean Thesis—there is a single threshold for all beliefs. That's a strong result. Too strong, I think. The next section shows why and how it should be generalized.

4 Lockeanism, generalized

Plan: §4.1 argues that dropping Extensionality_B is well-motivated and leads to more plausible results; §4.2 argues that we can drop Separability_B without loss.

4.1 Beyond Extensionality

Fixed Lockeanism is too strong: there are good reasons to think that the threshold for belief varies with proposition and context. Extensionality_B doesn't allow for this. So much the worse for Extensionality_B.

Why think the threshold varies? As Hawthorne et al. (2016) point out, saying what people think is a shifty endeavor—first some rough-and-ready data. I'm about to flip this coin three times; do I think it'll land heads at least once? Yup. (So the threshold must be *below* .875.) But a recent study concluded that there's a .9 chance that a new fundamental particle—the gloson—exists; do I think the world contains glosons? Not yet—we need more data. (So the threshold must be *above* .9.) Horses A, B, and C are racing, with betting ratios of .45, .30, and .25, respectively. Who do I think will win? Horse A, of course! (So the threshold must be *below* .5 (Hawthorne et al. 2016).) What's going on? Let's look at an example more carefully.

Peterson's class was harder than expected, so Steve is down in the dumps. How to cheer up? **Case 1:** Steve's bought 480 tickets in the next 1000-ticket lottery; there are 520 other players with 1 ticket each. Who do I think will win? I think Steve will win. Taking my ascription at face value, I believe *Steve wins* but not *Steve loses* despite being less than $\frac{1}{2}$ confident in the former and more than $\frac{1}{2}$ confident in the latter (.48 < .52). This calls out for *proposition*-dependence in the threshold: in this context *Steve wins* has a lower threshold for belief than *Steve loses*.

Case 2: Steve's bought 480 tickets, but Peterson's bought the remaining 520. Who do I think will win? I think Peterson (not Steve) will win. Taking my ascription at face value, in Case 2 I *don't* believe *Steve wins*, despite being just as confident in it (.48) as I was in Case 1. This calls out, further, for *context*-dependence in the threshold for believing this proposition—in Case 1 the threshold is below .48, while in Case 2 it is

not.¹⁴

(*Objection:* Perhaps these exchanges are felicitous for a pragmatic—not semantic—reason. Maybe I don’t, strictly speaking, think Steve will win in Case 1; but since the question ‘Who do you think will win?’ presupposes an answer, I accommodate. *Reply:* I *don’t* accommodate questions like ‘Who are you { $\begin{smallmatrix} \text{confident} \\ \text{sure} \end{smallmatrix} \}$ will win?’—instead I say, ‘I’m not { $\begin{smallmatrix} \text{confident} \\ \text{sure} \end{smallmatrix} \}$ of anyone!’ In contrast, I readily accommodate questions of the form ‘Who do you { $\begin{smallmatrix} \text{suspect} \\ \text{guess} \\ \text{think} \end{smallmatrix} \}$ will win?’ Diagnosis: the latter have weak enough presuppositions to be satisfied; the former don’t.)

We want to model and explain these natural language phenomena. Epistemic utility theory can help. First point: we should *expect* proposition- and context-dependence in the threshold for belief, given epistemic utility theory. Why? The threshold for believing something is generated from the value and disvalue of getting it right (*T*) and wrong (*F*). Yet it is independently plausible that the value of having a belief (an answer to a question) is a proposition- and context-dependent affair. A given inquiry is driven by a point, a purpose, a question—it is *directed* toward certain types of answers, and away from others. Thus different contexts of inquiry provide different *epistemic priorities*.

Example: Phyllis the Physicist is wondering whether black holes destroy information. I tell her that my garbage disposal destroys spoons. A truth, sure enough—but an unimportant one. It would be much more valuable (in this context) to learn that black holes don’t destroy information. However: across campus Dee the Detective is wondering how a murderer might destroy his murder weapon. I tell her that black holes don’t destroy information. A truth, sure enough—but an unimportant one. It would be much more valuable (in *this* context) to learn that my garbage disposal destroys spoons. Upshot: different contexts of inquiry give epistemic priority to different questions.

Incorporating this into epistemic utility theory requires making *A* proposition- and context-dependent.¹⁵ Does this make sense for a measure of accuracy? Yes (cf. Joyce 2013). Any measure of accuracy worthy of the name must have a preference for truths over falsehoods—but it needn’t be *indiscriminate* in that preference. Perhaps some questions (in some contexts) count more toward your overall accuracy than others. Recall the exam analogy: Rachael’s score is a measure of the accuracy of her answers—it goes up for correct answers, down for incorrect ones. But different questions might contribute more or less to her overall score than others; and the same question might contribute more or less on different exams.

Likewise: Rachael’s *A*-score is a measure of the accuracy of her beliefs—it goes up for true beliefs, down for false ones. But different propositions might contribute more or less to her overall *A*-score than others, and the same proposition might contribute more or less in different contexts. Example: Rachael writes a report on Happy Pill Pharmaceuticals. At a chemistry conference it’s praised for its accuracy; at an economics

¹⁴The cases are from Windschitl and Wells (1998), though they ask whether Steve will ‘probably’ win. (Cf. Yalcin 2010).

¹⁵Do these motivations force us to drop Extensionality_C for credences? Not obviously—though there may be contextual variation, I know of no data supporting proposition-dependence.

conference it's ridiculed for its inaccuracy. But the chemists and economists don't disagree about the facts, nor about the content of the report—they just have different priorities about *which* questions this kind of report should get right.

Upshot: epistemic utility theory provides reason to expect the proposition- and context-dependence in belief-ascriptions which we see in natural language.

More: it provides a concrete *model* for thinking about what happens in cases like Steve's lottery. *Context*-dependence really requires nothing new, so I won't wrap it in symbols. We simply acknowledge that A is not an 'ur-accuracy-metric' fixed for all time, but is instead generated within a given context by our priorities, interests, and questions. Since it will still obey the dynamics needed for Theorem 1 within any given context, we get a version of Lockeanism where thresholds for belief are generated by the (contextually-variable) values of T and F .

Proposition-dependence is more interesting: T and F must be *functions* that take a proposition p and output the value T_p of truly believing it and the disvalue F_p of falsely believing it (in a given context). That is, we now have a **proposition-dependent A** :

$$A(B, \{p\}, w) = \begin{cases} T_p & \text{if } p \text{ is true and believed} \\ F_p & \text{if } p \text{ is false and believed} \\ 0 & \text{if } p \text{ is not believed} \end{cases}$$

The upshot is a *variable* Lockean Thesis: for each proposition there's a threshold of credence $t = \frac{-F_p}{T_p - F_p}$ necessary and sufficient for belief.¹⁶ Formally:

Corollary 2. (Easwaran 2016, §G.2) If A is proposition-dependent and separable, a belief-state B maximizes expected accuracy iff, for all p :

- (i) If $p \in B$, then $C(p) \geq \frac{-F_p}{T_p - F_p}$
- (ii) If $p \notin B$, then $C(p) \leq \frac{-F_p}{T_p - F_p}$

The flexibility granted by Corollary 2 allows us to generate a concrete model of what happens in our lottery cases. Recall Case 1: Steve has 480 tickets and 520 other players have 1 ticket each; when asked, 'Who do you think will win?' I answer, 'I think Steve will win.' Proposal: *asking a question* (posing a new stage of inquiry) is a way of shifting epistemic priorities; in particular, it shifts the values of T_p and F_p for various propositions.

