



King's Research Portal

DOI: 10.1111/phpe.12064

Document Version Peer reviewed version

Link to publication record in King's Research Portal

Citation for published version (APA): Littlejohn, C. (2016). Who Cares What You Accurately Believe? *Philosophical Perspectives*, *29*(1), 217-248. Advance online publication. https://doi.org/10.1111/phpe.12064

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

•Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research. •You may not further distribute the material or use it for any profit-making activity or commercial gain •You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Who Cares What You Accurately Believe? Clayton Littlejohn King's College London clayton.littlejohn@kcl.ac.uk

0. Introduction

Some say that we should follow the evidence.¹ Some say that we should keep our stories straight. It can be hard to do both. Suppose you decided, never mind why, to memorize the phone book. As you commit each new entry to memory you might believe, quite reasonably, that each new entry is correct. Having just committed the last entry to memory, you discover the remnants of an errata slip. It tells you that an error was found in the book. Unfortunately the slip is torn and missing the piece that tells you the error's location. Because you have good evidence to believe there is an error in the book, that's what you believe. Because you have good evidence to believe each entry, you believe each entry. It's hard to see how the norms that govern belief could require you to believe everything the evidence supports and keep your story straight. If you follow one norm, you violate the other.

When conflicts between putative norms arise, we have to pick winners and losers. In trying to decide whether we should proportion our beliefs to the evidence or see to it that our beliefs are consistent we're liable to be knocked around by intuitions that will increasingly seem unreliable. One problem we face, then, in trying to state a theory of rational belief is that of finding a principled basis for choosing between putative norms.

A second problem we face in stating a theory of rational belief has to do with a specific kind of rational requirement. Some of the clearest examples of rational requirements are wide-scope requirements, requirements that require us

¹ I would like to thank Kristoffer Ahlstrom-Vij, Bill Brewer, Thomas Byrne, Jochen Briesen, Jennifer Carr, Charles Cote-Bouchard, Marian David, Jeff Dunn, Kenny Easwaran, Anna-Maria Asunta Eder, Claire Field, Sandy Goldberg, Jimmy Goodrich, John Hyman, Mark Eli Kalderon, Jason Konek, Eliot Michaelson, David Owens, Matthew Parrott, David Papineau, Sherri Roush, Florian Steinberger, Robyn Repko Waller, Ralph Wedgwood, Jake Wojtowicz, Jose Zalabardo and audiences at King's College London, the University of Konstanz, and the University of Leuven for discussing these issues with me. I should add that Julien Dutant, Branden Fitelson, John Hawthorne, and Richard Pettigrew each provided much more extensive feedback than anyone could reasonably ask for.

to have or avoid certain combinations of attitudes.² They might be requirements that tell us to keep our stories straight, keep our stories straight when our stories aren't very long, or see to it that we are probabilistically coherent. Although some of the clearest examples of rational requirements are wide-scope requirements, it isn't clear why there should be such requirements. These are requirements we'll conform to iff our attitudes exhibit the right patterns. Because the attitudes that exhibit these patterns might not be appropriate to the situation we're in, it isn't clear that we would always have good reason to have these attitudes. How, then, could rationality always require us to maintain these combinations of attitudes?

Our second problem is the problem of pretty patterns.³ If the ugliness of inconsistency or probabilistic incoherence doesn't explain the irrationality of small sets of inconsistent beliefs or the irrationality of probabilistic incoherence, what could? When assessing attitudes that are supposed to fit the facts, why should it matter if these attitudes fit together?

In this paper, I want to look at a consequentialist approach to epistemic norms because it seems to give us the resources needed to solve our two problems about rational belief. Starting from the assumption that accuracy or truth is an epistemic good that calls for promotion, we can use norms from decision-theory to vindicate probabilism.⁴ It seems we can explain why rational beliefs form pretty patterns without assigning any independent value to these beliefs' aesthetic properties. Moreover, we can use these same resources to

² While some authors might think that all such requirements are wide-scope, Brunero (2013) and Titelbaum (2015b) show us how to derive narrow-scope requirements from wide-scope requirements.

³ Wedgwood puts the problem this way when discussing views of practical rationality that posit wide-scope requirements, "According to the constructivists, the fact that one always 'should' make choices that meet ... internal conditions of coherence ... is not explained by the good external results to which such choices lead ... This is a rather surprising idea. Why on earth should such a thing matter purely for its own sake? Perhaps choices that do not meet these internal conditions of coherence ... are aesthetically unattractive in some way: they form a less pretty mental pattern than choices that do meet these conditions. But this hardly seems a sufficiently weighty consideration to explain why ... one 'should' never make choices that do not satisfy these internal conditions" (2003: 217).

⁴ See Joyce (1998, 2009, and MS) and Pettigrew (2013, 2016) for details.

explain why rational believers ought to tolerate certain kinds of inconsistency.⁵ If the accuracy-first framework really does help us get a handle on these problems, it would have an undeniable appeal.

Although the verdicts that our consequentialists try to vindicate certainly seem intuitively compelling and the apparent platitudes they argue from appear platitudinous, I have serious reservations about their approach to epistemic norms. In §1, I'll sketch the consequentialist arguments for probabilism and the normative Lockean view. In §2 I shall argue that the consequentialist approach to the norms for partial belief only look promising when consider some artificially constrained sets of options. When we lift certain restrictions, the implications of the view are far from intuitive. In §3 I shall argue that the value theory that consequentialist need to vindicate probabilism is unmotivated. While I won't argue that partial beliefs aren't potential bearers of epistemic goodness, I will argue that we don't yet have a good grip on what their good-making features might be. In §4 I shall discuss the consequentialist argument for the normative Lockean view. After presenting a series of objections I draw some general lessons about the epistemic good and the relations between full and partial belief in §5.

1. Veritism, Consequentialism, and Epistemic Norms

Let's start with a rough sketch of two consequentialist arguments for two epistemic norms. The first is an argument for probabilism, the view that says that rationality requires us to be probabilistically coherent.⁶ The second is an argument for a normative Lockean view, the view that it is rational to believe p iff it is rational to have a sufficiently high degree of confidence in p.⁷

The assumptions that figure in these arguments are pretty much what you'd expect. Because our accuracy-first epistemologists are consequentialists, they tell to evaluate the subject's options in terms of the values realized by the attitudes contained in those options and they'll then give us some norms that tell us what's rational to believe in light of these options and their ranking. The intended result is an argument for probabilism that rests on intuitively plausible claims about the value of accuracy.

⁵ See Easwaran and Fitelson (2015) and Fitelson (MS) for discussion of evidential norms and consistency. If interested in accuracy-first defenses of other norms, see Greaves' and Wallace's (2006) discussion of conditionalization.

⁶ For discussion, see Joyce (1998, 2009, MS) and Pettigrew (2013, 2016).

⁷ For discussion, see Dorst (MS), Easwaran (Forthcoming), and Fitelson (Forthcoming).

Let's start with the value theory. Our epistemic consequentialists are veritists. They think that there is one and only one fundamental epistemic good: accurate belief.⁸ The one interesting complication is that our consequentialists think there might be two kinds of accuracy and two kinds of epistemic good:

The relevant success criterion for full beliefs is wellknown and uncontroversial ... An epistemically rational agent must strive to hold a system of full beliefs that strikes the best attainable overall balance between the epistemic good or fully believing truths an the epistemic evil of fully believing falsehoods ... My position is that a rational partial believer must aim not simply to accept truths and reject falsehoods, but to hold partial beliefs that are gradationally accurate by adjusting the strengths of her opinions in a way that best maximizes her degree of confidence in truths while minimizing her degree of confidence in falsehoods. For the same reasons that a person should aim to hold full beliefs that are categorically accurate, so too should she aim to hold partial beliefs that are gradationally accurate (Joyce 1998: 579).

It is easy to see what accuracy for full belief amounts to. If you believe p outright, this state is categorically accurate iff p. To see how well you're doing epistemically we assign a positive value to your true beliefs, R, a negative value to your false beliefs, W, and the aggregate epistemic good realized by your full beliefs will be the sum. What if you don't believe p outright but have some partial belief concerning p? In this case we focus on the gradational accuracy of your belief:

The categorical good of fully believing truths is replaced by the gradational good of investing high credence in truths (the higher the better); the categorical evil of fully believing falsehoods is replaced by the gradational evil of investing high credence in falsehoods (the higher the worse) (Joyce MS).

To see how this works, consider a simple case. Suppose Agnes has opinions about p and $\sim p$. Her credence function, b, assigns .6 to p and .4 to $\sim p$. The ideal credence function, i, assigns 1 to truths and 0 to falsehoods in a world. We can see how well Agnes is doing in some world by measuring the distance between

⁸ For defenses of veritism, see Goldman (1999) and (possibly) Lynch (2004).

her credences and the ideal credences. The Brier score is the standard way to evaluate the accuracy of a subject's credences.⁹ For each of Agnes' opinions we subtract the value *i* assigns from the value *b* assigns and square it. We then sum these values to get a value that represents the total gradational inaccuracy of Agnes' credences. When comparing two scores, a lower score represents less distance from the ideal and less total epistemic disutility. As with the case of full belief, the aggregate good realized by a set of partial beliefs is simply the sum of the values associated with each partial belief. The main difference is that the scoring rules for partial belief tell us that a lower score is preferable to a higher score as higher scores represent a greater distance between your set of partial beliefs and the ideal.¹⁰

With the basic value theory before us, we can turn to the theory of rationality. In slogan form, epistemic consequentialists say that rational belief is good enough belief. If they are maximizers, they'll say that we're required to believe the best we can.¹¹ If we want to know whether having a full belief or partial belief concerning p is rationally permitted, we think about the epistemic value realized in feasible options where the subject has or lacks the full or partial belief concerning p. Assuming that our consequentialists are maximizers, rationality permits believing p if the options involving the belief in p does as well in terms of promoting epistemic value as the options that lack this belief. Belief is required if the best options all include it. It's forbidden if the best options lack it. Since the fundamental good is accuracy, the options should be ranked from best to worst in terms of the accuracy and inaccuracy of the attitudes contained in these options.

