REPUGNANT ACCURACY

Brian Talbot

Accuracy-first epistemology is an approach to formal epistemology which takes accuracy to be a measure of epistemic utility and attempts to vindicate norms of epistemic rationality by showing how conformity with them is beneficial. If accuracy-first epistemology can actually vindicate any epistemic norms, it must adopt a plausible account of epistemic value. Any such account must avoid the epistemic version of Derek Parfit's "repugnant conclusion." I argue that the only plausible way of doing so is to say that accurate credences in certain propositions have no, or almost no, epistemic value. I prove that this is incompatible with standard accuracy-first arguments for probabilism, and argue that there is no way for accuracy-first epistemology to show that all credences of all agents should be coherent.

1. Introduction

Accuracy-first epistemology takes accuracy to be a measure of epistemic utility, and attempts to derive and explain norms of epistemic rationality by looking at how conformity with putative norms affects the utility of one's overall belief state. For example, accuracy-first epistemology explains why we are rationally required to have credences that are probabilistically coherent by showing that, for any credal state with incoherent credences, there is a coherent version of it that is more accurate–and thus epistemically superior–in all possible worlds.[1] One of the attractions

[1] See, e.g., Joyce 1998.

of accuracy-first epistemology is that it can *vindicate* norms of rationality: it can show why these really are norms by showing how adherence to them is beneficial in terms of epistemic utility.[2] Accuracy-first epistemologists have not only proposed vindications of coherence requirements, but also of, for example, norms requiring conditionalization and norms for responding to disagreement.[3] In this paper I argue that, if accuracy-first epistemology has any hope of actually vindicating any norms of rationality, it must see accuracy about some topics as almost (or entirely) valueless. However, if it does, then the vindication of coherence as a norm of rationality does not succeed, nor do vindications of other norms of rationality which assume the coherence norm.

Before we start, let's establish some terminology. Accuracy-first epistemology focuses on rationality norms that apply to *credal states*, which are sets of *credences*. Credences can be understood best by contrasting them with so-called "full beliefs." Full beliefs are attitudes towards propositions that come in only a few states: belief, disbelief, or suspension of judgment. Credences, sometimes called "degrees of belief" or "degrees of confidence," are attitudes towards propositions that come in a larger range of states; we'll assume that these can be modeled by real numbers between 0 and 1 (it may be that human credences are less precise, but this won't affect my arguments).[4] Let's use the term *neutral* for whichever credal value separates the range of

---

[2] See e.g. Oddie 1997, Joyce 1998, Greaves & Wallace, 2006, Leitgeb and Pettigrew 2010, Pettigrew, 2016b.

[3] See Oddie 1997 and Greaves & Wallace 2006 on vindications of conditionalization, and e.g. Moss 2011, Staffel 2014, Levinstein 2015 on disagreement.

[4] Recent work extends some results of accuracy-first epistemology to full beliefs (e.g. Easwaran 2015). My arguments can, with some modifications, be applied to these extensions. In this

accurate credences in a proposition from the range of inaccurate ones (it may be that this value can differ in different cases). A credence in *p* is *accurate* when it is greater than neutral when *p* is true, or less than neutral when *p* is false. A greater-than-neutral credence in a false proposition, or a less-than-neutral credence in a true one, is *inaccurate*. To illustrate, consider a credence in the proposition that a particular coin flip came out heads. Intuitively, the neutral point for this credence is .5. If the coin did come up heads, then any credence above .5 is accurate; if the coin actually came up tails, then any credence above .5 is inaccurate. Accuracy and inaccuracy come in degrees: the closer an accurate credence in a true proposition is to 1, the more accurate it is, whereas the closer an inaccurate credence in a falsehood is to 1, the more inaccurate it is. Let's call accurate credences that are very close to the neutral point *minimally accurate* (I'll leave this term intentionally vague, as my arguments won't depend on any particular precisification). If our flipped coin came up heads, for example, a credence of .501 that it came up heads is minimally accurate.

It seems intuitive that it is epistemically better to be somewhat confident in a truth than to be completely uncertain about it, and that being more confident is even better. If the neutral point for a credence in a given proposition is typically the point that also represents maximum uncertainty about that proposition's truth–e.g. a .5 credence that a flipped coin came up heads–then we have intuitive support for the claim that accurate credences are better than neutral ones, and get better as they get more accurate.[5] We can find similar support for the claim that

---

paper, I'll only be discussing credences since this is what accuracy-first epistemology tends to focus on.

[5] I say "typically" here because there may be credences for which there is no unique or determinate value representing maximum uncertainty. Cases that have been used to present

inaccurate credences are worse than neutral ones. This gives us a plausible starting point for a theory of epistemic value (we'll consider some alternative views in section 2). Throughout section 2, I discuss a problem accuracy-first epistemology faces when more fully articulating its account of epistemic value, and I show how this problem must be solved. Accuracy-first epistemology uses its account of value to explain what it is to be rational. In section 3, I discuss the implications of this solution for accounts of epistemic rationality.

2. The epistemic repugnant conclusion

Accuracy-first epistemology faces a serious challenge when trying to articulate a plausible account of epistemic value. The challenge is related to one that was first discussed in the context of ethical consequentialism. It is perhaps unsurprising that similar issues arise for both accuracy-first epistemology and ethical consequentialism, since both try to explain oughts in terms of value. As we will see, however, because of disanalogies between ethics and epistemology, some responses to this problem that are sensible in one domain do not work in the other. I'll start by laying out the ethical version of the problem and then show how it translates to epistemology.