Model it thus. Suppose that before the question is asked the (dis)values of believing propositions of the form x wins and x loses ($= \neg(x$ wins)) are the same. Say $T_{x \text{ wins}} = T_{x \text{ loses}} = 1$ while $F_{x \text{ wins}} = F_{x \text{ loses}} = -2$, generating thresholds at $\frac{2}{1+2} = \frac{2}{3}$; therefore I believe neither that Steve will win nor that he'll lose (since $\frac{2}{3} > .48, .52$). Then you ask 'Who do you think will win?' The question presupposes an answer, so 'straightaway that presupposition comes into existence' (Lewis 1979, 339).

The question prioritizes answers of the form x wins; it's a way of saying, 'Don't worry too much about being wrong about who'll win—just tell me what you're inclined

¹⁶ *Worry*: is our Lockean result still substantive if it allows variation in the threshold? Yes—see §5.

to think.’ Formally, it makes the penalty $F_{x \text{ wins}}$ (of falsely believing that someone will win) less extreme. This has the effect of *lowering* the threshold for propositions of the form $x \text{ wins}$. Since the question presupposes an answer, it adjusts them (within reason) until it gets one. In particular, once $F_{x \text{ wins}}$ shrinks from -2 to $-.923$, the threshold becomes $\frac{-.923}{1+.923} < .48$ —I now believe *Steve wins*. According to epistemic utility theory, that’s how asking the question ‘Who do you think will win?’ can lead me to truly say ‘I think Steve will win’—despite the fact that I’m less than $\frac{1}{2}$ confident of it. That’s Case 1.

Case 2 is similar. Here Steve has 480 tickets but Peterson has the remaining 520. At first I don’t believe of either that they’ll win. When you ask, ‘Who do you think will win?’ the same mechanism kicks in, lowering $F_{x \text{ wins}}$ until it gets an answer. But this time around it only needs to be shifted until the threshold for $x \text{ wins}$ drops below .52, for then I can truly respond, ‘I think Peterson (not Steve) will win’—so there’s no need to drop the threshold below .48.

In fact, it *couldn’t* drop that low. If it did, I’d believe both that Peterson will win and that Steve will win (= Peterson *won’t* win). Yet epistemic utility theory won’t allow me to believe both p and $\neg p$ —even once we allow the threshold to fall below $\frac{1}{2}$.

Recall that this comes from the Conservativeness constraint that rational agents have ‘more to lose than to gain’ in becoming opinionated on a given proposition. Originally this said that the value of truly believing something is smaller than the disvalue of falsely believing it: $|T| < |F|$. But now that we have proposition-dependence, the value of getting p right or wrong (T_p or F_p) might be different than the value of getting its negation right or wrong ($T_{\neg p}$ or $F_{\neg p}$). Which one you’ll get will depend on whether p or $\neg p$ is true; our constraint should guarantee that either way you should be conservative. If p is true, becoming opinionated will yield either T_p or $F_{\neg p}$, depending on whether you believe p or $\neg p$. And if $\neg p$ is true, becoming opinionated will yield either $T_{\neg p}$ or F_p . Thus **Variable Conservativeness** requires *these* values to be coordinated: $|T_p| < |F_{\neg p}|$ and $|T_{\neg p}| < |F_p|$.

This guarantees that when an inquiry prioritizes p (makes T_p larger or F_p smaller), it thereby *deprioritizes* $\neg p$ (makes $F_{\neg p}$ larger or $T_{\neg p}$ smaller, respectively). Thus whenever the threshold for p drops below $\frac{1}{2}$, the threshold for $\neg p$ raises correspondingly *above* $\frac{1}{2}$. So even when believing with below $\frac{1}{2}$ probability, it’ll never be rational to believe both p and $\neg p$:

Proposition 3. Given Variable Conservativeness and Corollary 2, if the threshold for p is t , then the threshold for $\neg p$ is greater than $1 - t$. Thus it never maximizes expected accuracy to believe both p and $\neg p$.

Upshot: epistemic utility theory doesn’t merely motivate proposition- and context-dependence in the abstract—it provides a concrete model of particular cases of it, without leading to bad results. Further, we can use this model to help inform our semantic theorizing. Here I’ll just sketch a few applications.

Observation: Return to Case 1 where Steve has 480 tickets and 520 others have 1

ticket each. As already noted, if you ask ‘Who do you think will win?’ I’ll answer ‘I think Steve will win.’ But if instead you ask, ‘Do you think Steve will win or lose?’ I’ll answer ‘I think he’ll lose.’

Explanation: Different questions set different epistemic priorities for the context. In particular, asking ‘Do you think Steve will win or lose?’ sets equal (dis)value to getting *Steve wins* and *Steve loses* right (wrong), so $T_{\text{Steve wins}} = T_{\text{Steve loses}}$ and $F_{\text{Steve wins}} = F_{\text{Steve loses}}$. Since your question presupposes an answer, I believe the one with higher probability—namely, *Steve loses*.

Observation: In **Case 3** Steve has 5 lottery tickets and there are 995 other people with one ticket each. If you ask, ‘Who do you think will win?’ I’ll respond, ‘I have no idea!’ As Yalcin (2010) and Hawthorne et al. (2016) observe, this means that we can’t explain Case 1 (where Steve has 480 tickets) using the fact that he’s the *most* likely to win—for that holds in Case 3 as well.

Explanation: Epistemic utility theory predicts this. For it’s much *easier* to shift epistemic priorities to get the threshold for *x wins* down to .48 than to get the threshold down to .005—the former merely requires $|F_{x \text{ wins}}|$ to be a *bit* less than $|T_{x \text{ wins}}|$, while the latter requires it to be almost *two hundred times* less ($\frac{-F_p}{T_p - F_p} \leq .005 \Rightarrow |F_p| \leq \frac{|T_p|}{199}$). So in Case 3 if you ask simply, ‘Who do you think will win?’ it’s reasonable to assume that you’ve made a mistake about the likelihoods, rather than request such an extreme change in priorities.

Observation: But you *can* make clear that such extreme changes are what you’re asking for. If in Case 3 you ask, ‘I just want your best guess—who do you think will win?’ I’ll respond, ‘Steve, I guess.’

Explanation: Saying ‘I just want your best guess’ is a way of forcing the penalty $F_{x \text{ wins}}$ for a false belief very low, so I *do* accommodate the presupposition of your question.

Upshot: dropping Extensionality_B is well-motivated, and doing so fits epistemic utility theory with natural language considerations to offer a new picture of belief—a version of proposition- and context-dependent Lockeanism. On this picture, your beliefs are your *best guesses*—your best shots at the truth—in response to the epistemic priorities of your context. Moreover, our two approaches are mutually illuminating. The thresholds epistemic utility theory generates from the (dis)values of true (false) beliefs help to predict, model, and explain patterns in natural language belief-ascriptions. Conversely, the natural-language phenomena of presupposition accommodation and question-setting help to explain how the (dis)values of true (false) beliefs are modulated by contextual priorities. Though many questions remain, we can see that this new picture might benefit both epistemology and semantics.

4.2 Beyond Separability

Dropping Extensionality_B is progress, but the real bugbear is Separability_B—the assumption that the value of a belief-state is the sum of the value of its parts. It seems to

contradict epistemic holism (Joyce 2009, §5); so can holists avoid our Lockean picture? No. So long as we retain the core assumptions of epistemic utility theory, dropping Separability_B yields the same results. More precisely: assuming (1) a non-arbitrariness condition on beliefs, (2) Accuracy First, and (3) Maximize Expected Accuracy, we still get Lockeanism. (We must now get unavoidably technical; if you trust me on the details, feel free to skim/skip this section.)

Begin by granting the objector what he wants: A depends on Rachael's current belief-state B_0 (to capture the sense in which it matters what other beliefs she holds), and the value (accuracy) of any given belief-state B cannot be broken down into the value of the various propositions it contains. This blocks the straightforward route to Lockeanism we've been using.