When it comes to partial belief Agnes should aim to minimize total gradational inaccuracy. She should minimize the distance between her credences and the ideal. Bearing this in mind, here is a quick sketch of an argument for probabilism. Probabilism tells us that Agnes would be irrational if her credences in p and $\sim p$ did not sum to 1. It turns out that for every probabilistically incoherent credence function defined over a fixed set of propositions (e.g., p and

⁹ See Joyce (1998) and Leitgeb and Pettigrew (2010) for discussion of the merits of using the Brier score to measure a credence function's inaccuracy.

¹⁰ This difference will matter later. It isn't essential to the framework because, as we'll see, the basic framework can be modified so that the scoring rules for partial belief more closely resemble those for full belief, but it's going to be useful to start with the simple scoring rules to understand the difficulties discussed in the next section.

¹¹ Subject to the proviso that 'best' might be understood as subjectively best, prospectively best, objectively best, etc.

 $\sim p$) there is an alternative coherence credence function defined over the same set of propositions that would be less inaccurate in every possible world. Thus, if we're thinking of Agnes' options as the credences she could have in the propositions in this fixed set, we would think that the rationality of her credences would be determined by the gradational inaccuracy of the credences in these propositions. An option involving incoherent credences should rank lower than an option containing a coherent set. Moreover, no coherent credence function defined over a set of propositions would be weakly dominated by an alternative function defined over that same set of propositions.¹² Thus, using the Brier score to measure the total gradational inaccuracy of Agnes' credence function and some norms concerning dominance, we get the result that rationality requires her to be probabilistically coherent. Agnes isn't special in this regard. If rationality requires it from her, it requires it from us, too.

If you think that this is a plausible rationale for probabilism, it might seem that we've solved the pretty patterns problem. Any solution to the pretty patterns problem should satisfy two desiderata. First, in stating the solution we cannot assume that the patterns themselves are desirable. Because we don't want to fetishize patterns, we should focus on the total value realized by a subject's options and use our ranking of options to vindicate the idea that some putative pattern is rationally required. Second, the patterns should be pretty. The accuracy-first approach seems to satisfy both desiderata. Being probabilistically coherent is a pretty pattern and the argument just sketched seems to show that if your attitudes don't exhibit the right pattern there is an alternative set of nondominated attitudes that is better than your present set. Thus, if you care about accuracy, you should care about the pattern for reasons that have nothing to do with its aesthetic properties.

This should give the reader some sense of how the accuracy-firsters would go about justifying putative norms for credences, but what about norms for full belief? Consider an accuracy-first argument for a normative Lockean view, a view that says that it is rational to have a full belief in p iff it is rational to have a sufficiently high credence in p. Let's suppose that it is rational to believe p iff neither disbelieving p nor suspending on p will have a greater expected epistemic value than believing p. The relevant values are the positive values that attach to categorically accurate beliefs and the negative values that attach to categorically inaccurate beliefs. Let's say that each false belief realizes the same disvalue, W, and each true belief realizes the same value, R. Assuming that the subject's credences are probabilistically coherent, we can say that it would be rational for

¹² The formal details are discussed in Joyce (1998). The argument was inspired by work by di Finetti (1974).

the subject to believe p iff her credence in p is greater than or equal to W/R+W.¹³ If a subject fails to believe p when her credence in p is greater than this value, she fails to maximize expected epistemic accuracy. As the veritist sees things, this is a failure to maximize expected epistemic value.

This presentation has been brief. It doesn't cover the technical details. My aim here is to contest the significance of the formal results for debates about epistemic norms, not the results themselves.

2. Pretty Patterns and Partial Beliefs

Let's consider a view that combines gradational veritism with epistemic consequentialism (GVEC). Because GVEC is a consequentialist view, its proponents offer a theory of rational partial belief by arguing that partial beliefs by identifying some value these beliefs can realize, telling us how to rank options in terms of the realization of this value, and then by telling us how this ranking bears on the rationality of partial belief. Roughly, rational partial beliefs will be rational because of how they promote a good that calls for promotion. Because

¹³ See Dorst (MS) and Easwaran (Forthcoming: 13) for discussion of arguments for a normative Lockean view. There is an issue that we cannot explore here that probably deserves further consideration. In Littlejohn (2012), I explained that if we want a view on which we need to have stronger evidence for *p* than for ~*p* to rationally believe p, we need to assign greater weight to the disvalue associated with a false belief than we do to the value associated with a true belief. Should we take a similar line with the values associated with gradational accuracy and inaccuracy? That is, should we say that the magnitude of disvalue associated with a degree of gradational inaccuracy is greater than the magnitude of value associated with a comparable degree of gradational accuracy? If we do and we assign the same values or disvalues to each degree of gradational accuracy or inaccuracy, it will be difficult to see how it could be rational to have a credence of .5 in *p* when the evidence supports *p* and $\sim p$ equally. That state would include equal degrees of gradational accuracy and inaccuracy and would be an overall bad state to be in. We either need to assign the same magnitude of value and disvalue to equal degrees of gradational accuracy and inaccuracy so that the total value realized by a credence of .5 is 0 or we need to say that the same degrees of gradational inaccuracy get different values depending upon how much total inaccuracy a state involves. These complications are required to get the desirable result that having a credence of .5 in *p* is a state that has neutral value, but they seem to introduce a degree of unnaturalness into the underlying value theory. It forces us to try to explain why gradational accuracy has these features that categorical accuracy does not.

the view incorporates gradational veritism, it ranks options in terms of the gradational accuracy associated with a subject's partial beliefs. Its proponents see gradationally accurate and inaccurate partial beliefs as playing a theoretical role similar that of hedons and dolors in hedonistic utilitarianism.¹⁴

Remember that the crucial formal results for the argument for probabilism were these. First, when comparing credence functions defined over a fixed set of propositions a subject will be dominated if her credence function is not a probability function. Second, when we're comparing credence functions defined over a fixed set of propositions a subject will avoid being dominated if her credence function is a probability function. These results might seem to give the consequentialists what they need to vindicate probabilism because it seems that consequentialists would agree that it is irrational to opt for an option that's dominated and rational to opt for an option that's not dominated. Unfortunately, things aren't quite this simple. It isn't clear that the formal result establishes much of anything that a consequentialist would or should care about.

These results tell us nothing about credence functions defined over different sets of propositions.¹⁵ It's not clear that this set of options these scoring rules score is the set of options that consequentialists would care about. Maximizing consequentialists want to know whether an act or attitude is included in the best *feasible* options. Not only is it unclear whether each possible set of credences concerning a fixed set of propositions should count as a feasible option, it isn't clear why the feasible options should be limited to options in which the subject has credences concerning a single fixed set of propositions.¹⁶

¹⁴ See Pettigrew (2016: 10), for example. Much in the way that the value associated with hedons and dolors are thought by utilitarians to give moral norms their 'normative force', some writers (e.g., Talbot (2014)) suggest the value associated with accurate belief is supposed to serve as the source of normative force that addresses the kind of skepticism about the normative force of epistemic norms expressed by writers like Papineau (2013).

¹⁵ See Pettigrew (2016: 4).

¹⁶ A subject might know that she couldn't have certain credences owing to contingent facts about the causal structure of the world. If, for some reason, she knew that it would be nomologically impossible to have a probabilistically coherent set of credences, shouldn't a consequentialist say that rationality requires her to adopt the set that minimizes expected gradational inaccuracy and would be rational if she had these credences? If so, GVEC does *not* show that rationality invariably requires probabilistic coherence, only that it does so when the contingent causal structure of the world allows a subject to have such a set of credences. If a consequentialist in ethics purported to show that an equal

Remember the guiding thought in the argument for probabilism was this. Once we agree that the fundamental values are gradational accuracy and inaccuracy, we should agree that it would be irrational for a subject to have a set of credences if she can see that there is an alternative that involves less gradational inaccuracy or an alternative that is closer to the ideal. Recall Joyce's suggestion that rationality requires us to strive to hold a set of partial beliefs that have high degrees of gradational accuracy. Our scoring rules don't support this suggestion. They don't tell us that Agnes would be worse off for failing to have partial beliefs like the ones Joyce describes. What the scoring rules show is that for any set of credences over some fixed set of propositions Agnes could have it would be better for those opinions have low amounts of gradational inaccuracy. We can't use the scoring rules to identify any reason for Agnes to get in the partial belief game in the first place.¹⁷

If we want to show that Agnes really should aspire to have partial beliefs that have certain properties (e.g., they are coherent and they minimize gradational inaccuracy), we need to think of Agnes' available options as involving suspension and opinionation and we need a value theory that tells that Agnes could be better off for being opinionated.

If *b* is Agnes' credence function, she has a credence of .6 and .4 in *p* and $\sim p$ respectively. If Agnes' aim is to avoid being dominated by an alternative set of credences scored for their total gradational inaccuracy, she should avoid an incoherent credence function like *b*' that assigns .7 and .4 to *p* and $\sim p$ respectively. What about a function like *b*'' that assigns .6 to *p* and is undefined for $\sim p$? If Agnes' aim is to minimize gradational inaccuracy, we know *apriori* that *b*'' will contain less by way of gradational inaccuracy than *b*. As such, *b*'' would be closer to the ideal credence function in every possible world.¹⁸ If *b*'' is Agnes' credence

distribution of goods maximized utility but then conceded that their argument assumed for no particularly good reason that feasible options involving great amounts of total utility and an unequal distribution of these goods could be ignored for the purposes of ranking an agent's options, we wouldn't be impressed by this kind of consequentialist vindication of egalitarianism.

¹⁷ I suppose someone might argue that the partial beliefs are unavoidable in some sense, but it still seems strange to think that we should lament the fact that we're saddled with middling credences. At any rate, I think it's a useful exercise to take seriously Gibbard's suggestion that we can fruitfully think about these issues if, 'We ask whether those are the credences that, in light of her evidence, she most prefers to have, the ones that she would choose if she could choose her credences at will' (2007: 150) and offer our rational believer the option of opting out.