Not all possible people actually exist. In light of this, ethical consequentialists must tell us how adding new people to a world changes the overall utility of that world. Intuitively, when something is good, more of it is better. And, intuitively, a minimally decent life is at least

---

problems for the indifference principle might give us examples of these (see Carr 2015 for discussion). Although as Julia Staffel has pointed out to me, one might still think agents in such cases can have neutral credences over a range of options that they consider symmetrical, by assigning each equal credence. It's worth noting that none of my arguments in this paper depend on saying that the line dividing accurate from inaccurate credences should be the point the represents maximum uncertainty.

somewhat good. So, it seems initially plausible to say that adding new people to a world, when those people have at least minimally decent lives, increases the total utility of that world. But this leads to the "repugnant conclusion"–the conclusion that, for any seemingly good world, there is a better world containing no lives that are more than minimally decent (Parfit 1984). This repugnant conclusion is implausible, and ethical consequentialists must either give an account of value that denies it, or explain why the conclusion is not a *reductio* of their theory.

Similarly, the credal states of ordinary people do not contain credences about all propositions (Carr 2015, Pettigrew 2016a). For example, philosophy often introduces us to new questions, questions whose answers we had just never considered before, or whose answers we did not have the conceptual apparatus to entertain at all. In light of this, accuracy-first epistemology must tell us how adding new credences to a credal state changes its overall epistemic utility. If accurate credences are good, then *prima facie* more of them are better. As Jennifer Carr (2015) points out, this leads to the *epistemic repugnant conclusion* (ERC): the conclusion that, for any seemingly good (finite) credal state, there is a better (finite) state containing credences that are no more than minimally accurate (see also Pettigrew 2016a). Since I will be talking a lot about the ERC, it'll help to have labels to make its parts easier to refer to. Let's call the credal state with the large number of minimally accurate credences the *repugnant credal state* and the seemingly good state we compare it to the *attractive credal state.*

The ERC is implausible: it seems false that having the vaguest inkling of the truth about a vast number of things is better than knowing the truth about a smaller, but still decent, number of questions. To illustrate just how implausible the ERC is, let's consider an example. Consider an attractive credal state which contains only extremely high credences in all the wisdom that humanity will ever acquire. Assume "wisdom" is factive, so these are all high confidences in true propositions, and thus extremely accurate, and assume humanity acquires a finite but very large

amount of wisdom during our time in history. Contrast this with a repugnant state that contains nothing but a vast number of minimally accurate credences, each of which is about whether there is a particle in some arbitrary location in space and time (each credence is about a different location, so these are credences in distinct propositions).[6] If we accept the ERC, we would have to say that some repugnant state of this sort is superior to certainty about all human wisdom. But it is obviously false that such a repugnant state is better in any meaningful way than certainty about all human wisdom, no matter how large the repugnant credal state is. Perhaps this repugnant state has *more* of something than the credal state containing all human wisdom, but that does not make it *better*. One might worry that this confounds the practical with the epistemic, but to clear this up we can imagine an attractive state containing practically useless but still clearly valuable accurate credences–maximal confidences in elegant but impractical mathematic and logic theorems, or about deep issues in philosophy with no practical upshot. It is still obviously false that the repugnant state is superior to this attractive one.

This means that, if we accept the ERC, then we accept that what accuracy-first epistemology sees as epistemic utility is not something truly valuable. Accuracy-first epistemology is supposed to vindicate norms of rationality by showing how conformity with these norms is beneficial. If epistemic utility is not truly valuable, then the fact that conformity with certain norms of rationality increases epistemic utility does not vindicate these norms. Thus, if accuracy-first epistemology is to succeed, it must find a way of saying that the ERC is not true. I'll take this for granted now, and argue for it further in section 2.4.

[6] My thanks to Jacob Ross for suggesting this example.

2.1. Orthodox "solutions" to the ERC

Certain orthodox views in accuracy-first epistemology seem to suggest ways to deny the ERC. However, we'll see in this section that these "solutions," just like the ERC, prevent accuracy-first epistemology from vindicating norms of rationality. We will also see that accuracy-first epistemology is not really committed to these "solutions."

Many formal epistemologists favor Brier scores as measures of accuracy. The Brier score for a particular credence c in a true proposition is $(c\text{-}1)^2$; if the proposition is false, the score is $(c\text{-}0)^2$. The best Brier score a credal state can have is 0. It might seem natural to interpret Brier scores as negatively assessing credences for their distance from the truth, but not positively assessing credences for closeness to the truth. On this interpretation, credences can only be more or less bad, and are never positively valuable. Let's call this view *inaccuracy only*. If epistemic value really worked this way, then we would avoid the ERC–repugnant states would be worse than attractive states, since they contain more credences which are farther from the truth.

Another solution to the ERC is suggested by the following fact: proofs in accuracy-first epistemology often proceed by only allowing accuracy comparisons between credal states that contain credences toward all and only the same propositions (e.g. Joyce 1998, Pettigrew 2016b). In light of this, one might claim that the utilities of credal states that contain credences on any different propositions really are incomparable. Let's call this view *no comparisons*. If the no comparisons view were true, then repugnant states could not be better than attractive states, as repugnant states must contain credences on more propositions.

Both of these views give us clearly false accounts of epistemic utility (for reasons noted in Carr (2015)). So endorsing either would prevent accuracy-first epistemology from vindicating norms of rationality. This can be easily illustrated. Consider a credal state containing one maximally accurate credence, and no other credences (say it assigns credence 1 to one truth); let's

call this state *One*. Compare this to the omniscient credal state, which assigns 1 to all truths and 0 to all falsehoods. The omniscient state must be epistemically superior to *One*, but neither *inaccuracy only* or *no comparisons* can say that it is. Alternately, consider the credal state of an agent who, like us, cannot rule out the existence of an evil demon, but is otherwise as certain as they rationally can be in all contingent truths. Let's say this agent has extremely, but not maximally, accurate credences in all contingent truths, and no inaccurate credences. According to *no comparisons*, their credal state is no better or worse than *One*. Their state is worse than *One* according to *inaccuracy only*.[7] Neither of these claims are plausible. Accepting *no comparisons* or *inaccuracy only* solves the ERC, but still gives us an account of epistemic utility which fails to correspond to any recognizable good. Such an account of epistemic utility cannot be used to vindicate norms of rationality.