Nevertheless, we can be devious. Accuracy First ensures that the space of doxastic alternatives is open for exploration—it makes sense to compare the value of Rachel's current belief-state B_0 with another state B_1 that adds a belief in p ($B_1 = B_0 \cup \{p\}$). On our current supposition, the value of these two belief-states will depend on what her current beliefs are. Nevertheless, by the existence of an accuracy metric, there will be *some* function A_{B_0} that assigns values to them. The value of her current belief-state B_0 (relative to her beliefs B_0) is simply $A_{B_0}(B_0, \mathcal{P}, w)$. The value of B_1 can be divided into three cases. (i) If $B_1 = B_0$ (because $p \in B_0$), then its value is identical to B_0 's: $A_{B_0}(B_1, \mathcal{P}, w) = A_{B_0}(B_0, \mathcal{P}, w)$. (ii) If $B_1 \neq B_0$ and p is *true* then B_1 adds one true belief to B_0 , so B_1 will have a higher score: $A_{B_0}(B_1, \mathcal{P}, w) > A_{B_0}(B_0, \mathcal{P}, w)$. Since A_{B_0} is a real-valued function, there must be some $t > 0$ such that $A_{B_0}(B_1, \mathcal{P}, w) = A_{B_0}(B_0, \mathcal{P}, w) + t$. Write t as $T_p^{B_0}$ —the value, relative to B_0 , of adding a true belief in p . (iii) Finally, if $B_1 \neq B_0$ and p is *false* then B_1 adds one false belief to B_0 . By parallel reasoning, there will be some number $f < 0$ such that $A_{B_0}(B_1, \mathcal{P}, w) = A_{B_0}(B_0, \mathcal{P}, w) + f$. Write f as $F_p^{B_0}$ —the disvalue, relative to B_0 , of adding a false belief in p . Combining (i)-(iii), we've derived this special case accuracy-comparison between B_0 and B_1 :

$$A_{B_0}(B_1, \mathcal{P}, w) = \begin{cases} A_{B_0}(B_0, \mathcal{P}, w) & \text{if (i) } B_1 = B_0 \\ A_{B_0}(B_0, \mathcal{P}, w) + T_p^{B_0} & \text{if (ii) } B_1 \neq B_0 \text{ and } p \text{ is true} \\ A_{B_0}(B_0, \mathcal{P}, w) + F_p^{B_0} & \text{if (iii) } B_1 \neq B_0 \text{ and } p \text{ is false} \end{cases}$$

It's crucial to emphasize that this assumes only that A_{B_0} is well-defined. Since the value of the move from B_0 to B_1 depends on the beliefs Rachael already has, there's no way to string together applications of this constraint to recover Separability_B.

However, more deviousness. Accuracy First allows Rachael to form her final belief-state B_n *in sequence*—starting from nothing ($B_0 = \emptyset$) and deciding whether to add each proposition in some sequence p_1, \dots, p_n . (She needn't actually go through this process. The point is not that she can form beliefs one at a time, but that she can assess the value of each state B_{i+1} from the perspective of a less opinionated B_i . The process may be purely imaginative—but for illustrative purposes, I'll talk as if she actually goes through it.) By Maximize Expected Accuracy, her final state B_n will be rational only

if she maximizes her expectation at the previous stage B_{n-1} . Given our constraint on A_{B_i} , this leads to an (uninteresting) special case of Lockeanism: *if* she forms beliefs in this sequence, then Maximize Expected Accuracy ensures that the final proposition p_n has a threshold necessary and sufficient for belief. Formally:

Proposition 4. If A_{B_i} is belief- and proposition-dependent and B_1, \dots, B_n are formed on the sequence p_1, \dots, p_n , then B_n maximizes expected accuracy with respect to B_{n-1} only if:

- (i) If $p_n \in B_n$, then $C(p_n) \geq \frac{-F_{p_n}^{B_{n-1}}}{T_{p_n}^{B_{n-1}} - F_{p_n}^{B_{n-1}}}$
- (ii) If $p_n \notin B_n$, then $C(p_n) \leq \frac{-F_{p_n}^{B_{n-1}}}{T_{p_n}^{B_{n-1}} - F_{p_n}^{B_{n-1}}}$

Though this is *a* way for Rachael to form beliefs, it's an uninteresting special case. But—and here's the kicker—*it shouldn't matter* what sequence Rachael forms her beliefs in, or whether she does so in any sequence at all. In the special case where she forms her beliefs sequentially, how could renumbering the propositions make a difference to what it's rational to believe? That'd be like changing the numbering of the questions on an exam and thereby changing which answers are rational to give. That shouldn't happen; hence the following constraint:

Sequence-Irrelevance: Holding fixed all other factors, rational constraints on beliefs are not affected by the sequence (if any) in which they are formed.

Sequence-Irrelevance is stated in terms of rationality—not expected accuracy. Everyone should accept it. To deny it is to allow that rational belief is to a surprising extent an arbitrary affair—it depends on Rachael's choices about how and when and whether to consider various questions. Suppose it's not rational for her to believe in God, yet she thinks she'd be happier if she did. If Sequence-Irrelevance fails, it's possible that she could design a sequence that *makes* it rational for her to believe. That's not possible. So Sequence-Irrelevance is true.

But now it turns out that denying Separability_B has gained us nothing: Sequence-Irrelevance turns the uninteresting Proposition 4 into a substantive Lockean result.

Proposition 5. Given Sequence-Irrelevance, even if A_{B_i} is belief- and proposition-dependent, there are belief-*independent* values T_p and F_p for each proposition such that B is rational only if:

- (i) If $p \in B$, then $C(p) \geq \frac{-F_p}{T_p - F_p}$
- (ii) If $p \notin B$, then $C(p) \leq \frac{-F_p}{T_p - F_p}$

That is, we return to a variable (normative) Lockean Thesis: Rachael is rational only if, for each proposition p , she believes p iff p is above a given threshold of credence.

The connection between epistemic utility theory and the Lockean Thesis is robust. Even a minimal version of the theory—which relies only on the existence of A , Maximize Ex-

pected Accuracy, and Accuracy First—leads to Lockeanism. Moreover, recall that these are all *core* assumptions of the theory: the first comes from the existence of an accuracy-metric, the second from the norm to promote accuracy by your own lights, and the third from the attempt to systematically derive rational norms from such considerations. Of course, once we drop Separability_B the result also relies on Sequence-Irrelevance. But this assumption is plausible; and—perhaps more importantly—it shows that *dropping Separability_B is not the way to respect holism*. After all, no holist ever thought that the insight of their position is that the *order* in which you consider propositions affects what you ought to believe. So *if* there’s a tension between epistemic utility theory and holism—and I’m not saying that there is—it’s somewhere in the core assumptions of the theory, not in Separability_B (contra Joyce 2009, §5).

Upshot: the bare epistemic utility framework—the simple idea that beliefs are valuable if true and disvaluable if false—is enough to generate our new Lockean picture.

5 Significance?

It’s been a long road. We took a fresh look at belief through natural language—Lockeanism looked natural. We developed an approach to the issue through epistemic utility theory—Lockeanism looked inevitable. We realized that the threshold for belief varies with proposition and context—and found that epistemic utility theory explains why. We showed that these results are robust—they follow from the core of the epistemic utility approach. We are almost home.