¹⁸ The sum of two (non-zero) distances is always greater than one summand.

function, Agnes doesn't satisfy the demands imposed by probabilism. If Agnes' has the credences associated with b'', she is probabilistically incoherent. Thus, it seems there is a real problem with the consequentialist argument for probabilism. While Agnes won't be dominated by an incoherent alternative to b defined for p and $\sim p$, she knows *apriori* that there is an incoherent alternative to b that has less by way of gradational inaccuracy in every possible world. She can move closer to the ideal in every possible world by simply 'dropping' a non-extremal credence.¹⁹ If she should prefer less gradational inaccuracy to more, shouldn't she prefer b'' to b? Consequentialists have to think that it's irrational to prefer an acknowledged worse state of affairs to one that's better, so wouldn't it be irrational for Agnes to 'choose' b over b''?

GVEC faces the problem of partial opinionation. Suppose the proponents of GVEC want to allow that there can be rationally acceptable credences that include some non-extremal credences. Among other things, they want to say that this set is rational only if it is probabilistically coherent and it isn't dominated by an alternative. Let's suppose that the set of credences associated with b is one such collection. When a gradational veritist proposes that this is a set of rationally acceptable credences, we can find an alternative set that differs from *b* only insofar as it doesn't include any opinion about $\sim p$ (e.g., b''). This alternative is probabilistically incoherent but it involves less gradational inaccuracy in every possible world than b. The problem isn't just with b, mind you. We can do this with any coherent set of credences that involves non-extremal credences in a proposition and its negation. Thus, we get the troubling result that *if* there are any rationally acceptable non-extremal credences, there will be a rationally acceptable set of credences that is probabilistically incoherent that we know apriori involves less gradational inaccuracy in every possible world. Thus, if the more expansive set of coherent credences were rationally permitted, the incoherent contracted set would be rationally permitted. (This follows from the consequentialist idea that if one option is rationally permitted and another ranks at least as well in terms of the ranking value this second one must be permitted, Moreover, if rationality requires us not to have credences that are too.) dominated in this way, no rationally acceptable set of credences involves nonextremal credence. (We can always move closer to the ideal by jettisoning another non-extremal credence.) This can't be a consequence that proponents of GVEC would be happy with.

¹⁹ I've learned from Robyn Repko Waller that this was essentially Curious George's insight about golf. If the golfer's aim is just to avoid a high score, it's best not to play.

Unless GVEC is modified, its proponents are saddled with two unattractive options. If they think that some rationally acceptable sets of credences include non-extremal credences, they'll have to say that some very ugly sets of credences are also rationally permitted. They'll lose their argument for probabilism. If, however, they think that ugly sets of credences are irrational, they get the troubling skeptical result that it is always irrational to have nonextremal credences. Maybe they can still get the result that a rational set of credences will always be probabilistically coherent, but this feels like an uninteresting result if all sets of credences involving non-extremal credences are deemed irrational.²⁰

In discussion, people have suggested two potential solutions. The first involves reformulating the norm. When it comes to evaluating certain sets of options over fixed sets of propositions, it might make sense to work under the assumption that the fundamental epistemic norm is one that enjoins us to minimize gradational accuracy, but when we expand the set of options to include suspension we need to introduce a different norm. The second focuses on the value theory. The gradational veritist might add to the initial value theory to explain why it's sometimes good to have partial beliefs.

2.1 Average or Total Consequentialism?

Here's a natural diagnosis of the problem of partial opinionation. The problem arises because we're assuming that options should be ranked in terms of distance from an ideal or in terms of the minimization of gradational inaccuracy. The problem only arises if we assume that a subject should be worried about *total* gradational accuracy, but the problem does not arise if the norm that matters is one that enjoins us to minimize *average* inaccuracy. Instead of scoring a subject's credences by summing the squared distances between the subject's credences and the ideal, we can score a subject's credences by summing the squared

²⁰ Of course, if they think that having extremal credence in a proposition amounts to having a full belief that isn't a partial belief, the upshot is that there are no epistemic norms that permit partial beliefs. One of the attractions of the present project was that it was supposed to help us see how it could be rational to have high degrees of credence in inconsistent sets of propositions. I thought that many people were troubled by Ryan's (1991) defense of the consistency norm for full belief, but it seems that we're led to embrace something much more radical if the skeptical line is taken, for now it seems that the reason that we shouldn't have inconsistent sets of beliefs in cases like our original preface is simply because we shouldn't have any beliefs (partial or full) except in the cases where extremal credence is rationally permitted.

distances between these credences and then dividing by the number of opinions. When we compare *b* and *b*^{''} we get the result we want. If Agnes had *b* as her credence function she would be coherent, she could potentially be rational, and her credences would have a lower average inaccuracy than they would if she had *b*^{''} as her credence function.²¹

This is a neat fix, but the resulting view is unsatisfactory. When we combine gradational veritism with average epistemic consequentialism (GVAEC), the resulting view won't vindicate an important intuition about the virtues of well-founded opinion. The people I love tell me that I should have more opinions. I don't have views about things like life on other planets, Jeremy Corbyn, or the housing market. I've been told that there's something attractive about David's mix of curiosity and confidence. His evidence always supports his attitudes and he seems to have opinions about everything (e.g., politics, the morality of cheating in sport, the number of colors, and the best places to fish in Essex). If you compare our opinions in terms of average inaccuracy our attitudes do equally well, but wouldn't David be doing better if he's more opinionated than timid thinkers like me?

GVAEC trades off one problem for another. If it solves the problem of partial opinionation, it won't solve the problem of population (i.e., the problem of explaining why we should strive to have more well-founded partial beliefs).²² If we want our accuracy-first view to look anything like hedonism, it better tell us that, other things equal, it's better to have more partial beliefs with high degrees of gradational accuracy. It should treat such these non-ideal credences as hedons, not dolors. Unfortunately, GVAEC doesn't deliver that result. If the aim is to minimize *average* gradational inaccuracy, I couldn't have a reason to take on new opinions like David does. By hypothesis, we do equally well in terms of average gradational inaccuracy.

GVAEC also generates a skeptical worry. If the fundamental norm enjoins us to miminimze average gradational inaccuracy, we face considerable rational pressure to avoid non-extremal credences. If we followed the narrator's advice from the *Meditations* by believing only what was indubitable and having only extremal credences in certainties, we would be do better minimizing average gradational inaccuracy than we would if we had non-extremal credences concerning contingent matters of fact.

²¹ Eleanor Knox first suggested this response to me in conversation. Pettigrew (MS) discusses the merits of this view in connection to a different problem I will discuss momentarily.

²² Richard Pettigrew pointed out that the population problem was first discussed in Carr (2015). See his (MS) for his response.

Finally, GVAEC doesn't solve the problem of pretty patterns. It might seem to vindicate probabilism because we fail to minimize average gradational inaccuracy if we are probabilistically incoherent, but it doesn't satisfy our two desiderata. We don't want to fetishize patterns, so we want to rank options in terms of their total intrinsic value. GVAEC doesn't do that. It ranks options in terms of the average inaccuracy. It seems fetishistic precisely because it ranks options without an eye towards total value.

2.2 Carrots and Sticks

Nothing good could come from moving from GVEC to GVAEC so let's focus on the former view. We won't get a satisfactory response to the problem of partial opinionation unless we attend to the details of gradational veritism and enrich the scoring rules.

Let's consider a second diagnosis of the problem of partial opinionation. The arguments for probabilism assume that sets of coherent but non-extremal credences are always worse than the ideal. GVEC then characterizes the rational believer's aim as that of minimizing inaccuracy. Thus, there's always an incentive to 'drop' non-extremal credences and minimize your number of partial beliefs. Whenever you drop some non-zero distance from the ideal, you 'move' closer to the ideal and whenever you take on some partial belief you move some non-zero distance from the ideal.

Since the arguments for probabilism were designed only to show something about sets of credence functions defined for fixed sets of propositions, nothing really turned on whether it would be good to have the credences associated with these credence functions. Once we allow for the possibility of suspension, however, it matters whether there's any potential upside to having partial beliefs. It seems natural to think that the value realized by suspension is 0. What the gradational veritists have to do is show that Less than Zero is false and give us an account of gradational accuracy that tells us that there's some reason for Agnes to get into the partial belief game:

Less than Zero: The total value realized by the ideal set of credences in a world is 0 or greater than 0 but the value associated with some non-extremal credence is always less than 0.

Until they show that partial beliefs are potentially good states to be in, we're stuck choosing between ugly patterns and skeptical pressures.

Let's consider making two changes to the gradational veritist's value theory. First, suppose that some non-extremal credences are good states to be in because they have more gradational accuracy than some valueless state. Second, suppose there is a neutral state, a credence that has no epistemic value whatever and was thus of comparable value to suspension. The natural suggestion is a credence of .5.²³ We could then say that any credence with a higher degree of gradational accuracy than a credence of .5 would be a good state to be in and that any credence with a higher degree of gradational inaccuracy would be a bad state to be in. As your credence in a truth increases, things get better. As your credence in a falsehood increases, things get worse. We now have our carrots and sticks. We can now see why it's good to be like David. Each partial belief he has that has more gradational accuracy than the neutral credence would be a good state to be in. The more states like this, the better. Once we introduce a potential upside to partial belief, we can explain why partial opinionation is bad. Initially, it looked as if GVEC was committed to the unfortunate result that it would be better to have b'' as your credence function because the credences associated with b. Now it

²³ Note that 'neutral' can be understood in two ways. First, we might say that it is neutral in the sense that it contains equal measures of accuracy and inaccuracy. Second, we might say that it is neutral in the sense that it contains no value. It's this second reading that's operative in the discussion. I'll assume here for no good reason that they come to the same thing. When it comes to categorical accuracy, an epistemic state that contains equal amounts of categorical accuracy and inaccuracy is one that the epistemic consequentialist should say is overall worse than a state that involves suspending on the two propositions. This raises an interesting question that would need to be explored by gradational veritists if they wished to develop this view: should we say that equal degrees of categorical accuracy and inaccuracy balance out? If so, it seems that there is a surprising difference between the weights of the comparative values realized by categorical accuracy and inaccuracy and gradational accuracy and inaccuracy. If we had a view on which each degree of gradational accuracy and each degree of gradational inaccuracy realized the same amount of epistemic value or disvalue and we said that in terms of its weight the value associated with a degree of gradational inaccuracy exceeded the weight of the value associated with a comparable degree of gradational accuracy, we would get the result that a state like a credence of .5 that contained equal measures of gradational accuracy and inaccuracy would be a bad state to be in however the world happened to be. There are ways of avoiding this result (e.g., assigning variable amounts of value to comparable degrees of accuracy or inaccuracy or insisting that the value and disvalue associated with a credence of .5 washes out) but it seems rather ad hoc to posit features of a value theory to get the pleasing formal results in the absence of any independent justification to think that there are these surprising complications.

looks as if *b*" is potentially a bad state to be in. If *p* is true, having a credence less than .5 in $\sim p$ would be a good state to be in.²⁴

For all its attractions, this modification doesn't remove the difficulties for GVEC. We needed to introduce a neutral credence to make sense of the idea that some alternatives to suspension are potentially good or bad states to be in. This was supposed to mitigate the rational pressure against opinionation. Once we introduce neutral credences, we lose our argument for probabilism. Consequentialists think that if we know *apriori* that two options realize the same objective value it's impossible for one state to be permitted and the other forbidden. Suppose Agnes has credences of .5 in *p* and in ~*p*. If we compare her credences to a contracted set where she has a credence of .5 in *p* and no credence concerning ~*p*, we see that these credences have the same objective value in every possible world. It would be irrational for Agnes to have a credence of .5 in *p* and no credence in ~*p* but the modified view doesn't deliver this result. The problem of partial opinionation has been moved, not removed.