One might wonder if I'm too fast to dismiss these views, given that *no comparisons* and *inaccuracy only* are so similar to ideas accepted by many formal epistemologists. We should be careful here, however. These views I have just dismissed are views about epistemic utility, while the views in orthodox accuracy-first epistemology they are related to are about rationality. Orthodox accuracy-first epistemology's claims about rationality can still make sense even though these parallel claims about value are clearly false. To see what I mean, let's consider *no comparisons*. This view is inspired by assumptions made in proofs intended to give us norms of

[7] Further, *inaccuracy only* entails that whenever an agent formed a new credence about any contingent claim, but was not certain about it, their new credal state would be weakly dominated by their prior state: the new state would have no epistemic benefits, but in some worlds where the contingent claim is false it would incur an epistemic cost (Carr 2015). This would be incompatible with certain proofs vindicating coherence (Carr 2015, Pettigrew 2016a).

rationality. To work, these proofs need to say that, when determining whether credal state *S* is rational, we only compare its utility to the utilities of other states that contain credences on all and only the same propositions as *S*.[8] This is a weaker claim than *no comparisons* makes – *no comparisons* says that these comparisons *cannot* be made. There is nothing clearly implausible about denying *no comparisons* while still holding the view about rationality that is built into standard proofs in accuracy-first epistemology. We see something similar in ethical consequentialism. On standard consequentialist views, we can compare the utilities of any state of affairs we like, including states of affairs that we cannot realize by our actions; but, when determining what a person ought to do, we only appeal to comparisons between utilities of states of affairs that that person could bring about. So, the sorts of assumptions that are made in accuracy-first epistemology that look a bit like the *no comparisons* view do not actually require the *no comparisons* view. That's good, because the *no comparisons* view is not at all plausible. Similarly, we can sensibly use Brier scores to determine what is rational without adopting the *inaccuracy only* view. We can mathematically transform Brier scores so that accurate credences are given positive values, thus denying *inaccuracy only*, without surrendering the features that make Brier scores useful in vindicating rationality norms (see e.g. Pettigrew 2016a). Even if this were not so, Brier scores could give us a model of epistemic utility which was useful for determining what is rational even though they did not capture all features of epistemic value. This is because, if one thinks that determinations of rationality should not involve comparisons between credal states containing credences on different propositions, then when one uses accuracy to determine what is rational, it does not matter whether or not one's accuracy measure assigns positive value to accurate credences. After all, when we compare the epistemic values of credal states containing

[8] See Pettigrew (2016a) for a defense of this restriction.

credences on all and only the same propositions, the positive epistemic value of additional accurate credences is irrelevant. So, we should not take formal epistemologists' use of Brier scores to determine what is rational as a commitment to a view of epistemic value such as *inaccuracy only*. This is good because philosophers have known that views like *inaccuracy only* are highly implausible since at least William James' time.[9]

That said, some ethicists do adopt views somewhat analogous to *no comparisons* – they argue that certain goods are incommensurable (e.g. Chang 1997). Perhaps, then, I need to say more to show why incommensurability is implausible in the epistemic domain, given its plausibility in the ethical domain. The accuracy-first proofs vindicating coherence rely on seeing the total accuracy of a credal state as the sum of the accuracies of the credences contained in the credal state (e.g. Pettigrew 2016b). So, for any two credences that belong to the same credal state, accuracy-first epistemology sees changes in their accuracies as comparable, and accuracy-first epistemology also sees the utilities of credal states as composed of the utilities of their component credences. If we adopted *no comparisons,* we would be saying that, when we have two credal states containing credences on all and only the same propositions, their accuracies and the accuracies of all of their member credences are straightforwardly comparable, but adding a single new credence to one of these somehow makes comparisons impossible. That is implausible. Further, value comparisons clearly are possible between some credal states containing different

[9] James' *The Will to Believe* discusses how implausible it is to say that the epistemic goal is just avoiding falsehood.

propositions. We see this when we consider *One* and the omniscient credal state.[10] Denying that these two states can be compared is just as implausible as accepting the ERC. So we cannot say that the epistemic utilities of credal states containing credences on different propositions can *never* be compared. At the very least, we can make these comparisons when the states are hugely different from each other in value. Given this, we have no grounds yet to rule out a vast enough repugnant state whose value can be compared to (and is superior to) that of all human wisdom. Conversely, any plausible theory of epistemic value must say that all human wisdom is vastly superior to, and thus comparable to, the repugnant state containing credences about arbitrary space-time locations. So, even if it is plausible that the values of some credal states are incomparable to one another, the states that most clearly illustrate the ERC are not such states. The completely generalized version of *no comparisons* solves the ERC but still gets epistemic value wrong. A more restricted view of epistemic incommensurability will not solve the ERC.

2.2. Other "solutions" to the ERC

Let's consider some other ways of trying to avoid the ERC. I'll only briefly explain why each fails, as these failures are well documented in the literature on the ethical version of the repugnant conclusion. Discussing each failure will help us understand what a solution to the ERC must look like.