But wait. Does something smell fishy? A normative Lockean Thesis says that high credence is necessary and sufficient for rational belief. For a *fixed* threshold this claim is clear and substantive. It is also false—hence our move to variable Lockeanism. But has this move sapped our conclusion of content? We say that for each proposition-context pair, there’s a threshold necessary and sufficient for rational belief. Does that say anything at all? Give me *any* profile of beliefs you like, and it looks like I’ll be able to give you a function from contexts-and-propositions to thresholds that fits it. You believe *It’ll rain and snow* but not *It’ll snow*? No problem: set the threshold for the conjunct higher than that for the conjunction. You believe in glosons on Mondays, but not on Tuesdays? No problem: set the threshold low for Monday-contexts and high for Tuesday-contexts. Etc. Conclusion: our (normative) variable Lockean Thesis is no thesis at all! Right?

Wrong. Granted, *on its own* normative variable Lockeanism says very little. But our Lockeanism doesn’t stand alone—so it says quite a bit.

Reason number one: as we’ve seen, epistemic utility theory and natural language considerations provide a story for *how* the Lockean thresholds get set. They come from the (dis)values of true (false) beliefs, which in turn are modulated by contextual priorities and subject to formal constraints. The picture that emerges is that your beliefs are your *best guesses* at the truth in response to the epistemic priorities of your context.

If this is correct, beliefs play a fundamentally different role than many have thought.

For Rachael’s best guess—what she thinks—is that it’ll rain (based on the 70% chance); that you’ll like the starfruit (since most people do); and that Steve will win the lottery (since he has 480 tickets). But in doing so she doesn’t *rule out* that it’ll be sunny (contra Hintikka 1962); she has merely statistical evidence (contra Buchak 2014; Staffel 2016); she wouldn’t assert ‘It’ll rain’ (contra Kaplan 1996); she doesn’t aim at knowledge (contra Williamson 2000); it wouldn’t be *abnormal* if you don’t like starfruit (contra Stalnaker 2006; Smith 2016; Goodman and Salow 2017); she wouldn’t use Steve’s winnings as a *premise* in her reasoning (contra Frankish 2009; Ross and Schroeder 2014); and conditionalizing on these beliefs would radically change her behavior (contra Weatherston 2005). Upshot: our normative, variable Lockean Thesis forces us to rethink the role of beliefs in assertion, reasoning, and inquiry.

So what we’ve said is significant. And we can say more.

5.1 Normative metaphysics

Reason number two: our normative result sets up a metaphysical conclusion.

Begin with a parable. Billy is an up-and-coming philosopher of mind who tells us that there’s a heretofore undiscussed, irreducible mental attitude had by most agents—he calls it ‘thredence’. When asked to explain, he replies:

‘The main effect of a thredence that p is to make it reasonable to take bets on p that have roughly 2 to 1 odds; to say ‘probably $\neg p$ ’; etc. In fact, you have a thredence that p iff you have a credence that p in the [.3, .4) range.’

Billy is silly, I hope you agree? Thredence is not a new attitude at all—Billy has simply coined a new term for picking out a particular range of *credence*.

The point of the parable is a methodological one about positing and individuating mental attitudes; call it my **Pragmatist Premise**. It says that the primary reason for positing such attitudes is the explanation, prediction, and rationalization of the dynamics of rational agents. It’s a form of functionalism: mental attitudes earn their keep in our metaphysics of mind by playing a functional role in the life of an agent. After all, it’s not as if we’re going to open up the heads of agents and find beliefs, credences, and desires. We know exactly what we’ll find—neurons and synapses; the point of attitude ascriptions is to characterize the dynamics these underlying mechanisms give rise to. The pragmatic picture is one on which:

Rational creatures are essentially agents. Representational mental states should be understood primarily in terms of the role that they play in the characterization and explanation of action... And, according to this picture, our conceptions of beliefs and of attitudes pro and con are conceptions of states which explain why a rational agent does what he does. (Stalnaker 1984, 4)

So the problem with Billy’s ‘thredences’ is that positing them as an irreducible attitude doesn’t *do* anything for us. Facts about thredences can simply be ‘read off’ of facts about credences: if you’re between .3-.4 confident, you have a thredence; and if not, then not. We can give all the same explanations of the dynamics of rational agents

in terms of credences, without the middle-man. Conclusion: there is no irreducible attitude of so-called ‘thredence.’

You know where this is headed. Given our Lockean results, the Pragmatist Premise allows us to conclude something similar about beliefs. Epistemic utility theory entails that—at least for rational agents—positing beliefs as an irreducible attitude doesn’t *do* anything for us. Facts about beliefs can be simply ‘read off’ of facts about credal states and contextual priorities: if you’re at least so-and-so confident, then you believe; and if not, then not. We can give all the same explanations of the dynamics of rational agents in terms of credences, without the middle-man. Conclusion: at least for rational agents, there is no irreducible attitude of so-called ‘belief.’ This is how we’ll try to pull a metaphysical rabbit out of a normative hat. I anticipate resistance.

Objection: The result holds only for ideal agents. Non-ideal ones still need beliefs.

Reply: Not so. Granted, ideal agents are easy. Their credences and utilities are numerically precise, and they always maximize expectations. Granted, non-ideal agents are messy. Their credences and utilities are rough, and they sometimes make poor decisions—even by their own lights. But how to model, understand, and explain such imperfections is a wide-open question. If we had compelling reason to invoke (irreducible) beliefs to explain the quirks of non-ideal agents, then my reductive proposal would be a non-starter. But we don’t: beliefs are neither necessary nor sufficient to do the needed explanatory work. Not *sufficient* because beliefs can’t explain all our non-ideal behavior: witness the fact that my indifference to a bet is maintained even after you add \$1 to the payout. Sharp credences plus beliefs won’t explain that—to do so, we’ll need to say I have (something like) imprecise credences. But once we make *that* move, adding beliefs to non-ideal agents is no longer even *necessary* to explain my non-ideal behavior.

Here’s why. Invoking imprecise credences is an instance of a much more general strategy for modeling and explaining non-ideal agents: using *collections* of ideal doxastic states to do so. One form this takes is fragmentation: when Steve seems to think both that Main Street runs north-south and that it runs east-west, model him as having two fragmented doxastic states—one with each belief, but neither with both (Lewis 1982; Stalnaker 1984; Rayo 2013). A second form is imprecise attitudes: an agent’s degrees of belief and values are represented with *sets* of sharp credence and utility functions (e.g. Joyce 2010; Hare 2010). We can easily imagine variations on this theme; the general idea is that our models and explanations of non-ideal agents should be continuous with those of ideal ones. This approach has much more flexibility than simply adding beliefs to non-ideal agents.

Moreover, it vindicates our metaphysical reduction. At a high level: since our idealized explanations don’t invoke beliefs, neither should our non-idealized ones. More concretely: say we’re using a set \mathcal{D} of ideal doxastic states to represent my non-ideal one. Though it can be tricky to figure out how exactly to infer the properties of my non-ideal state from the members of \mathcal{D} , if *all* of the members of \mathcal{D} have some (appro-

priate¹⁷) property, then my state does as well. (Example: all members of \mathcal{D} are more confident of p than q ; therefore I am more confident of p than q .) But given that our Lockean result holds for the ideal states in \mathcal{D} , the Pragmatist Premise entails that *none* of those states will have beliefs. So there are no (irreducible) beliefs in \mathcal{D} , which is what models and explains my non-ideal state. Thus there are no (irreducible) beliefs in me. The reduction should extend to non-ideal agents.

Objection: Beliefs can't be reduced to credences, since they have different normative contours. A belief is either completely right (if true) or completely wrong (if false), whereas the state of having credence above a threshold always falls in between (cf. Fantl and McGrath 2010, 141).

Reply: Differing normative contours do not imply non-reduction—they may simply reflect different modes of evaluating the same underlying state. Example: the best *grade* you can get on the exam is an *A*; the best *score* you can get is a 100. Rachael got an *A* by getting a 98. In one sense, she did as well as possible—the options are *A, B, C...*, and she got an *A*. In another sense, she could have done better—the options are 100, 99, 98, ..., and she got a 98. But this doesn't show that grades are irreducible to scores! It just shows that, given an exam, we can evaluate the outcome in two different ways: By asking, 'Did she get the highest grade?' Or by asking, 'Did she get the highest score?'