Because the current formulation of GVEC assigns positive value to a wide range of credences, it's the only view thus far that resembles hedonistic utilitarianism. In some ways, this might seem to be an advantage in dealing with the population problem and the problem of partial opinionation. Unfortunately, this new value theory generates a skeptical worry of its own. Suppose Agnes decided to study American history. We might imagine that things go quite well for her. She reads the literature and follows the arguments where they lead. The result is a set of opinions with high degrees of gradational accuracy. This should be enough to ensure that her opinions are rational, but nothing I've said tells us anything about where this set of opinions ranks amongst her alternatives. Had she used her mental energy and resources differently, Agnes could have learned more by dedicating herself to settling questions that would be easier to settle. She could have done mental math and traced out the obvious logical consequences of obvious things, for example. Had she done that instead, she might have had a greater number of opinions about these matters and no opinions about American history.

This fact about comparative value matters to consequentialists so it should matter to epistemic consequentialists.²⁵ We have to remember that gradational

²⁴ Julien Dutant and Matthew Parrott first suggested this response in conversation, but neither endorsed it. Richard Pettigrew also suggested a response along these lines and coined the phrase 'neutral credence'.

²⁵ In Littlejohn (2012: 47), I raise this as a problem for views that tell us to maximize the epistemic value realized by full beliefs. In Littlejohn (forthcoming b), I discuss the relationship between this objection to certain forms of epistemic

accuracy, like death, can be bad because it deprives us of some greater good. However accurate Agnes' current attitudes are, they can only be rational if they compare favorably to alternatives ranked in terms of total gradational accuracy. In thinking about alternatives as consequentialists should, we are not thinking of alternatives as alternative opinions over fixed sets of propositions but alternative feasible options, accessible possibilities in which Agnes has some opinions and forgoes others for uninteresting reasons like limits to computational capacity, the need to attend to some matters and ignore others, and the time it takes to settle the questions that concern her. To my mind, the fact that Agnes could have achieved greater amounts of total gradational accuracy by choosing to do mental math instead of investing her time in the careful but demanding field of American history tells us nothing about the rational status of her opinions about the causes of the Civil War. This is an anti-consequentialist point. It indicates that the rational status of her opinions is determined independently from consideration of the total value contained in the alternative options.

The best formulation of GVEC isn't good enough. Once GVEC takes on the resources it needs to address the problems of partial opinionation, it resembles consequentialist views that say that no matter how good an option is it can fail to be good enough because of the opportunity cost associated with making it actual. I don't see how any further permutation will avoid some combination of the problems discussed here, so we should conclude that CVEC doesn't solve the pretty patterns problem or vindicate probabilism. Once we try to understand what's wrong with Agnes if she's partially opinionated or what's good about David who follows Joyce's advice in seeking partial beliefs with high degrees of gradational accuracy, we see that the view's implications aren't terribly attractive.

3. Two Kinds of Accuracy

When presented with purported counterexamples some consequentialists try to beat their critics with cleverness. After some tinkering and modifying, they'll try to show that their view now vindicates their critics' intuitions. Some prefer to outSmart their critics. They'll try to convince you that it's a virtue of their view that it has some seemingly troubling implication. If armed with a good value theory, this might seem like a good strategy for the epistemic consequentialist. If gradational accuracy really is a fundamental epistemic good and it really does call for promotion, maybe Agnes *is* being unreasonable for wasting her limited resources on history. Maybe we should all just stick to mental math.

consequentialism and Ross' (1930) criticism of the utilitarian conception of beneficence.

This kind of move is only tempting if we think that the underlying value theory is sound. Is it? Joyce suggests that we should aim to have gradationally accurate partial beliefs and categorically accurate full beliefs for the very same reason. This suggests that we might value these kinds of accuracy for the very same reason. Do we? I don't think so. I have to confess that I find much of the discussion of epistemic value opaque. Because these discussions often start from assumptions that I don't accept, I've tried to identify the strongest arguments for veritism. The best arguments I've found all come from Lynch's work on the value of truth. In this section, I'll argue that these considerations support parts of categorical veritism but don't support gradational veritism. Until further arguments are found, we should think of gradational veritism as an unmotivated motivator. If GVEC's implications are as bad as they seem, we shouldn't feel pressured into accepting them on the grounds that they're supported by a good theory of epistemic value.

In the vast literature on the Meno Problem the assumption that true belief is good is often taken for granted.²⁶ It is fair to ask, though, why we should think that accuracy is a promotional good (i.e., a good that calls for promotion).²⁷ To motivate the idea that truth matters, Lynch asks us to reflect on our aversion to life in the experience machine. In spite of how pleasant that life would be, something is missing:

> In preferring not to live in either the vat or the Russell world, I do not simply prefer that the world be a certain way. My preference involves my beliefs and their proper functioning, so to speak. For not only do I not want to live in a world where I am a brain in a vat, I also don't want to live in a world where I am not so deceived, but believe that I am. That is, if such and such is the case, I want to believe that it is, and if I believe that it is, I want it to be the case. We can put this by saying that I want my beliefs and reality to be a certain way—I want my beliefs to track reality, to

²⁶ For gentle introductions, see Hyman (2015) and Kvanvig (2003).

²⁷ Not every appropriate use of 'good' corresponds to a promotional good. Even if you think that the notion of good simpliciter is coherent and think of this kind of good and good-for are promotional goods, there is also the attributive use and normative use. These latter two uses do not pick out promotional goods, things that by their very nature call for promotion. See Baron (1995) and Zimmerman (2015) for helpful discussion.

"accord with how the world actually is"—which is to say I want them to be true (2004: 18).

In an earlier passage he rejects the suggestion that the value that attaches to accurate belief is just a matter of their instrumental value on the grounds that this doesn't capture what's disturbing about the thought that we might be deceived and have no inkling of it. Even if things go equally well for us if we're deceived or informed, there is something bad about living a life in which there's an undetectable gap between appearance and reality.

There are two points to take from this. The negative point is this. While it would be good to have beliefs that get you to San Jose if that's where you want to be, the value that interests us doesn't derive from the way in which certain beliefs happen to knock us down the paths that take us where we want to be.²⁸ The positive point is that the primary reason that accurate beliefs are good (when they are good) is that these beliefs track reality and so enable us to keep in touch with parts of reality that we're concerned with when we rely on our beliefs. Specifically, they track the parts of reality that consist of facts, the facts that we're curious about and the facts that we hope figure in reasoning. In other words, if you recognize that a subject's reasons for feeling things, doing things, and believing things consist of certain facts that the subject has in mind, the reason that accurate beliefs are good (when they are good) is that puts these reasons into our possession or enables us to be guided by them in feeling things, thinking things, or doing things.

Since a belief cannot track reality without being accurate, we can see why accuracy is a condition on good belief. If we care about being guided by reasons, it makes sense to care about whether our beliefs can serve as such guides. We can get to San Jose in many ways. We can follow a map or we can be kidnapped and driven there. We don't need accurate beliefs to get to San Jose, but we *do* need accurate beliefs to be guided by certain considerations. You cannot be relieved that you've made it to San Jose or happy that you didn't take a bus that dumps you off in Berkeley if you don't know that you're in San Jose and don't know that you're not in Berkeley.

This best explains why some accurate beliefs realize the fundamental epistemic good:

Reasons and Epistemic Goodness

²⁸ Hyman (2015) stresses this point in his discussion of the Meno problem.

²⁹ For a defense of this conception of reasons as facts, see Alvarez (2010), Hyman (2015), Littlejohn (2012), Unger (1975), and Williamson (2000).

Beliefs are valuable (when they are) because these beliefs track the part of reality that consist of facts and thereby enable us to feel things for reasons, believe things for reasons, and do things for reasons that consist of these facts. They are a source of potential motivating reasons (i.e., reasons that can be the subject's reason for feeling, believing, and doing things).

While this account provides a neat explanation as to why some accurate beliefs are good (i.e., the accurate beliefs that track reality), the account doesn't support gradational veritism. It doesn't cover partial belief. Unless some partial belief is, *inter alia*, a full belief, it cannot play the functional role that full beliefs are supposed to. Thus, the grounds for thinking of some accurate full beliefs as good give us no reason to think of partial beliefs as potentially desirable states to be in. Excessive focus on the links between belief and bodily movement has obscured this. While both a full belief and a partial belief can lead you to take an umbrella, you need a full belief to regret that it's raining, be happy that it's not hailing, or be surprised that it's not sunny outside.³⁰

To believe, feel, or do something for a reason, the fact you have in mind has to capture the light in which the relevant response struck you as appropriate or fitting. It's only possible for the relevant fact to capture the light in which you took the response to be appropriate when you fully believe the fact to be a fact. (You cannot be angry with your neighbor for the reason that they took your paper if you do not fully believe they took it but merely assign some middling credence to this hypothesis.) Once we see that full belief is required for feeling things for reasons that consist of facts, it's easy to see that similar points hold for action and belief. You can take an umbrella if you suspect that it will rain or fear that it will rain, but you cannot take the umbrella for the reason that it will rain if you don't have the full belief that it will. If all you care about is staying dry, you might think that little turns on whether you have a full belief or a partial one, but we can choose better cases. If you are punishing someone, you had better fully believe that they did it. If not, your reason for punishing them couldn't be that they committed the relevant misdeed.