One might claim that the utility of a credal state is the average of the accuracies of the credences in that state. This avoids the ERC, since the average accuracy of repugnant states is much lower than that of attractive states. But it runs into the same problem as *no comparisons* and

[10] One might protest that this comparison only makes sense because the omniscient credal state contains *One* in it. Consider the slightly less omniscient credal state which has no credence in the proposition that *One* contains a credence in; this is still clearly superior to *One*.

inaccuracy only: it says that the omniscient credal state is no better than *One*, and that *One* is superior to a credal state containing non-maximal, but extremely accurate, credences in all contingent truths (see Carr 2015 and Pettigrew 2016a).[11]

One might try to address the ERC by saying that only very accurate credences are valuable, and any credence that is less than very accurate is disvaluable.[12] If this were so, then any repugnant state would be disvaluable, since it will consist of just minimally accurate credences. This view of utility has its own problems (as established in e.g. Arrhenius (2000) and Pettigrew (2016a)). To illustrate, let's consider a view that sets the point above which credences have positive value at .9. Consider now some arbitrary credal state *A*. Construct *A** by adding to

---

[11] There are versions of ethical consequentialism that focus on average welfare. However, these tend to say only that we *ought* to maximize average welfare, and don't make the much less plausible claim that a state of affairs is *only better* if it has a higher average welfare than some other state (see, e.g. Rawls 2009, Parfit 1984). To the extent that there is something plausible about seeing the utility of a world-state as the average of the welfare of the people it contains, this plausibility does not extend to the parallel view in epistemology. Whatever plausibility the ethical claim has comes from the intuitive importance of human equality or of just/fair distributions of welfare among people. There is nothing plausibly important about a parallel kind of justice or fairness in distributions of accuracy among propositions (although, as Liam Bright has pointed out to me, there can be cases in which considerations of moral justice seem to have implications for how accuracy should be distributed over credences).

[12] Alternately, we might describe this as raising the neutral point–the point at which credences become accurate–higher, so that the credences in the repugnant credal state no longer count as accurate. The objection I am about to give applies to this description of the "solution" as well.

A a maximally *inaccurate* credence—e.g. credence 1 in a falsehood. Construct *A*** by adding instead some number of .89 credences in truths to *A*. Since credences below .9 are disvaluable, these .89 credences reduce the value of *A*** as compared to *A*. What's more, there is some number of these .89 credences that we can add to *A*** so that *A*** is worse than *A**. But it cannot be better to add a maximally inaccurate credence to a credal state than to add many credences that are very close to the truth.[13] Similar problems will arise for any version of this solution that avoids the ERC. So, raising the level of accuracy at which credences become positively valuable solves the ERC but again gives us a false account of epistemic value.

Another way of solving the ERC is to say that credences have diminishing marginal utility; if this utility diminishes quickly enough, repugnant states might never be superior to intuitively good attractive states.[14] However, this still gives us a false account of epistemic value.

---

[13] I am not sure I endorse this as a universal generalization. I suspect that it may depend on the topic. E.g. it might sometimes be good in some ways to expand the set of questions we have considered, even if we endorse false answers to these questions (this is suggested by ideas in Carr 2015). But we can give versions of the examples I've given using propositions and credences that avoid this possibility.

[14] There's two things that one could have in mind when thinking of the accuracy of credal states as having diminishing marginal utility. One is that the total accuracy of a credal state has diminishing marginal utility: each bit of accuracy we add to the system contributes less and less utility. This is similar to how the marginal utility of money tends to diminish for individuals. This form of diminishing marginal utility won't solve the ERC. This is because we can construct a repugnant and attractive state that have the same total accuracy. If total accuracy had diminishing marginal utility, then these would be equally valuable. If diminishing marginal

To show how, I'll try to give an intuitively-accessible overview of results proved by Arrhenius (2000) in the context of the ethical repugnant conclusion (see also Pettigrew (2016a) for application of these results in the context of the epistemic repugnant conclusion). If credences have diminishing marginal value, then the contribution of each credence to the total utility of the credal state it is in is a function of the (in)accuracy of the credence and the size of the credal state.[15] As credal states get larger, each individual credence makes a smaller contribution to the state's total utility than it would have made in a smaller credal state (holding its accuracy fixed). To see why this is a problem, imagine a credal state C made up of all very accurate credences. Now consider an alternative credal state C^*, which contains the same credences as C plus some additional credences, which are less accurate than those in C but individually do have positive epistemic value. If credences have diminishing marginal utility, there will be some versions of C and C^* such that the total utility of C^* is less than that of C, even though C^* is just C plus

---

utility is to solve the ERC, we need a view that says that spreading the total accuracy of a credal state over more credences means that the accuracy contributes less to the utility of the state than it would if concentrated in fewer credences. This is analogous to the diminishing marginal utility views sometimes offered in response to the ethical repugnant conclusion. This is the sort of view I'll discuss in the main text.

[15] The notion of diminishing marginal utility I am using says that as a credal state gets bigger, each credence in it contributes less to the total utility. One might consider instead a time-relative notion, which says that as a credal state gets bigger, each new credence we add to it contributes less. On such a view, an agent who starts with a repugnant credal state, and then learns all of human wisdom, benefits less than an agent who starts with an empty credal state and then learns all human wisdom. This is clearly false (see Parfit (1984) for a similar point about ethical utility).

additional, valuable, credences. This is because the amount that the accuracy of each credence that is also in *C* contributes to the total utility of *C** will have diminished more than can be made up for by the additional valuable credences. In fact, there will be some *C*** which is *C* plus some *dis*valuable credences, and which has a greater total utility than *C**. This is because, if we add fewer "new" credences to *C*** than we did to *C**, then the contribution of the "original" credences (the ones shared by *C*) will be diminished less in *C*** than in *C**. If we pick our credal states carefully, the disutility of the inaccurate credences added to *C*** will not be worse than the diminishment of utility caused by the additional credences in *C**. So, if credences have diminishing marginal value, then there are cases where adding individually valuable credences to a credal state makes that state worse, and cases where adding disvaluable credences to a credal state is better than adding valuable ones. We thus have to reject the diminishing marginal utility solution to the ERC.