Likewise: the most accurate *belief* you can have is a belief that it'll rain; the most accurate *credence* you can have is credence 1. Rachael believes it'll rain by having a credence of .98. In one sense, she's as accurate as possible—the options are believing *Rain*, \neg *Rain*, or neither, and she believes *Rain*. In another sense, she could have been more accurate—the options are credence 1, .99, .98, ..., and she has .98. But this doesn't show that beliefs are irreducible to credal states. It just shows that, given a doxastic state, we can evaluate its accuracy in two different ways: By asking, 'Did she have the most accurate belief?' Or by asking, 'Did she have the most accurate credence?'

Objection: Beliefs play all sorts of theoretical roles that high credence can't. By *ruling out* possibilities, they simplify and guide our activities of assertion, reasoning, and inquiry (Weisberg 2017; Holton 2017).

Reply: Do they? My main claim is about beliefs—that is, about the attitude that the word 'belief' refers to in the natural language we're using to have this debate. The picture that emerges: you believe whatever you're sufficiently confident in—and in some contexts 'sufficient' need not be very confident at all. Your beliefs are your *best guesses* (in context). So natural-language and epistemic utility theory suggest that belief is not, after all, what fills these theoretical roles.

That's all a traditional, Bayesian Lockean will say. Many do. You might think that's all *any* Lockean can say—namely, that knowledge-theoretic and traditional epistemology must fall by the wayside. But you'd be wrong. Less radical Lockeans must simply show how credal states can play the desired theoretical roles. We require an attitude toward p that...

¹⁷'Appropriate' to rule out, say, the property of being ideal.

- (1) ...must be deductively closed and consistent;
- (2) ...rules out $\neg p$ (Hintikka 1962);
- (3) ...demands more than statistical evidence (Buchak 2014);
- (4) ...warrants assertion in a context of inquiry (Kaplan 1996);
- (5) ...aims at knowledge (Williamson 2000);
- (6) ...would make you surprised if $\neg p$ (Stalnaker 2006; Smith 2016);
- (7) ...warrants premising that p (Frankish 2009; Ross and Schroeder 2014); and
- (8) ...warrants acting as if p (Weatherson 2005; Ganson 2008).

Granted: Lockean belief cannot play these roles. But here's a state that can: being sure, or being certain, or having credence 1, or taking for granted, or taking yourself to know—all the same (context-sensitive) state; all the same (context-sensitive) roles. Though I can't adequately defend that claim here, I follow the proposals of Clarke (2013, 2017) and Greco (2017)—modulo their use of 'belief'.

To get a feel for the position, consider: (1) Probability one is deductively closed and consistent. (2) If you rule out $\neg p$, then all the open possibilities are p -worlds—you're sure of p . (3) Mere statistics cannot lead to certainty. (4) 'It'll rain, but I'm not sure it will' has a Moore-paradoxical ring to it.¹⁸ And if you assert 'It'll rain' you update the common ground to include only rain-worlds (Stalnaker 1978); but the common ground in a context of inquiry is the set of possibilities you distribute credence over—so you must be sure it'll rain. (5) Another Moore-paradoxical line: 'I don't know if it'll rain, but I'm sure it will.' (6) Nothing says 'Surprise!' like finding out that $\neg p$ when you were certain that p . (7)/(8) And what things can you use as premises—or conditionalize on—without going beyond your information? The things you're already certain of. Upshot: Lockeanism *does* have the resources to fill the desired theoretical roles. We need not look beyond credence—we simply need look to the top of the scale.

In fact, there's a stronger reading of our results. *Question:* why is certainty suited to play our theoretical roles, while belief is not? The contrast is that belief falls on a coarsely categorized scale—you either believe p , or $\neg p$, or neither—whereas certainty falls on a *continuous* scale of credences. This is crucial. For no matter the underlying values, on a coarse scale there are inevitably *threshold points* at which it becomes worth the risk to 'plump' for the strongest attitude. On a continuous scale, in contrast, you can always hedge your bets—no plumping is needed.

Conjecture: this structural contrast is fully general. Take the attitude—whatever it is—that plays theoretical roles (1)-(8). Call it 'Xing that p '. As we've seen, Xing that p cannot reduce to being sufficiently confident of some proposition q . But our results suggest that if Xing falls on a *coarsely* categorized scale, then there's no way around it—when you're sufficiently confident of q , it'll be rational to 'plump' for Xing that p . Therefore Xing that p and Xing that $\neg p$ cannot fall on a coarse scale: they must fall on the extreme ends of a *continuous* range of attitudes—a range of attitudes eerily reminiscent of degrees of confidence.

¹⁸Notice: 'I $\left\{ \begin{smallmatrix} \text{think} \\ \text{believe} \end{smallmatrix} \right\}$ it'll rain, but I'm not sure it will' sounds fine.

6 Beyond Belief

Home, at last. The first claim I've argued for is a conditional: *if* you like the basics of epistemic utility theory, then you're committed to some version of Lockeanism. The second is that this result meshes with natural language to offer a proposition- and context-sensitive Lockean picture: your beliefs are your best guesses in response to the epistemic priorities of your context. The third is that it further supports a metaphysical reduction of belief to credence: your beliefs *just are* the things you're sufficiently confident in.

I find these results compelling. The reason is that we have these two disparate sources of considerations—natural language and epistemic utility theory—that converge and cohere around a particular picture of belief. On its own, neither would force our hand. Natural language might be thought too commonsensical to address theoretical questions about our epistemic life. Epistemic utility theory might be thought too theoretical to address commonsense questions about what we 'think.' It is the confluence of these two sources that warrants our attention—and, perhaps, our conviction.

Suppose it gets them. Why care about beliefs, on the resulting picture? In one sense, we philosophers should care less than we have, for it turns out beliefs don't play the key theoretical roles that they are often thought to—credences have supplanted them. Instead, the philosopher's notion of 'outright belief' is a mongrel, mixing the natural language state of 'sufficient confidence' with the philosophically important state of *being sure*. Much confusion has arisen, I think, from invoking one state to play both roles. At the very least, then, we should care about beliefs so that we know what they are not.

In another sense, though, we should care about beliefs just as much as we always have. We *do* care about what people think—about what they're sufficiently confident in—and that's no mystery. Consider: being 'flat' reduces to being flat enough (in context); being 'busy' reduces to being busy enough (in context); and being 'heavy' reduces to being heavy enough (in context). But none of these concepts are less important for that—to the contrary! We care about whether the roads we drive on are flat, the semesters we teach are busy, and the boxes we lift are heavy. Sure, there are more extreme questions we could ask—Are they *completely* flat? Are they *extremely* busy? Are they *very* heavy? But usually we don't—'flat' is flat enough. Likewise: believing a proposition reduces to being sufficiently confident in it (in context). We care about whether the forecaster is confident of rain, the congresswoman believes in climate change, and our colleagues think we smell. Sure, there are more extreme questions we could ask—Is he *certain* it'll rain? Is she *convinced* of climate change? Are they *sure* we smell? But usually we don't—'sufficiently confident' is confident enough. It's no more mysterious why we talk about (contextually variable) sufficient confidence than why we talk about (contextually variable) sufficient flatness. And it's no less obvious why we care.¹⁹

¹⁹For much helpful discussion and feedback, thanks to Alex Byrne, Kenny Easwaran, Branden Fitelson, Cosmo Grant, Caspar Hare, Samia Hesni, Sophie Horowitz, Pavel Janda, Matthias Jenny, Richard Pettigrew, Milo Phillips-Brown, Agustín Rayo, Bernhard Salow, Ginger Schultheis, Kieran Setiya, Brad