Categorically accurate beliefs are good (when they are good) because they play a distinctive role in providing us with potential motivating reasons (i.e., the things that can be our reasons for V-ing). The ability to provide us with these reasons explains the value of the full beliefs that are good in the fundamental

³⁰ See Adler (2002), Buchak (2014), Gordon (1987), and Owens (2013) for discussion of the relationship between full belief and emotion or blame.

way. Their accuracy is *not* the fundamental good-making feature. Once we see that full belief plays a distinctive role in providing us with potential motivating reasons, a role that mere partial belief cannot play, we can see that gradational veritists cannot assume that gradational and categorical accuracy are desirable for the same reasons. Until we better understand why partial beliefs matter, we shouldn't assume that they realize some good simply by virtue of their gradational accuracy and that this accuracy calls for promotion.

4. Categorical Veritism and the Lockean View

The focus of our discussion now shifts to a consequentialist approach to norms for full belief. As we've seen, there is a plausible line of argument for the normative Lockean view, the view that says that it's rational to fully believe a proposition iff it's rational to have a sufficiently high degree of confidence in that proposition. If, as categorical veritists believe, the cardinal epistemic good is true belief and the cardinal epistemic evil is false belief, the Lockean view might seem very attractive. It would seem to be motivated by this seemingly compelling line of argument:

> ... a rational agent should be doing well by her own lights, in a particular way: roughly speaking, she should follow the epistemic rule that she rationally takes to be most truth-conducive. It would be irrational, the thought goes, to regard some epistemic rule as more truth-conducive than one's own, but not adopt it (Horowitz 2014: 43).

This captures the fundamental consequentialist rationale for recognizing an epistemic norm. Once we fix the right values for true and false belief, we can compare the expected categorical accuracy of each set and it seems that you have to prefer the ones that score better than your current set.

In these next two sections we'll see that it's not always rational to believe the things that Lockeans take to be rational. It's also not always desirable to believe in line with CVEC. The fundamental consequentialist rationale for recognizing an epistemic norm is unsound.

4.1 The Problem of Constraints

On the Lockean view, it wouldn't be rational to fail to believe p if you rationally have a sufficiently high degree of confidence in p. The view thus denies that there can be constraints that prohibit full belief if they do not, *inter alia*, tell us to keep our credence below some threshold. The problem of constraints is the problem posed by the existence of a special kind of normative constraint on full belief that tells us that it's irrational to have a full belief even when rationality

permits or requires a sufficiently high degree of confidence in the target proposition.

Let's start with lottery propositions. Obviously it is rational to have a high degree of confidence in lottery propositions, but what about full belief? You know that the drawing was held yesterday, but it's not rational to fully believe that your ticket lost if you are going on the statistical information alone.³¹ To get this verdict, the Lockean would need to either argue that it's not rational to have a high degree of confidence in a lottery proposition or assign weights to the value of truth and falsity such that it's irrational to take the risk associated with full belief in spite of the fact that it's rational to have a high degree of confidence. Neither strategy is the slightest bit promising. You wouldn't want to ignore the statistical information in proportioning your degree of belief to the evidence. You wouldn't want to assign too much weight to disvalue associated with false belief because it's rational to believe the things you read in the paper even though these beliefs have lower expected categorical accuracy than lottery beliefs. It looks as if the constraint against believing lottery propositions is a counterexample to the Lockean view.

Not everyone sees this counterexample as a counterexample. A standard response is to concede that we cannot know that lottery propositions but insist that we can rationally believe them. I disagree, but set that aside. Suppose we cannot know lottery propositions. We can use this concession against the Lockean. Consider the proposition expressed by, 'This ticket lost but I don't know if it did'. If we cannot know that our ticket is a loser, this is very likely to be true. It is also a Moorean absurdity. It isn't just bad to say, it would be irrational to believe.³²

The Lockean tries to dodge the first counterexample by saying that it's rational to have a high degree of confidence in lottery propositions *and* to believe them outright. If it is rational to have this high degree of confidence in lottery propositions, it better be rational to have a high degree of confidence in this Moorean absurdity. This conjunction is the conjunction of the lottery proposition and an apriori truth about the subject's epistemic position. If it is rational to have a sufficiently high degree of confidence in the first conjunct, the same holds true for the conjunction. (There is no reason to think that if the support for a lottery proposition is sufficiently strong to make it rational to fully believe the proposition that the conjunction that includes the lottery proposition and the

³¹ On this point, I agree with Bird (2007) and Nelkin (2000), though I prefer Bird's explanation.

³² For discussion of Moorean absurdity and the rationality of belief, see Adler (2002) and Sorensen (1988). For dissent, see McGlynn (2013).

negative appraisal will have a sufficiently lower probability to fall below some line so that it's not sufficiently probable on the evidence.) Thus, the simple Lockean view implies mistakenly that it would be rational to believe that your ticket lost and that you don't know that it did.³³

If it is irrational to believe both p and that p is not known, we have another counterexample to the Lockean view. There should be constraints that rule out believing a proposition whilst believing that it is irrational to believe that proposition and constraints that rule out believing conjunctions where one conjunct represents the entire conjunction as irrational.³⁴ In terms of expected accuracy, these negative epistemic self-appraisals all do very well. If it's rational to have high credence in the first, the Lockean should regard each of the following as rational:

This ticket lost.

This ticket lost but I don't know if it did.

This ticket lost but I don't know if it did. It's irrational to believe that.

I don't know that this losing ticket is a loser and it's irrational for me to believe this.

Someone who didn't believe these things should see that conforming to constraints that prohibited these beliefs results in her doing worse given the aim of maximizing expected categorical accuracy than she would if she believed these propositions. A more permissive set of epistemic rules would do better from the perspective of truth-conducivity, but it hardly seems irrational to stick to rules that prohibit these negative epistemic self-appraisals. Thus, we should be suspicious of the fundamental consequentialist rationale for recognizing epistemic norms. Someone who sees that she's failing to maximize expected accuracy should sometimes be unmoved by that fact. The consequentialist rationale rests on some mistake about the good, some mistake about the relationship between the good and the rational, or both.

4.2 Foot Stomping

³³ Adler (2002) goes so far as to say that we *cannot* believe such things, but I submit that it would be irrational to believe them if we could *and* it would be irrational to both believe p and to separately believe that p is the kind of thing we cannot know when one fails to see the connection between the first-order belief and the negative assessment.

³⁴ For discussions of epistemic akrasia and defenses of this constraint, see Littlejohn (forthcoming), Titelbaum (2015), and Smithies (2012).

Sets of full belief that involve the negative epistemic self-appraisals do better in terms of expected categorical accuracy than those that lacked them, but it's hard to see how, from the perspective of rationality, it's desirable to have these categorically accurate full beliefs. It's not desirable from the perspective of rationality to believe things you know to be irrational even if this representation of your irrationality is accurate. As Foot (1985) observed, there is something to the consequentialist thought that the right or rational has to do with what's better or best. It does seem irrational to prefer an acknowledged lesser good to a greater good, but if the objections presented in §4.2.1 and §4.2.2 have the force I take them to, they show that the states of affairs that CVEC takes to be best or most desirable needn't be particularly good. Once we see this, we can see what's wrong with the consequentialist rationale sketched above. If we don't see anything particularly good in the states of affairs that proponents of CVEC insist are the best, we aren't guilty of preferring some lesser good to a greater one. We prefer the good and we can see that the good isn't simply a function of accuracy. The 'Foot stomping' objections are intended to show that the relationships between epistemic goodness and categorical accuracy are more complicated than CVEC suggests.35

4.2.1 When is Categorical Accuracy Beneficial? When is Categorical Inaccuracy Harmful?

According to CVEC categorical accuracy is always good to some extent. If a belief is true, it is *pro tanto* objectively good because of its categorical accuracy. If a set of beliefs is true, it would be *pro tanto* objectively good because of the categorical accuracy of its members. If a belief is sufficiently probable on the evidence, it would be a good candidate for being a rational belief because of its prospective goodness or expected categorical accuracy. Remember, though, that our categorical veritists are also consequentialists, so they'll allow that objective and prospective goodness are not the sole determinants of rational status. In this section we'll see that proponents of CVEC make three mistakes about the value of categorical accuracy. Their view doesn't allow for worthless accuracy, doesn't allow for the right kind of harmful accuracy, and falsely predicts that certain kinds of desirable beliefs are irrational.

Let's start with a case of worthless accuracy:

³⁵ I'm describing these objections as 'Foot stomping' in honor of Philippa Foot who rightly saw that there was often little good to be found in the states of affairs that the consequentialists took to be desirable and saw that seeing this was the key to seeing what's wrong with consequentialist reasoning. See Foot (1985) for discussion.

Gettier's Experience Machine: Gettier gets a good deal on Nozick's experience machine at a garage sale. He decides to step in for a while and see how it works. It works exceptionally well, so well in fact that he forgets that he's in it. It would please Gettier greatly if the Cubs made it to World Series. It seems to him just as if he's watching the game and watching them clinch the playoffs. As it happens, the Cubs just then did make it to the playoffs. A smile stretches across Gettier's face. 'They won', he says, 'They finally did it.'³⁶

Gettier's reason for being happy isn't that the Cubs won. The cost of being in the machine is that you've lost touch with reality. His belief that the Cubs won is accurate, but being in that state while being in the machine doesn't put him in a position to be happy for the reason that the Cubs won. That cannot guide him rationally in his feelings, thoughts, or deeds. Knowing what we know, we see that Gettier has lost track of reality. I don't see much good that comes from being in the state of mind that he's in even though it is accurate. If the point of full belief is to provide potential motivating reasons, this accurate belief fails and is a bad belief.