The lessons we learn from considering these "solutions" are that credences should not be disvaluable if they are not inaccurate, that adding good credences to credal states sometimes makes those credal states better, and that good credences do not have diminishing marginal utility. In the next section, we will consider an account of epistemic value that is consistent with these lessons and also avoids the ERC.

2.3. A workable account of epistemic value

It is widely accepted that some things are not particularly worth knowing (see e.g. Goldman 1999, Alston 2005, or Grimm 2009). For example, accurate beliefs about the number of dust motes on the desk in front of me, or about the phone number some stranger's

grandparents had in 1972, seem typically to have extremely little value, if any.[16] This common-sense claim about epistemic value opens the door to a solution to the ERC.

Let's distinguish *interesting* and *boring topics*; credences about such topics are *interesting credences* and *boring credences* respectively. Accurate boring credences have significantly less value than accurate interesting credences. How much less? There are two possibilities that can help us avoid the ERC. First, we could say that accurate boring credences have infinitesimal epistemic value–they have some positive epistemic utility, but so little that no amount of it will ever equal the positive epistemic value of any accurate interesting credence. On this view, inaccurate boring credences would also have infinitesimal epistemic disutility (for a discussion of why, see the Appendix). Alternately, we could say that boring credences have no epistemic utility or disutility whatsoever.[17] Consider now the example I used to motivate the repugnance of the ERC, which compares the utility of a credal state containing all human wisdom with that of a repugnant state containing only minimally accurate credences about whether arbitrary locations in space/time contain a particle. The latter are intuitively boring credences, and this allows us to say that the attractive state in this example is epistemically more valuable than the repugnant one: either the boring credences in the repugnant state have no value, or they have some positive value but their total value can never equal that of a state containing interesting credences. Note that, to avoid

---

[16] While it is widely accepted the accurate credences on such topics have relatively little value, when we focus on single credences, it is hard to be sure just how little that value is. The ERC allows us to narrow the possible values these could have down to just two options.

[17] One argument for the infinitesimal value version of this view is that omniscience is an ideal, which cannot be easily explained if boring credences have no value (see Kvanvig 2008 for discussion of omniscience as an ideal).

the ERC, accurate boring credences cannot have more than infinitesimal value (assuming we reject the "solutions" to the ERC discussed in the previous sections), or else some version of this repugnant state would be superior to all of human wisdom.

An account of epistemic value that says boring credences have no, or infinitesimal, epistemic value gives us a plausible verdict about cases where attractive credal states are made of interesting credences and repugnant states of boring credences. What does this account of value say about comparisons between other sorts of attractive and repugnant states? Consider a case where the "attractive" state has only boring credences, and the repugnant state interesting ones. The repugnant state will be superior. But that seems right: having very accurate views about worthless topics is inferior to having at least some sense of the truth about worthwhile topics. Consider a case where both the attractive and boring state contain only boring credences. If boring credences have no value, then neither state is better, which seems plausible to me: knowing quite a bit about a few topics not worth knowing about does not seem better than knowing just a little about a vast range of such topics. If boring credences have infinitesimal value, then the repugnant state is better, but only infinitesimally, and this does not seem particularly problematic. However, when both the attractive and repugnant state are made up of interesting credences, we get the ERC all over again. We might respond to this by saying that the set of all interesting propositions is limited in size–it is small enough that there cannot be a repugnant state made up of enough interesting credences to make it superior to any clearly very good attractive state. Or we might instead accept this limited version of the ERC, and say that any (finite) good credal state is inferior to some vast set of minimally accurate credences in interesting propositions. This would not be such a bad thing to say, as the cases that make the ERC seem most implausible are those in which the repugnant state consists of boring credences.

If we accept that accurate boring credences have no, or infinitesimal, epistemic value, there are many additional questions that we must eventually answer. As just discussed, we have to determine if the set of interesting credences is large or small, so as to determine whether we can entirely, or just partly, avoid the ERC. We have to figure out what distinguishes interesting from boring topics, and whether there is a sharp or a vague distinction. To do this, we would have to determine if interestingness is "purely" epistemic, or if it has some connection to practical considerations. We would have to figure out if interestingness is contingent or necessary, or if some propositions are necessarily interesting and others only contingently. And, if interestingness is contingent, what does interestingness vary with? I will set these questions aside for the remainder of this paper, as nothing I say in the rest of this paper depends on their answers.

All I need to say going forward is this: if accuracy-first epistemology can succeed at all, it must adopt an account of epistemic value that distinguishes interesting from boring topics, and says credences on boring topics have no, or infinitesimal, epistemic value and disvalue. This view is antecedently plausible, it avoids the clearly implausible versions of the ERC, and it also avoids the problems faced by other alleged solutions to the ERC. To make that last point clear, note that when we rejected other "solutions" to the ERC, we learned that credences should not be disvaluable if they are not inaccurate, that having more good credences should sometimes be better than having fewer, and that good credences do not have diminishing marginal utility. All of these are consistent with the view that boring credences have no, or infinitesimal, epistemic value and disvalue.

For the rest of this paper, I refer to the account of epistemic value I have just sketched as one that *makes the interesting/boring distinction*. Any plausible account of epistemic value that sees value as related to accuracy must make this distinction.

2.4. Why can't we accept the ERC?

In part 3 of this paper, we will see that making the interesting/boring distinction raises problems for the vindication of coherence. Before we go on to see what those problems are, we should make sure we really understand why accuracy-first epistemology has to make the interesting/boring distinction, and why it cannot instead simply bite the bullet and accept that the ERC is true.