7 Appendix

Theorem 1. (Easwaran 2016, §3.2) If A is separable and extensional, a belief-state B maximizes expected accuracy iff, for all p :

- (i) If $p \in B$, then $C(p) \geq \frac{-F}{T-F}$
- (ii) If $p \notin B$, then $C(p) \leq \frac{-F}{T-F}$

Proof. Recall that Extensionality _{B} ensures that our scoring rule for B for a single proposition $\{p\}$ gives T , F , or 0 depending on whether p is truly believed, falsely believed, or not believed. We can combine this with Separability _{B} to recover a scoring rule for an arbitrary set of propositions P . In particular, we have $A(B, P, w) = \sum_{p_i \in P} A(B, \{p_i\}, w)$, i.e. the sum of T s, F s, and 0s got by our local scoring rule for each p_i . Letting \bar{w} be the set of true propositions at w ($\bar{w} = \{p \subseteq W | w \in p\}$), we can partition P into those the propositions believed truly at w ($P \cap B \cap \bar{w}$), those believed falsely at w ($(P \cap B) - \bar{w}$), and those not believed ($P - B$). Each member of $P \cap B \cap \bar{w}$ adds T value, each member of $(P \cap B) - \bar{w}$ adds F disvalue, and each member of $P - B$ adds 0. Hence our global scoring rule:

$$A(B, P, w) = T|P \cap B \cap \bar{w}| + F|(P \cap B) - \bar{w}|$$

We use this in the ensuing proof.

\Rightarrow : Suppose B maximizes expected accuracy (for the whole algebra \mathcal{P}). Taking an arbitrary p , first suppose (i) $p \in B$. Then for $B^- = B - \{p\}$ we have $EA(B, \mathcal{P}) \geq EA(B^-, \mathcal{P})$. Now,

$$EA(B, \mathcal{P}) = \sum_{w_p \in p} (C(w_p) \cdot A(B, \mathcal{P}, w_p)) + \sum_{w_{\bar{p}} \in \neg p} (C(w_{\bar{p}}) \cdot A(B, \mathcal{P}, w_{\bar{p}}))$$

Since $p \in B$, it gets T for p at w_p -worlds and F for p at $w_{\bar{p}}$ -worlds:

$$\begin{aligned} &= \sum_{w_p \in p} (C(w_p) \cdot (A(B, \mathcal{P} - \{p\}, w_p) + T)) + \sum_{w_{\bar{p}} \in \neg p} (C(w_{\bar{p}}) \cdot (A(B, \mathcal{P} - \{p\}, w_{\bar{p}}) + F)) \\ &= EA(B, \mathcal{P} - \{p\}) + C(p)T + C(\neg p)F \end{aligned} \tag{1}$$

Note that since B and B^- agree on all propositions in $\mathcal{P} - \{p\}$, we have $EA(B, \mathcal{P} - \{p\}) = EA(B^-, \mathcal{P} - \{p\})$. And since $p \notin B^-$, $A(B^-, \{p\}, w) = 0$ at all worlds, hence $EA(B^-, \mathcal{P} - \{p\}) = EA(B^-, \mathcal{P})$. Thus (2) $EA(B, \mathcal{P} - \{p\}) = EA(B^-, \mathcal{P})$. Substituting into (1) yields:

$$EA(B, \mathcal{P}) = EA(B^-, \mathcal{P}) + C(p)T + C(\neg p)F$$

Skow, Julia Staffell, Bob Stalnaker, Quinn White, Roger White, and Steve Yablo, as well as audiences at MIT, the 2015 Central APA, and the 2015 Epistemic Consequentialism conference at Kent.

Since $EA(B, \mathcal{P}) \geq EA(B^-, \mathcal{P})$, we know $C(p)T + C(\neg p)F \geq 0$. This holds iff $C(p)T + (1 - C(p))F \geq 0$, which simplifies to $C(p) \geq \frac{-F}{T-F}$, as desired.

Next, if (ii) $p \notin B$, then parallel reasoning yields that for $B^+ = B \cup \{p\}$:

$$EA(B^+, \mathcal{P}) = EA(B, \mathcal{P}) + C(p)T + C(\neg p)F$$

Since $EA(B^+, \mathcal{P}) \leq EA(B, \mathcal{P})$, we know $C(p)T + C(\neg p)F \leq 0$, and parallel reasoning gives the result that $C(p) \leq \frac{-F}{T-F}$, as desired.

\Leftarrow : Suppose (i) and (ii) hold and, for reductio, that B does not maximize expected accuracy. So there is a B^* such that $EA(B^*, \mathcal{P}) > EA(B, \mathcal{P})$. By Separability_B there must be a p such that $EA(B^*, \{p\}) > EA(B, \{p\})$. If p is in both or neither of B, B^* this is impossible. So either (1) $p \in B^*$ and $p \notin B$, or (2) $p \notin B^*$ and $p \in B$. If (1) since $p \notin B$ then by (ii) $C(p) \leq \frac{-F}{T-F}$. But by supposition $0 = EA(B, \{p\}) < EA(B^*, \{p\}) = C(p)T + C(\neg p)F$, which implies that $C(p) > \frac{-F}{T-F}$. Contradiction. If (2) $p \notin B^*$ and $p \in B$, then by (i) $C(p) \geq \frac{-F}{T-F}$; yet $0 = EA(B^*, \{p\}) > EA(B, \{p\}) = C(p)T + C(\neg p)F$, which implies that $C(p) < \frac{-F}{T-F}$. Contradiction again; so we reject our hypothesis. \square

Corollary 2. (Easwaran 2016, §G.2) If A is proposition-dependent and separable, a belief-state B maximizes expected accuracy iff, for all p :

- (i) If $p \in B$, then $C(p) \geq \frac{-F_p}{T_p - F_p}$
- (ii) If $p \notin B$, then $C(p) \leq \frac{-F_p}{T_p - F_p}$

Proof. We begin by partitioning the propositions \mathcal{P} into equivalence-classes based on the values of T and F : $\mathcal{P} = E_1 \cup \dots \cup E_m$, where the $p, q \in E_i$ iff $T_p = T_q$ and $F_p = F_q$. We then partition B into correlative equivalence-classes $B = B_1 \cup \dots \cup B_m$ where $p \in B_i$ iff $p \in B$ and $p \in E_i$. Given this, our accuracy function A will be constant across propositions within each E_i and B_i , therefore each E_i and B_i corresponds to a total set of propositions and a total belief-state used in Theorem 1. By Separability_B, B maximizes expected accuracy for \mathcal{P} iff each B_i maximizes expected accuracy for E_i ; so m applications of Theorem 1 yields the result. \square

Proposition 3. Given Variable Conservativeness and Corollary 2, if the threshold for p is t , then the threshold for $\neg p$ is greater than $1 - t$. Thus it never maximizes expected accuracy to believe both p and $\neg p$.