CVEC predicts that beliefs that are *pro tanto* good because of their categorical accuracy might also be harmful in a way that bears on their rational status. Consider a case that they'd describe as a case of harmful accuracy:

Trump: For years I have believed that I would never believe that Trump would win the GOP primary and never know that Trump would win the GOP primary. Upon opening the paper, I learn that he won. This

³⁶ We might be able to construct similar examples to show that there's worthless gradational inaccuracy. The key to constructing such an example would be to describe a set of credences that have the good-making features that gradational veritists key in on while cutting the credences off from the sorts of systems that should be involved in their rational maintenance and revision. If, say, the evidence that merits having a certain credence in a proposition is provided by something external to the credences themselves (e.g., experiences or apparent memories) we could test intuitions about the comparative goodness of two identical sets of credences where one set fits with experience and memory and the other does not. I suspect that most people will see one set as better than another and that suggests that there's more to the good of good partial belief than gradational veritism suggests.

belief, while accurate, renders two standing beliefs mistaken.

While my belief about Trump's success in the primary has a high degree of expected categorical accuracy, it also adds to the inaccuracy of my total epistemic state since it renders two standing beliefs false. The attitude makes a direct contribution to accuracy and indirect contribution to inaccuracy. In this case, it seems that the harmful accuracy isn't harmful in a way that matters, not to rationality. Even if the belief ensures that I'm worse off overall in terms of categorical accuracy, knowledge of this fact doesn't undermine the belief's rational standing. The fact that the belief detracts from total accuracy simply doesn't matter at all to the assessment of the belief.

CVEC predicts that some beliefs that are *pro tanto* bad because their categorical inaccuracy might be beneficial in a way that bears on their rational status:

Flawed and Flawless: Agnes completes her latest work of non-fiction. Her fact-checker tells her that she found precisely one error in the manuscript but cannot recall where the mistake was. Agnes decides to note this in the book's preface and sends it off to the publisher. Consider two versions of the case. Flawed: The fact checker was right. There was an error on p. 237. Flawless: The fact checker was wrong. There was no error in the body of the book.

The preface sentence is part of Agnes' book. Let's suppose that Agnes believes each proposition expressed by the sentences in the book. In Flawed, the preface sentence expresses a true proposition, so it adds to the total amount of categorical accuracy. What about in Flawless? In Flawless, Agnes could not believe the proposition expressed by the preface sentence because there is no such proposition (and no such propositional attitude). If there were, it would either have to be accurate or inaccurate. It couldn't be the former, for then it would be the latter. It cannot be the latter, for then it would be the former. The accuracy contained in the body of the book functions as one of Prior's blocks and thus prevents Agnes from getting things right in the preface.³⁷

Finally, how bad should a belief set be if it contains falsehoods? It might seem that each false belief should realize some negative value, each true belief

³⁷ For a discussion of these blocks, see Prior (1961). I first learned about these blocks and their significance from John Hawthorne's presentation on the surprise exam paradox. It was Marian David who pointed out that the preface sentence would be a liar in cases like Flawless.

should realize some positive value, and the total value realized by the belief set is determined by simply adding up the positive value and subtracting out the negative value. Thus, when comparing two belief sets with equal numbers of true and false belief, they should contain equal amounts of categorical goodness and badness.³⁸ If these subjects' beliefs concern different propositions but agree in terms of their expected categorical accuracy, these subjects should count as equally rational. This overlooks something important, though, which is the significance of the subject matter of these subjects' beliefs. Suppose Agnes gets some non-normative facts right in w1 and some facts about what rationality requires of her wrong in w1. Suppose Agnes gets the non-normative facts wrong in w^2 but the facts about what rationality requires of her right. Even if we stipulate that Agnes has the same number of true beliefs and false beliefs in these worlds so that the only difference here concerns the subject matter of her beliefs, it seems that Agnes might have a greater number of rational beliefs in w2 than $w1.^{39}$ Mistakes about rationality matter to rationality in ways that mistakes about matters of fact do not, but it is hard to see how the Lockean might do justice to this fact if we stipulate, as it seems we can, that Agnes' evidence warranted having high credence in all her beliefs in these two worlds.⁴⁰ This suggests that

³⁸ This is not completely uncontroversial for reasons Sylvan (2013) discusses but this fits with the kinds of scoring rules used in the standard accuracy-first arguments.

³⁹ These kinds of cases put interesting pressures on the intuitions that underwrite the Immodesty Thesis, "The credences recommended by your own epistemic rule, given a body of evidence, should uniquely maximize expected accuracy for you" (Horowitz 2014: 43). Suppose this holds for both partial belief and full belief. Think about the cases of negative self-appraisal where it seems that rationality requires high credence in both p and the proposition that it's irrational to believe p. You should think that if you don't take that combination of full beliefs on that you would fail by your own lights to maximize expected categorical inaccuracy. You should also think that if you take those attitudes on that you would be irrational. Remember that the motivation for the Immodesty Thesis has to do with doing well by your own lights. It seems you wouldn't be doing well by your own lights if you add the attitudes or refrain if doing well is cashed out in veritist terms. If you don't add the attitudes, you don't love truth enough. if you add the attitudes because of a love of truth, you seem to suffer from a strange kind of misology.

⁴⁰ Lasonen-Aarnio (MS) and Field (2015) think that there can be rational 'mismatch' between beliefs about the requirements of rationality and the first-order attitudes that rationality permits, requires, or forbids.

it's a mistake to think that all that matters to rationality are considerations of expected categorical accuracy.⁴¹

The Foot stomping objection comes to this. When we see the various ways in which a state might be accurate or might make a contribution to accuracy or inaccuracy, it seems that there is often nothing bad in what CVEC classifies as overall bad and nothing good in what CVEC classifies as overall good or *pro tanto* good. The direct inaccuracy of our beliefs matters, but indirect inaccuracy sometimes does not, not when it comes to assessing the particular belief. The direct accuracy of our beliefs sometimes matters, but not always, and when it matters it matters in a way that indirect accuracy does not. There is something wrong with the consequentialist concern for aggregate accuracy and the role that belief plays in bringing it about. Thus, we might acknowledge that other belief sets do better in terms of overall goodness as characterized by categorical veritism and be unmoved by this.

4.2.2 Goodness and Grain

I want to highlight a further way in which it is a mistake to fixate on accuracy in the way that accuracy-first epistemologists have. We'll look at a worry about granularity. Accuracy is too coarse-grained to do the work CVEC needs it to do.

In the value theory there are two goods to consider, goods that play different roles. According to categorical veritism, actual or objective epistemic goodness is a function of the accuracy or inaccuracy of a subject's full beliefs and it determines how well things are going for the subject. Expected categorical goodness is a function of the probability of categorical accuracy or inaccuracy and it determines what's rational for a subject to believe, according to CVEC. Let's consider the categorical veritist's approach to objective epistemic goodness. Let's suppose that propositions and propositional attitudes are individuated in a Fregean way. Beliefs in two propositions that differ only at the level of sense will count as distinct beliefs.

⁴¹ For arguments that there cannot be rational false beliefs about the requirements of rationality, see Littlejohn (forthcoming) and Titelbaum (2015). As Dorst (MS) notes, a Lockean needn't have a single, fixed threshold for all beliefs for the Lockean might assign different amounts of disvalue to different falsehoods on the basis of the beliefs' subject matters. This has the effect of raising the threshold of evidential support needed to rationally believe the relevant propositions, but I fear that this strategy raises the bar too high to be plausible and I fear that there is no particularly good independent motivation for thinking that the disvalue associated with false beliefs about rationality is greater than the disvalue associated with, say, false beliefs about astronomy or the law.

Accuracy is assessed at the level of reference, not sense. The belief that Hesperus shines is accurate iff the belief that Phosphorus shines is. Suppose Agnes doesn't realize that Hesperus is Phosphorus, believes that Hesperus shines in the morning, and believes that Phosphorus shines in the evening. Let's suppose that these beliefs constitute knowledge and are supported by the evidence. In short, let's suppose that Agnes is doing very well, epistemically. Suppose that Agnes' twin is nearly identical to Agnes. She has all of Agnes' evidence concerning the propositions that Hesperus shines in the morning and that Phosphorus shines in the morning. Both propositions have high degrees of expected categorical accuracy. Suppose, however, that her mental life differs from Agnes' *only* in these two respects: instead of believing that Hesperus shines in the morning she believes that Phosphorus does and instead of believing that Phosphorus shines in the evening she believes that Hesperus does.

While Agnes' and Agnes' twin have beliefs that score equally well in terms of objective categorical accuracy, they don't do equally well in terms of objective epistemic goodness. In spite of the fact that their beliefs concerning Hesperus and Phosphorus are equally accurate representations of that heavenly body, Agnes' attitudes are preferable because they are intelligible to Agnes in a way that her twins' attitudes couldn't be. Objective epistemic goodness depends upon accuracy *and* the way that things are presented. The twin's beliefs are perfectly accurate but they don't seem to be particularly good states to be in. Modes of presentation matter and they don't just matter to rationality or expected goodness. This raises a difficult question for the accuracy-first epistemologists. Why should something like mode of presentation or sense matter for objective epistemic value if the cardinal epistemic good is accuracy or truth?⁴²

⁴² We'll see that a rival account offers a straightforward explanation. Proponents of accuracy-first epistemology might want to employ a fine-grained notion of a truth and say that 'Hesperus shines' and Phosphorus shines' express different truths to try to get a handle on these kinds of problems but difficult questions would remain. Why, we might ask, should these believing these truths differ in value if they don't differ with respect to their accuracy? Even if their view of truth and the individuation of truths allows them to say that Agnes had different truths in mind in the examples just described the veritist thinks that the number of truths and falsehoods determines the total value of a subject's epistemic state. Thus, they cannot use some fine-grained notion of truth to try to vindicate the intuition that Agnes would be worse off in terms of objective or actual value if her beliefs were altered in the way described. We'll see that the alternative value theory proposed in the next section provides a natural answer to this question and handles many of the other difficulties that arise for veritism as well.