Many ethicists do accept that the ethical version of the repugnant conclusion is true. They do so because, while the ethical repugnant conclusion is implausible, there is no theory of ethical value which avoids it without taking on what they see as even more implausible commitments (e.g. Tännsjö 2002, Huemer 2008). These implausible commitments are parallel to those I discuss in sections 2.1 and 2.2. Epistemologists, on the other hand, cannot accept the ERC, since we can avoid both the ERC and these implausible commitments by making the interesting/boring distinction. This is not an option in ethics. The ethical version of the interesting/boring distinctions would be the view that certain human lives are just not very important–that these lives can be maximally good from the perspective of the person living them, but either contribute nothing, or contribute only an infinitesimal amount, to the overall utility of the world. This is untenable. And yet the epistemic version of the claim *is* independently plausible–we have reasons to accept it even before considering the ERC. It also avoids the problematic entailments of other solutions to the ERC. If what we want is a plausible theory of epistemic value, we cannot accept the ERC given that we can make the interesting/boring distinction.[18]

---

[18] Pettigrew (2016a) does advocate biting the bullet on the ERC. His argument for this is parallel to arguments for biting the bullet in ethics: he claims there is no way of avoiding the ERC that

Some may not be convinced. If (as we will see) making the interesting/boring distinction causes problems when we try to use our theory of epistemic value as the basis for a theory of rationality, why isn't it worth it to accept the ERC? To see why it is not, it will be helpful first to consider two projects that one *could* undertake, but that I don't think accuracy-first epistemologists typically mean to be undertaking.

We could invent a measure, and then determine what rules one would have to follow to maximize this measure. For example, we could say *PaintDryTime* is the measure of how many minutes one spends watching paint dry. We could then determine the rules for behavior such that conformity with them maximizes *PaintDryTime*, and call these *PaintDry-rationality*. But showing that conformity with the rules of *PaintDry-rationality* maximized *PaintDryTime* would not vindicate these rules: it would not show that *PaintDry-rationality* is normative in any meaningful way. Another project one could undertake is to consider some set of rules and then find a measure such that conformity with those rules maximized that measure. But finding such a measure would not, in and of itself, vindicate these rules. That is because, for any putative norms we like, we can find *some* measure such that conformity with these rules maximizes that measure (assuming the rules are internally consistent). This is the upshot of the literature on "consequentializing" deontological moral theories. Work in this literature shows that almost any deontological moral theory can be turned into a consequentialist one by giving the right account

---

does not have equally implausible consequences. However, he does not consider the interesting/boring distinction. Because of this, his argument does not show that biting the bullet is the best response to the ERC.

of value (e.g. Oddie & Milne 1991, Drier 1993, Louise 2004, Portmore 2007).[19] The same will be true for non-moral systems of norms. If the fact that conformity with some set of norms maximized some score or another were enough to vindicate those norms, then almost any system of norms could be vindicated.

One project of accuracy-first epistemology is to vindicate criteria of rationality–to show that they are appropriate criteria–by showing how conformity with them relates to the good. The points in the previous paragraph show us that this project requires more than merely showing how conformity with the criteria gives us more of something. Rather, vindication of criteria of rationality requires showing that they give us more of something that is recognizably good. Compare the credal state containing a vast number of barely accurate credences in claims about arbitrary space/time coordinates to the credal state containing all human wisdom. The former has more of something, but not more of anything recognizably good. This is obvious by itself, and bolstered by the independent plausibility of the claim that boring credences are significantly less valuable than interesting credences, even if they are equally accurate. So, our evidence tells us that the epistemic value of accurate credences is tied to the topic the credences are about. The need to avoid the ERC gives us clear constraints on what we can say about this connection. To bite the bullet, and say that the ERC is true, is to give up on the vindication of

[19] There is some debate about whether "consequentializing" is possible for all moral theories (see Brown 2011). However, the features of particular moral theories that make them un-consequentializable are not features that are possessed by any mainstream theories of epistemic rationality.

rationality constraints. If we accept the ERC, we can only vindicate rationality constraints in the same sense that the norms of *PaintDry-rationality* can be vindicated.[20]

3. Boring credences and the vindication of coherence

In the remainder of this paper, we will see that, if accuracy-first epistemology makes the interesting/boring distinction, it cannot vindicate coherence as a universal norm of rationality. *Probabilism* is the view that the credences in a rational credal state must be consistent with the probability axioms. The probability axioms formalize the notion of coherence as applied to credences, and so a vindication of coherence as a constraint on rationality requires an argument for probabilism. We will see that, given the interesting/boring distinction, the vindication of probabilism only works in a limited range of cases. This is a serious issue for two reasons. For one, the vindication of coherence is a core goal for accuracy-first epistemology. Further, many of the accuracy-based arguments used to vindicate other norms of rationality assume that credal states should be probabilistically coherent.[21] If rational agents are not always required to be coherent, then these arguments do not fully vindicate these other norms.

---

[20] Of course, people might find themselves in a situation in which they do care about how many minutes they spend watching paint dry. One can say that *PaintDry-rationality* is vindicated for such situations. Similarly, norms of rationality that do not make the interesting/boring distinction might be vindicated for the very rare situations in which accurate credences are important independent of their topic. However, the mere possibility that we might be in such situations does not vindicate these norms more generally, as we'll see in section 3.