Proof. Recall that Variable Conservativeness states that **(a)** $|T_p| < |F_{\neg p}|$ and **(b)** $|T_{\neg p}| < |F_p|$. By Corollary 2, the threshold t_p at which it maximizes expected accuracy to believe p is $\frac{-F_p}{T_p - F_p}$, and the threshold $t_{\bar{p}}$ at which it maximizes expected accuracy to

believe $\neg p$ is $\frac{-F_{\neg p}}{T_{\neg p} - F_{\neg p}}$. We want to show that $t_{\bar{p}} > 1 - t_p$. This holds iff

$$\begin{aligned}
& \frac{-F_{\neg p}}{T_{\neg p} - F_{\neg p}} > 1 - \frac{-F_p}{T_p - F_p} \\
\Leftrightarrow & \frac{-F_{\neg p}T_p + F_{\neg p}F_p}{T_{\neg p} - F_{\neg p}} > T_p \\
\Leftrightarrow & -F_{\neg p}T_p + F_{\neg p}F_p > T_pT_{\neg p} - F_{\neg p}T_p \\
\Leftrightarrow & F_{\neg p}F_p > T_pT_{\neg p}
\end{aligned} \tag{1}$$

Since $F_{\neg p}$ and F_p are both negative, and T_p $T_{\neg p}$ are both positive, (1) holds iff

$$|F_{\neg p}F_p| > |T_pT_{\neg p}| \tag{2}$$

By (a) above we have $|F_{\neg p}| > |T_p|$, and by (b) we have $|F_p| > |T_{\neg p}|$, thus (2) holds, as desired. That is, if the threshold for p is t_p , then the threshold $t_{\bar{p}}$ for $\neg p$ is greater than $1 - t_p$. *Corollary:* Since $C(\neg p) = 1 - C(p)$, if $C(p)$ is above the threshold t_p to suffice for belief, then $C(\neg p)$ is *below* the threshold $t_{\bar{p}}$ to suffice for belief, and vice versa. Therefore it never maximizes expected accuracy to believe both p and $\neg p$. \square

Proposition 4. If A_{B_i} is belief- and proposition-dependent and B_1, \dots, B_n are formed on the sequence p_1, \dots, p_n , then B_n maximizes expected accuracy with respect to B_{n-1} only if:

- (i) If $p_n \in B_n$, then $C(p_n) \geq \frac{-F_{p_n}^{B_{n-1}}}{T_{p_n}^{B_{n-1}} - F_{p_n}^{B_{n-1}}}$
- (ii) If $p_n \notin B_n$, then $C(p_n) \leq \frac{-F_{p_n}^{B_{n-1}}}{T_{p_n}^{B_{n-1}} - F_{p_n}^{B_{n-1}}}$

Proof. For convenience relabel B_{n-1} as B^- , and suppose B_n maximizes expected A_{B^-} . Relabel p_n as p , and first suppose (i) $p \in B_n$. We have:

$$\begin{aligned}
EA_{B^-}(B_n, \mathcal{P}) & \geq EA_{B^-}(B^-, \mathcal{P}) \\
\Leftrightarrow \sum_w [C(w)A_{B^-}(B_n, \mathcal{P}, w)] & \geq \sum_w [C(w)A_{B^-}(B^-, \mathcal{P}, w)]
\end{aligned}$$

Recall that at a p -world w_p , $A_{B^-}(B_n, \mathcal{P}, w_p) = A_{B^-}(B^-, \mathcal{P}, w_p) + T_p^{B^-}$, and similarly at a $\neg p$ -world $w_{\bar{p}}$: $A_{B^-}(B_n, \mathcal{P}, w_{\bar{p}}) = A_{B^-}(B^-, \mathcal{P}, w_{\bar{p}}) + F_p^{B^-}$. Thus we have the ugliest

equations in this paper:

$$\begin{aligned}
& \Leftrightarrow \sum_{w_p} C(w_p)[A_{B^-}(B^-, \mathcal{P}, w_p) + T_p^{B^-}] + \sum_{w_{\bar{p}}} C(w_{\bar{p}})[A_{B^-}(B^-, \mathcal{P}, w_{\bar{p}}) + F_p^{B^-}] \\
& \geq \sum_w (C(w)A_{B^-}(B^-, \mathcal{P}, w)) \\
& \Leftrightarrow \sum_w (C(w)A_{B^-}(B^-, \mathcal{P}, w)) + (C(p)T_p^{B^-} + C(\neg p)F_p^{B^-}) \\
& \geq \sum_w (C(w)A_{B^-}(B^-, \mathcal{P}, w)) \\
& \Leftrightarrow C(p)T_p^{B^-} + C(\neg p)F_p^{B^-} \geq 0 \\
& \Leftrightarrow C(p) \geq \frac{-F_p^{B^-}}{T_p^{B^-} - F_p^{B^-}}
\end{aligned}$$

as desired. If we suppose it (ii) $p \notin B_n$, then by parallel reasoning (via parallel ugliness)

we'll arrive at $C(p) \leq \frac{-F_p^{B^-}}{T_p^{B^-} - F_p^{B^-}}$. \square

Proposition 5. Given Sequence-Irrelevance, even if A_{B_i} is belief- and proposition-dependent, there are belief-independent values T_p and F_p for each proposition such that B is rational only if:

- (i) If $p \in B$, then $C(p) \geq \frac{-F_p}{T_p - F_p}$
- (ii) If $p \notin B$, then $C(p) \leq \frac{-F_p}{T_p - F_p}$

Proof. Take an arbitrary proposition p ; we find a T_p and F_p for which (i) and (ii) apply to all rational belief-states. Consider a sequence p_1, \dots, p_{n-1}, p that puts p the end and consider any $(n-1)$ -stage belief-state B^- . By Proposition 4 and Maximize Expected Accuracy, it is a rational constraint on the final belief-state B_n that both (iii) if $p \in B_n$ then $C(p) \geq \frac{-F_p^{B^-}}{T_p^{B^-} - F_p^{B^-}}$ and (iv) if $p \notin B_n$ then $C(p) \leq \frac{-F_p^{B^-}}{T_p^{B^-} - F_p^{B^-}}$. But then by Sequence-Irrelevance, this rational constraint on B_n must be a constraint on *all* rational belief-states: they all obey (iii) and (iv). That means if we set $T_p = T_p^{B^-}$ and $F_p = F_p^{B^-}$, the desired results (i) and (ii) hold. Since p was arbitrary, the same operation will work for every proposition. \square

References

- Berker, Selim, 2013. ‘The Rejection of Epistemic Consequentialism’. *Philosophical Issues*, 23(1):363–387.
- Briggs, Rachael R.A and Pettigrew, Richard, 2016. ‘An Accuracy-Dominance Argument for Conditionalization’. Manuscript.

- Buchak, Lara, 2014. 'Belief, Credence, and Norms'. *Philosophical Studies*, 169(2):285–311.
- Carr, Jennifer, 2015. 'Chancy Accuracy and Imprecise Credence'. *Philosophical Perspectives*, 29:67–81.
- , 2017. 'Accuracy or Coherence?' *Philosophy and Phenomenological Research*, To Appear.
- Clarke, Roger, 2013. 'Belief is Credence One (In Context)'. *Philosophers' Imprint*, 13(11).
- , 2017. 'Assertion, Belief, and Context'. *Synthese*.
- De Bona, Glauber and Staffel, Julia, 2016. 'Graded Incoherence for Accuracy-Firsters'. *Philosophy of Science*, To Appear.
- de Finetti, Bruno, 1974. *Theory of Probability*. John Wiley and Sons.
- Demey, Lorenz, 2013. 'Contemporary Epistemic Logic and the Lockean Thesis'. *Foundations of Science*, 18(4):599–610.
- Easwaran, Kenny, 2013. 'Expected Accuracy Supports Conditionalization – and Conglomerability and Reflection'. *Philosophy of Science*, 80(1):119–142.
- , 2016. 'Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities'. *Nous*, 50(4):816–853.
- Easwaran, Kenny and Fitelson, Branden, 2015. 'Accuracy, Coherence, and Evidence'. *Oxford Studies in Epistemology*, 5:61–96.
- Fantl, Jeremy and McGrath, Matthew, 2010. *Knowledge in an Uncertain World*. Oxford University Press.
- Foley, Richard, 2009. 'Beliefs, Degrees of Belief, and the Lockean Thesis'. In Franz Huber and Christoph Schmidt-Petri, eds., *Degrees of Belief*, 37–47. Springer.
- Frankish, Keith, 2009. 'Partial Belief and Flat-Out Belief'. In Franz Huber and Christoph Schmidt-Petri, eds., *Degrees of Belief*, 75–93. Springer.
- Friedman, Jane, 2013. 'Suspended judgment'. *Philosophical Studies*, 162(2):165–181.
- , 2017. 'Why Suspend Judging?' *Nous*, 51(2):302–326.
- Ganson, Dorit, 2008. 'Evidentialism and Pragmatic Constraints on Outright Belief'. *Philosophical Studies*, 139(3):441–458.
- Gibbard, Alan, 2008. 'Rational Credence and the Value of Truth'. In Tamar Szabó Gendler and John Hawthorne, eds., *Oxford Studies in Epistemology*, volume 2. Oxford University Press.