This problem is masked to some extent because expected, not actual, categorical accuracy is what matters most directly to rationality and expected categorical accuracy is the sort of thing that will be sensitive to modes of presentation. Still, there are problems with granularity and expected categorical accuracy that arise for CVEC. Recall the cases involving conjunctions of lottery propositions and negative self-appraisals. It is irrational to believe lottery propositions and highly irrational to believe the conjunction of such a proposition with a negative epistemic self-appraisal (e.g., to believe that this ticket is a loser and that I don't know that it is). It seems worse from the point of view of rationality, however, for Agnes to believe that her ticket is a loser and that she doesn't know that it is than it would be for Agnes to believe this about others who hold tickets for the same lottery.⁴³ Agnes shouldn't believe on the basis of purely statistical information that her friend's ticket will lose and that her friend doesn't know this, but it doesn't seem that this belief that involves a negative epistemic appraisal of someone else is as deeply and strikingly irrational as beliefs that involve negative appraisal of her currently held attitudes. In spite of a striking difference in how irrational these beliefs would be they do not differ with respect to their expected categorical accuracy. The properties that determine whether a belief constitutes a Moorean absurdity do not supervene upon the properties that determine actual or expected categorical accuracy.

4.3 Trade-Offs

We finally we come to the familiar problem of trade-offs.⁴⁴ We've seen a version of the problem in the course of discussing GVEC. Once we find a version of the

⁴⁴ The earliest discussions of trade-offs were Firth (1981) and Jenkins (2007). In Littlejohn (2012), I argued that the trade-off problem was a problem for Goldman's rule-consequentialist arguments for reliabilism. The trade-off problem is also discussed in Berker (2013) and Greaves (2013). For responses to

⁴³ There is another striking difference to note. Consider the rational status of Agnes' belief that her mother's ticket lost and that her mother's belief that it lost doesn't constitute knowledge. Suppose it is based, in part, upon her mother's testimony and what she read in the paper. This belief about her mother's belief and ticket would have a high degree of expected categorical accuracy but it might have a lower degree of expected categorical accuracy than a belief she has about a ticket she holds for some other lottery and her belief in her own inability to know that this ticket is a loser. The properties that make for the irrationality of the belief include those that determine whether the belief is about your own beliefs, not just its expected categorical accuracy.

view that introduced the incentives necessary to address the problems of partial opinionation and population, we saw that the view would deliver the wrong verdict in trade-off cases by classifying rational beliefs as irrational because of their opportunity costs. Even if some belief constitutes knowledge, the opportunity cost associated with holding this belief might be so great that CVEC would classify the belief as irrational on the grounds that suspension or disbelief has a greater expected amount of categorical accuracy. (Maybe knowing how clever you are will keep you out of Athena's private library.) It seems that however bad a belief might be, the expected boost in categorical accuracy might be so great that CVEC would classify the belief as rationally required. (Maybe the key to living long enough to benefit from time in Athena's library is the belief that Zeus is the greatest conceivable being.)

Trade-offs arise because consequentialists are concerned with the total value associated with options where an option is understood as the world that would or could be realized if a subject were to form a belief. In turn they hold that there's an important difference between *pro tanto* harms and benefits that come in the form of the categorical accuracy or inaccuracy of a full belief and the overall harms and benefits that we only determine by calculating the total categorical accuracy associated with belief, disbelief, and suspension contained in an option. These cases have been discussed elsewhere, but I think that some implications of these cases has been overlooked. These implications have to do with CVEC's attitude towards pretty patterns and putative principles.

Suppose a subject knows that her belief in *p* would be both a belief that is highly probable *and* a belief that would prevent the believer from maximizing actual categorical accuracy (e.g., the benefits associated with believing that a lottery ticket is a loser is good enough that the way to maximize total categorical accuracy would be to refrain from believing the ticket will win.) Such cases show that if it's possible to have high credence without full belief, this is what rationality would require according to CVEC. Thus, CVEC implies that there are possible situations in which it's rational to have high credence without full belief, a situation that would be a counterexample to the normative Lockean thesis. If there's an argument from CVEC to the conclusion that there are counterexamples to the Lockean thesis, we shouldn't turn to consequentialism to find a vindication for the Lockean thesis.

Suppose Agnes knows p and knows that some of p's consequences will be revealed if a card is flipped over. Agnes also knows that, owing to the strange causal structure of the world, if she believes these consequences the total

Berker and Greaves, see Ahlstrom-Vij and Dunn (2014) and Konek and Levinstein (MS) respectively.

categorical accuracy of her beliefs will be worse than they'd be if she believed p and failed to believe some of the consequences of p written on the card. If she were to believe what's written on the card, these beliefs would do well in terms of its categorical accuracy but CVEC implies that she should refrain from holding them. We can use cases like this to cause trouble for closure. If it's possible to believe p and refrain from believing p's consequences, this combination might do better than alternatives in terms of expected categorical accuracy and would thus be rationally preferable to belief sets that involved both the belief in p and its consequences. CVEC implies that there can be situations in which rational belief isn't closed under known entailment. Tinkering with the causal structure of the world and the costs and benefits associated with being in certain states of mind, we should be able to use CVEC to generate counterexamples to principles of all sorts and show that the prettiness of a pattern is an unreliable guide to the rational standing of the beliefs that figure in it.

5. Three Mistakes

I have identified a variety of problems with the accuracy-first approach to norms for partial and full belief. What should we take from this? I think proponents of the accuracy-first approach make two mistakes about the fundamental epistemic good and are liable to make a further mistake about the relationships between partial and full belief.

It is a mistake to think of accuracy as the cardinal epistemic good. It is merely required for the realization of that good because, as we saw, some accuracy is worthless. In §4 we saw that two epistemic states that do equally well in terms of accuracy might not do equally well in terms of their total objective epistemic value. We also saw that certain kinds of purely accidental connections between belief and fact meant that beliefs that matched the facts weren't good because they failed to track reality. In §3 we saw that the strongest argument for categorical veritism was one that focused on the value that beliefs realize when they track reality. A belief provides a potential motivating reason iff it tracks reality. An accurate belief can fail to track reality in the right way, but not if it constitutes knowledge. It turns out that the best argument for veritism is actually an argument for thinking that knowledge, not mere true belief, is the fundamental epistemic good. This is because a belief provides a potential motivating reason iff it constitutes knowledge. Some accurate beliefs don't do what they're supposed to.

It's easy to make sense of the points about epistemic goodness in §4 if we replace veritism with *conscientiaism*, the view that knowledge is the cardinal epistemic good. Conscientiasm also helps us make sense of the counterexamples to the normative Lockean thesis. The recipe for counterexamples should now be

obvious. We look for beliefs that do well in terms of expected categorical accuracy. These would be propositions that the Lockean and proponents of CVEC would classify as rationally permitted. (It would if we bracket trade-offs.) We then focus on the subset of these beliefs that are 'known unknowns' and find our counterexamples. ⁴⁵ We know *apriori* that we cannot know lottery propositions, so it's not rational to believe them. We know *apriori* that we cannot know the negative epistemic self-appraisals and Moorean absurdities, so we cannot rationally believe them.⁴⁶ We know it matters if we know *apriori* that some belief doesn't constitute knowledge because that means we know *apriori* that these are bad beliefs, beliefs that cannot do what they're supposed to.

It might be thought that an epistemic consequentialist could avoid these difficulties by embracing conscientiaism, but they cannot. Conscientiaism isn't for consequentialists. Epistemic consequentialism assumes that the fundamental epistemic good is a promotional good. It would be perverse to apply decision-theoretic norms to goods that aren't promotional goods (e.g., things that are attributively good like good assassins or good toasters or things that are normatively good (i.e., good in the sense that something is good by virtue of being right, appropriate, or fitting).) I think we have two good pieces of evidence that the fundamental epistemic good isn't a promotional good.

First, if the fundamental epistemic good were a promotional good, it would be irrational to believe when it's improbable that a belief realizes the relevant good. Think about the way in which CVEC implies that if p is improbable, it is irrational to believe p. If we modify the value theory by swapping knowledge for true belief and belief that fails to constitute knowledge for false belief, would we want to say that if it is improbable that p is known that it is irrational to believe p? No, not if it's possible to know p when it isn't probable that you know p.⁴⁷ As Flannery O'Connor might have said, no man with a good

⁴⁵ This is Sutton's (2007) memorable terminology. The key thing to notice is that knowledge requires both high probability and further conditions that we know needn't be met when beliefs have high probability. The judgment that you're not in a position to know p doesn't require you to think that the reason for failing to know p is that p isn't sufficiently probable on your evidence. It might be that you don't know because success wouldn't be attributable to ability, your belief wouldn't be safe, your belief wouldn't be sensitive, etc.

⁴⁶ We know *apriori* that we cannot know < p and I don't know p > because of the factivity of 'knows'. Assuming that knowledge of p implies rational belief in p (for mature humans, at any rate) we can see why we cannot know < p and it's not rational for me to believe p >.

⁴⁷ See Williamson (2011).

belief needs to be justified. If you know, that's good enough for a rational belief. Rationality doesn't require that a belief is probably an instance of knowledge, only that it is.

Second, we shouldn't forget that nearly everyone agrees that trade-offs are perverse. If the fundamental epistemic good were a promotional good, it would be perverse *not* to accept trade-offs. Refusing to accept a trade-off would be like refusing to accept an acknowledged greater good for a lesser one. That isn't how we see things when we think about Agnes' knowledge of history. We don't think that it is irrational for Agnes to have the beliefs that constitute knowledge even when we recognize that possessing this knowledge comes at a cost and precludes her knowing lots of things that she could have learned by doing mental math.⁴⁸

These points about the importance of knowledge and the kind of good knowledge isn't are sufficient to show that the accuracy-first project is fundamentally flawed, but I want to close by making a point about full and partial belief. Suppose everything went as planned. All the counterexamples to the normative Lockean view work. We accept that the fundamental epistemic good is knowledge because the realization of that good requires a belief that tracks reality, not just a belief that happens to match it. We now have the materials for an argument against a simple version of the metaphysical Lockean view, the view that says that a subject has a full belief iff she has a sufficiently high degree of confidence.⁴⁹

On this view, a subject couldn't refrain from believing p if she had a sufficiently high degree of confidence in p because her belief in p just is this high degree of confidence in p. If we suppose that 'ought' implies 'can', the examples discussed above that show that a subject often ought to refrain from believing even when she rationally has a sufficiently high degree of confidence show that it's possible to have high credence in a proposition that isn't believed. When a

⁴⁸ In Littlejohn (forthcoming b), I provide further arguments for conscientiaism and for a teleological and non-consequentialist approach to epistemic norms for justification and for rationality. On this view, a belief would be justified if it did what beliefs are supposed to do (i.e., provide us with potential motivating reasons) and would be rational if formed in such a way that it or some similar belief would constitute knowledge. For discussion of the link between knowledge and rationality, see Bird (2007). One nice feature of Bird's account is that it treats the preface and the lottery differently. Lottery beliefs are irrational because they are known unknowns and preface beliefs are sometimes rational because they are potential cases of knowledge. The accuracy-first approach seems to treat them similarly.