[21] For example, Greaves and Wallace's (2006) vindication of conditionalization makes this assumption, as does Levinstein's (2015) discussion of disagreement. Moss' (2011) arguments for

3.1. Dominance arguments for probabilism

Let's start by seeing how the interesting/boring distinction undermines the vindication of coherence that is standard in the accuracy-first literature. The standard argument for probabilism is to show that non-probabilistic credal states are *accuracy dominated* by probabilistic ones. This means that, for any incoherent credal state, there is at least one coherent credal state, which has credences on all and only the same propositions, that is more accurate at every possible world; and, for any coherent credal state, there is no alternative state (which has credences on all and only the same propositions) that is at least as accurate in all possible worlds and more accurate in some (Joyce 1998, Pettigrew 2016b). If arguments along these lines are to vindicate probabilism, what really needs to be the case is that any incoherent credal state is *utility dominated* by some coherent state–it needs to be the case that, for any incoherent credal state, there is an alternative coherent state that is *better*, and not just more accurate, no matter what.

---

certain norms governing epistemic compromises seem to involve this assumption as well (see, e.g. Moss' footnote 17).

The results I discuss below also raise problems for some arguments that do not rely on agents having coherent credences. For example, DeBona and Staffel (forthcoming) discuss agents who start out incoherent, and show how these agents can make their credal state more accurate by reducing (even if not eliminating) incoherence in particular ways. My results show that their arguments cannot be universally applied: their results require that completely eliminating incoherence in the way they describe always improves the credal state, and my results entail that this is not always true.

Accounts of epistemic value which make the interesting/boring distinction make this impossible.[22]

It will be helpful to start with an example. I will return to this example throughout the rest of the paper to illustrate various arguments. My example involves credal states which assign credences to only propositions A and B. A and B are contingent. B is a disjunction of A and some other proposition; we'll ignore this other proposition to keep things simple. *Inc* is a credal state which assigns .7 to A and .4 to B, so it is incoherent. Let's contrast *Inc* with some particular coherent credal state *Coh*. I'll pick a *Coh* that would dominate *Inc* according to standard accounts of accuracy-first epistemology. Let's say *Coh* assigns .55 to A and .55 to B.[23] Assume now that

---

[22] In their discussions of the ERC, Jennifer Carr (2015) and Richard Pettigrew (2016a) also discuss a problem for dominance arguments for probabilism. If, when we ask whether a given credal state is rational, we compare it to alternative states containing credences in different propositions, then any non-omniscient state will be utility dominated by some larger state. This is incompatible with standard dominance arguments for probabilism, which require rational states to be undominated. In response, Pettigrew argues that determinations of whether a credal state is rational should not involve comparisons to credal states containing credences in different propositions, even if the utilities of these states are comparable. I am not sure I agree with that, but my arguments to follow are compatible with his claims. All of the utility comparisons I discuss in the rest of this paper are between credal states with credences on all and only the same propositions.

[23] Here is why this particular *Coh* dominates this particular *Inc*. It is standard to calculate accuracies using Brier scores; lower is better. To calculate the Brier score for a particular credence in a proposition at a world, we take the credence and subtract the truth value of the

there are possible worlds in which A is interesting and true while B is boring. In such worlds, *Inc* is better than *Coh*. Its credence in A is .15 closer to the truth. *Coh* would have a more accurate credence in B, but in worlds where B is boring, this will be almost worthless, and cannot make up for *Inc*'s superior credence in A. So, *Inc* is not dominated by *Coh*.

We can generalize this, and show that there are sets of propositions such that no incoherent assignment of credences to them will be utility dominated by any coherent assignment of credences to them. I prove this in the Appendix, and I'll just sketch the proof here. Take any incoherent assignment of credences to two contingent propositions A and B, picking an A and B such that A can be interesting in worlds where B is boring. Compare this incoherent state to any coherent assignment of credences to A and B. Either the incoherent and coherent states assign the same credence to A, or different credences. Assume that the incoherent and coherent credal states assign different credences to A. Since A is contingent, there is a world in which the incoherent state is more accurate with regards to A. If A interesting in that world and B boring, then the coherent state cannot dominate the incoherent one–the superior value of the coherent credence in B in this world cannot make up the additional value the incoherent state gets from its credence in A. What if the incoherent and coherent state assign the same credence to A? Then they must assign different credences to B. If B is possibly interesting, or if boring credences have infinitesimal epistemic value, then the incoherent state is better than the coherent in some world– since B is contingent, the incoherent state will assign a more accurate credence to B in some

---

proposition at the world (1 for true propositions, 0 for false), and square the result. The score for a credal state at a world is the sum of the scores for all the credences it contains. In worlds where A and B are both true, *Inc* scores .45 and *Coh* .405. In worlds where A is false and B true, *Inc* scores .85 and *Coh* .505. In worlds where both A and B are false, *Inc* scores .65 and *Coh* .605.

world. If B is necessarily boring, and boring credences have no epistemic value, then the coherent state is never better than the incoherent one. Either way, the incoherent state is not utility dominated by the coherent one. So, all we need to generate counterexamples to the dominance argument is that there are contingent propositions which are logically related (so credences in them can be incoherent), such that a credence in one can be interesting in worlds where a credence in the other is boring. I prove in the Appendix that there must be propositions like this. In fact, we should expect these sorts of problem cases to be relatively common, since interesting truths seem quite often logically related to potentially boring truths.[24]

Thus, if accuracy-first epistemology adopts an account of epistemic value which makes the interesting/boring distinction, then it cannot use standard dominance arguments to fully vindicate probabilism. We can use dominance arguments to vindicate coherence requirements on sets of propositions which are all necessarily interesting, or (if boring credences have infinitesimal but non-zero value) all necessarily boring. But these sets are rare, and we would like to vindicate a universal requirement of coherence–one that applies to all credences of all agents at all times. Let's consider other ways of trying to do so.