- Goodman, Jeremy and Salow, Bernhard, 2017. ‘Taking a Chance on KK’. *Philosophical Studies*, To Appear.
- Greaves, Hilary, 2014. ‘Epistemic Decision Theory’. *Mind*, 122(488):915–952.
- Greaves, Hilary and Wallace, David, 2006. ‘Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility’. *Mind*, 115(459):607–632.
- Greco, Daniel, 2017. ‘Cognitive Mobile Homes’. *Mind*, 126(501):93–121.
- Hare, Caspar, 2010. ‘Take the Sugar’. *Analysis*, 70(2):237–247.
- Hawthorne, John, Rothschild, Daniel, and Spectre, Levi, 2016. ‘Belief is Weak’. *Philosophical Studies*, 173(5):1393–1404.
- Hempel, Carl G, 1962. ‘Deductive-Nomological vs. Statistical Explanation’. *Minnesota Studies in the Philosophy of Science*, 3:98–169.
- Hintikka, Jaako, 1962. *Knowledge and Belief*. Cornell University Press.
- Holton, Richard, 2017. ‘Intention as a Model for Belief’. In Manuel Vargas and Gideon Yaffe, eds., *Rational and Social Agency: Essays on the Philosophy of Michael Bratman*. Oxford University Press.
- Horowitz, Sophie, 2017. ‘Accuracy and Educated Guesses’. In *Oxford Studies in Epistemology*. Oxford University Press.
- James, William, 1897. *The Will to Believe*. Longmans, Green, and Co.
- Joyce, James M, 1998. ‘A Nonpragmatic Vindication of Probabilism’. *Philosophy of Science*, 65(4):575–603.
- , 2009. ‘Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief’. In Franz Huber and Christoph Schmidt-Petri, eds., *Degrees of Belief*, 263–297. Springer.
- , 2010. ‘A Defense of Imprecise Credences in Inference and Decision Making’. *Philosophical Perspectives*, 24(1):282–323.
- , 2013. ‘Why Evidentialists Need Not Worry about the Accuracy Argument for Probabilism’. Manuscript.
- Kaplan, Mark, 1996. *Decision Theory as Philosophy*. Cambridge University Press.
- Leitgeb, Hannes, 2013. ‘Reducing Belief Simpliciter to Degrees of Belief’. *Annals of Pure and Applied Logic*, 164:1338–1389.
- , 2014. ‘The Stability Theory of Belief’. *Philosophical Review*, 123(2):131–171.

- Leitgeb, Hannes and Pettigrew, Richard, 2010. ‘An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy’. *Philosophy of Science*, 77(2):236–272.
- Levi, Isaac, 1967. *Gambling with Truth*. The MIT Press.
- Lewis, David, 1979. ‘Scorekeeping in a Language Game’. *Journal of Philosophical Logic*, 8(3):339–359.
- , 1982. ‘Logic for Equivocators’. *Noûs*, 16(3):431–441.
- Littlejohn, Clayton, 2015. ‘Who Cares What You Accurately Believe?’ *Philosophical Perspectives*, 29(1):217–248.
- Maher, Patrick, 1993. *Betting on Theories*. Cambridge University Press.
- Oddie, Graham, 1997. ‘Conditionalization, Cogency, and Cognitive Value’. *The British Journal for the Philosophy of Science*, 48(4):533–541.
- Pedersen, Arthur and Glymour, Clark, 2012. ‘What Language Dependence Problem? A Reply for Joyce to Fitelson on Joyce’. *Philosophy of Science*, 79(4):561–574.
- Pettigrew, Richard, 2012. ‘Accuracy Domination Arguments and Credences as Estimates of Truth-Values’. Manuscript.
- , 2013a. ‘A New Epistemic Utility Argument for the Principal Principle’. *Episteme*, 10(1):19–35.
- , 2013b. ‘Epistemic Utility and Norms for Credences’. *Philosophy Compass*, 8(10):897–908.
- , 2015. ‘Accuracy and the Credence-Belief Connection’. *Philosopher’s Imprint*, 15(16).
- , 2016a. *Accuracy and the Laws of Credence*. Oxford University Press.
- , 2016b. ‘Accuracy, Risk, and the Principle of Indifference’. *Philosophy and Phenomenological Research*, 92(1):35–59.
- , 2016c. ‘Jamesian Epistemology Formalized: An Explication of ‘The Will to Believe’’. *Episteme*, 13(3):253–268.
- Predd, J, Seringer, R, Lieb, E H, Osherson, D, Poor, H V, and Kulkarni, S, 2009. ‘Probabilistic coherence and proper scoring rules’. *IEEE Transactions on Information Theory*, 55(10):4786–4792.
- Rayo, Agustín, 2013. *The Construction of Logical Space*. Oxford University Press.
- Rosenkrantz, Roger D, 1981. *Foundations and Applications of Inductive Probability*. Ridgeview Press.

- Ross, Jacob and Schroeder, Mark, 2014. ‘Belief, Credence, and Pragmatic Encroachment’. *Philosophy and Phenomenological Research*, 88(2):259–288.
- Schoenfield, Miriam, 2016. ‘An Accuracy Based Approach to Higher Order Evidence’. *Philosophy and Phenomenological Research*, To Appear.
- , 2017. ‘Conditionalization Does Not (In General) Maximize Expected Accuracy’. *Mind*, To Appear.
- Shear, Ted and Fitelson, Branden, 2017. ‘Two Approaches to Belief Revision’. Manuscript.
- Smith, Martin, 2016. *Between Probability and Certainty: What Justifies Belief*. Oxford University Press Uk.
- Staffel, Julia, 2016. ‘Beliefs, Buses and Lotteries: Why Rational Belief Can’t Be Stably High Credence’. *Philosophical Studies*, 173:1721–1734.
- , 2017. ‘Accuracy for Believers’. *Episteme*, 14(1):39–48.
- Stalnaker, Robert, 1978. ‘Assertion’. *Syntax and Semantics*, 9:315–332.
- , 1984. *Inquiry*. Cambridge University Press.
- , 2006. ‘On the Logics of Knowledge and Belief’. *Philosophical Studies*, 128(1):169–199.
- Sturgeon, Scott, 2008. ‘Reason and the Grain of Belief’. *Noûs*, 42(1):139–165.
- Weatherson, Brian, 2005. ‘Can We Do Without Pragmatic Encroachment?’ *Philosophical Perspectives*, 19(1):417–443.
- Weisberg, Jonathan, 2017. ‘Belief in Psychontology’. *Philosopher’s Imprint*, To Appear.
- Williamson, Timothy, 2000. *Knowledge and its Limits*. Oxford University Press.
- Windschitl, P D and Wells, G L, 1998. ‘The alternative-outcomes effect’. *Journal of Personality and Social Psychology*, 75(6):1411–1423.
- Yalcin, Seth, 2010. ‘Probability Operators’. *Philosophy Compass*, 5(11):916–937.