⁴⁹ See Sturgeon (2008) for a helpful discussion of the metaphysical Lockean thesis.

prospective belief is flagged as an unknown, it's possible to both have high credence in p and fail to have a full belief in p. If the belief would be a known unknown, it would be right to have high credence without full belief because the full belief couldn't do what it's supposed to do and the high credence might well do what it's supposed to do.⁵⁰

Many epistemologists see full belief and partial belief as playing similar roles (e.g., pushing you around in certain ways). If you start to think of partial and full belief this way, it's not surprising that you'd be tempted to accept a version of the metaphysical Lockean picture. The picture is much less attractive, however, once we identify a point or purpose for full belief that mere partial belief cannot serve. A full belief is supposed to track the facts so that it can provide you with a potential motivating reason, a fact that can be your reason for believing, feeling, or doing something. No partial belief can do that.⁵¹ Because of this, it makes sense to have a norm that, *inter alia*, enjoins us not to have a full belief in *p* when we know we couldn't know p.⁵² However high our confidence gets in *p*, if we know that we cannot know, we know that a full belief could not do what it's supposed to do. The partial belief, however, could presumably fulfill its function when the strength of this belief is proportional to the probability of the target proposition.

While we haven't yet specified the point or purpose of partial belief, this suggests that we should talk about points and purposes of full belief and partial belief respectively, acknowledge that they answer to divergent normative standards, and recognize that a partial belief however strongly held needn't constitute a full belief. Partial beliefs come into their own when we have middling credences or have to reason about known unknowns. Full beliefs do their work by tracking facts so that they can serve as our reasons for believing, feeling, and doing what we do. This proposal allows for a kind of unity in the cases where it is an open possibility that the subject could come to know. In these

⁵⁰ For further discussion of the possibility that the relation between partial and full belief might be more complicated than the simple Lockean picture presupposes, see Buchak (2014), Locke (2014), and Staffel (forthcoming).

⁵¹ Consider Adler's (2002: 17) observation that mild resentment is the result of a full belief that someone did something a little bit bad, not the product of having some partial belief with middling credence that someone did something unspeakably bad.

⁵² I take this norm to be derivative from the more fundamental norm that enjoins you to refrain from believing p when your belief in p wouldn't constitute knowledge but the present account doesn't assume that this is the fundamental norm of belief.

cases it's possible that a sufficiently strong partial belief might constitute a full belief, but attending to the cases where the subject appreciates that she's not in a position to know helps us recognize the differing purposes of partial and full belief and the need to reject the Lockean suggestion that the norms for full belief simply grow out of the norms for partial belief. The aim of full belief is knowledge. The aim of partial belief is something else, something it can attain when we know we cannot know. Partial belief and full belief need norms of their own because they serve different purposes.

References

Adler, Jonathan. 2002. Belief's Own Ethics. MIT Press.

Ahlstrom-Vij, Kristoffer and Jeffrey Dunn. 2014. A Defence of Epistemic Consequentialism. *The Philosophical Quarterly* 64: 541-51.

Alvarez, Maria. 2010. Kinds of Reasons. Oxford University Press.

Baron, Marcia. 1995. Kantian Ethics Almost Without Apology. Cornell University Press.

- Berker, Selim. 2013. The Rejection of Epistemic Consequentialism. *Philosophical Issues* 23: 363-387.
- Bird, Alex. 2007. Justified Judging. *Philosophy and Phenomenological Resarch* 74: 81-110.
- Brunero, John. 2013. Rational Akrasia. Organon F 20, 546–566.
- Buchak, Laura. 2014. Belief, Credence, and Norms. *Philosophical Studies* 169: 285-311.
- Carr, Jennifer. 2015. Epistemic Expansions. Res Philosophica 92: 217-36.
- Di Finetti, Bruno. 1974. Theory of Probability, Volume 1. John Wiley and Sons

Dorst, Kevin. MS. Lockeans Maximize Expected Accuracy.

Easwaran, Kenny.

Firth, Roderick. 1981. Epistemic Merit, Intrinsic and Instrumental. *Proceedings and Addresses of the American Philosophical Association* 55: 5-23.

Fitelson, Branden. MS. Coherence.

- Easwaran, Kenny and Branden Fitelson. 2015. Accuracy, Coherence, and Evidence. In Tamar Szabo Gendler and John Hawthorne (eds.). *Oxford Studies in Epistemology*: Vol. 5. Oxford University Press: 61-96.
- Field, Claire. 2015. *Rational Incoherence: How to Make Mistakes about what Rationality Requires*. Master of Philosophical Studies Thesis, King's College London.
- Foot, Philippa. 1985. Utilitarianism and the Virtues. Mind 94: 196–209.
- Gibbard, Alan. 2007. Rational Credence and the Value of Truth. In T. Gendler and J. Hawthorne (ed.), *Oxford Studies in Epistemology* 2. Oxford University Press, pp. 143-64.

Goldman, Alvin. 1999. Knowledge in a Social World. Oxford University Press.

Gordon, Robert. 1987. The Structure of Emotion. Cambridge University Press.

- Greaves, Hilary and David Wallace. 2006. Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind* 115: 607-32.
- Greaves, Hilary. 2013. Epistemic Decision Theory. Mind 122: 915-52.
- Horowitz, Sophie. 2014. Immoderately Rational. Philosophical Studies 167: 41-56.

Hyman, John. 2015. Action, Knowledge, and Will. Oxford University Press.

- Jenkins, Carrie. 2007. Entitlement and Rationality. Synthese 157: 25-45.
- Joyce, James. 1998. A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65: 575-603.
- _____. 2009. Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In F. Huber and C. Schmidt-Petri (ed.), *Degrees of Belief.* Springer, pp. 263-97.
- _____. MS. Why Evidentialists Need not Worry About the Accuracy Argument for Probabilism.
- Konek, Jason and Ben Levinstein. MS. The Foundations of Epistemic Decision Theory.
- Kvanvig, Jonathan. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Lasonen-Aarnio, Maria. MS. Enkrasia or Evidentialism? Learning to Love Mismatch.
- Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge University Press.
- Littlejohn, Clayton. Forthcoming. Stop Making Sense? On a Puzzle about Rationality. *Philosophy and Phenomenological Research*.
- Littlejohn, Clayton. Forthcoming b. The Right in the Good: A Defense of Teleological Non-Consequentialism in Epistemology. In K. Ahlstrom-Vij and J. Dunn (ed.), *Epistemic Consequentialism*. Oxford University Press.
- Locke, Dustin. 2014. The Decision-Theoretic Lockean Thesis. Inquiry 57: 28-54.
- Lynch, Michael. 2004. True to Life. MIT University Press.
- McGlynn, Aidan. 2013. Believing Things Unknown. Nous 47: 385-407.
- Nelkin, Dana. 2000. The Lottery Paradox, Knowledge, and Rationality. *Philosophical Review* 109: 373–409.
- Owens, David. 2013. Value and Epistemic Normativity. *Teorema* 33: 35-58.
- Pettigrew, Richard. 2013. Epistemic Utility and Norms for Credences. *Philosophy Compass* 8 (10): 897-908.
- Pettigrew, Richard. 2016. Accuracy and the Laws of Credence. Oxford University Press.

- Pettigrew, Richard. MS. The Population Ethics of Belief: In Search of Epistemic Theory X.
- Prior, Arthur. 1961. On a Family of Paradoxes. *Notre Dame Journal of Formal Logic* 2: 16-32.
- Ross, W.D. 1930. The Right in the Good. Oxford University Press.
- Ryan, Sharon. 1996. The Epistemic Virtues of Consistency. *Synthese* 109: 121-41.
- Smithies, Declan. 2012. Moore's Paradox and the Accessibility of Justification. *Philosophy and Phenomenological Research* 85: 273-300.
- Sorensen, Roy. 1988. Blindspots. Oxford University Press.
- Staffel, Julia. Forthcoming. Beliefs, Buses, and Lotteries: Why Rational belief Can't Be Stably High Credence. *Philosophical Studies*.
- Sturgeon, Scott. 2008. Reason and the Grain of Belief. Nous 42: 139-165.
- Sutton, Jonathan. 2007. Without Justification. MIT University Press.
- Sylvan, Kurt. 2013. Truth Monism without Teleology. *Thought* 1: 161-9.
- Talbot, Brian. 2014. Truth-Promoting Non-Evidential Reasons for Belief. *Philosophical Studies* 168: 599-618.
- Titelbaum, Mike. 2015. Rationality's Fixed Point (Or: In Defense of Right Reason). In J. Hawthorne (ed.) Oxford Studies in EpistemologyOxford Studies in Epistemology 5. Oxford University Press.
- Titelbaum, Mike. 2015b. How to Derive a Narrow-Scope Requirement from Wide-Scope Requirements. *Philosophical Studies* 175: 535-42.
- Unger, Peter. 1975. Ignorance. Oxford University Press.
- Wedgwood, Ralph. 2003. Choosing Rationally and Choosing Correctly. In S. Stroud and C. Tappolet (ed.), Weakness of Will and Practical Irrationality. Oxford University Press, pp. 201-29.
- Williamson, Timothy. 2000. Knowledge and its Limits. Oxford University Press.
- Williamson, Timothy. 2011. Improbable Knowing. In T. Dougherty (ed.), *Evidentialism and its Discontents*. Oxford University Press, pp. 147-64.
- Zimmerman, Michael. 2015. Value and Normativity. In I. Hirose and J. Olson (ed.) *The Oxford Handbook of Value Theory*. Oxford University Press, pp. 13-29.