3.2. Expected utility arguments

One potential response to the argument I've just made is to "downplay" the worlds that undermine arguments for probabilism. If it is sufficiently unlikely that B is boring, so worlds where *Inc* does better than *Coh* were also unlikely, then the expected benefits of *Coh* in other,

[24] Note that nothing about this proof relies on comparing incoherent credal states to coherent ones. So, given the interesting/boring distinction, there will be pairs of propositions such that no assignment of credences to them can dominate any other assignment. This is important if one wanted to use dominance arguments to vindicate norms not involving coherence.

more likely, worlds might outweigh the expected benefits of *Inc* in the worlds where *Inc* does better. Another response is to appeal to the instrumental value of accurate boring credences. Accurate boring credences might be more likely than inaccurate ones to lead to accurate interesting credences. If so, one might say that *Inc* is intrinsically better than *Coh* in some worlds, but the superior instrumental value of *Coh* makes up for this.[25] These responses both appeal to expected value. It is transparent that the first does, but more explanation is needed about the second. We know that not every accurate boring credence will help us form accurate interesting credences. If the chance that some credence will be helpful in such a way is vanishingly small, then it does not make any sense to prefer it to a more intrinsically valuable credence. So appeal to the instrumental value of accurate boring credences must really be appeal to their expected instrumental value.

It turns out that we cannot appeal to expected value to solve the problems I raised for dominance arguments. To calculate the relevant expected values, we need to appeal to the probability that some proposition is interesting, or that it has instrumental value. What probabilities do we use?[26] We could use the probabilities encoded in the credences that the agent in question actually assigns. Some agents will assign these probabilities in wildly incoherent ways, and in these cases we'd be forced to endorse whatever bizarre claims about what is rational that

---

[25] This idea is inspired by ideas in Craig, 1990, Grimm 2009, and Clifford's "The Ethics of Belief."

[26] Objective chances won't always help. Even if there are objective chances that we are in worlds where the proposition in question is interesting or useful, these will sometimes be far too low to make coherence expectedly better than incoherence.

comes out of these.[27] That won't do. So, we have to appeal to the probabilities or credences that would be rational for the agent to assign. Presumably, these hypothetical credences should be coherent. But we run into problems explaining why they should be coherent, because many of them are boring, and so we cannot use dominance arguments to explain the need for coherence.

To see why propositions about epistemic value must often be boring, consider a repugnant credal state: e.g. a state containing a vast number of minimally accurate credences about whether particles occupy arbitrary locations in space and time. Now consider a second credal state which just consists of minimally accurate credences about the epistemic value (either intrinsic or instrumental) of the propositions in the first repugnant state. The second state is just as repugnant: if the first state is not epistemically superior to all human wisdom, neither is the second. So, to avoid the ERC, we have to say that the credences in this second state are boring.

If we want to appeal to expected utility to explain coherence requirements on possibly boring credences, we have to appeal to the credences an agent rationally could have about the value of these boring credences. Since these credences about value are potentially boring, to explain the rationality constraints on them, we have to appeal to the expected value of coherence with regard to them. This requires appeal to the credences an agent could rationally have about the value of these credences about value, and so forth. Thus, trying to vindicate rationality constraints on possibly boring credences by appeal to expected utility either licenses incoherence, or leads to infinite regress.

---

[27] Some think that expected utilities of incoherent credences are undefined; if so, then so much the worse for this attempt to vindicate coherence.

3.3. Indirect consequentialism

If accuracy-first epistemology can vindicate probabilism, it has to show that coherence is beneficial. The dominance arguments were supposed to show, roughly, that incoherent credal states could be improved no matter what by making them coherent. But we've seen that this is not always the case. Using the same sorts of examples, we can show that, for certain sets of propositions, any coherent assignment of credences to those propositions will be worse in some world than some incoherent assignment, and sometimes that world will be the actual one. So we cannot universally vindicate probabilism by appeal to the actual benefits of coherent credal states. We also cannot vindicate probabilism by appeal to the expected benefits of coherent credal states, as we've just seen. Thus, universal vindications of probabilism cannot appeal to the guaranteed benefits of coherence, nor its actual benefits, nor its expected benefits. To see what options accuracy-first epistemology might have left, let's consider how ethical consequentialists have dealt with parallel problems.

Intuitively, we should keep our promises, even in cases where keeping one's promise has a somewhat lower actual utility, or lower expected utility, than breaking it. We clearly cannot vindicate a promise-keeping norm by appeal to the actual or expected utility of promise keeping in such cases, just as we cannot always vindicate probabilism by appeal to the actual or expected utility of particular coherent credal states. To vindicate a promise-keeping norm, many ethical consequentialists appeal to indirect forms of consequentialism, such as rule consequentialism. Direct consequentialism explains what action one ought to take by pointing to the consequences of that action; this is analogous to accuracy-first epistemology's explanation of what credal states are rational by pointing to the utility of those states. Indirect consequentialism in ethics says that what we ought to do is based not on the consequences of actions, but on the consequences of some other thing relevant to action. *Rule consequentialism* is a version of indirect consequentialism

which says, roughly, that one ought to do what the utility maximizing rule says to do.[28] We need to be careful in formulating this view. In some sense, the best rule to follow is a rule which mimics the advice of direct consequentialism and says, "Maximize (expected) value." Of course, indirect consequentialists do not want to say that this is the rule we should follow, since this would reduce their view to direct consequentialism, and would not vindicate a promise keeping norm. To avoid this, indirect ethical consequentialists say that we can't reliably follow a rule which says to always maximize utility. If we tried, we'd make mistakes, and the costs of those mistakes would (it is claimed) outweigh the benefits of cases in which we correctly followed this rule. Because of this, rule consequentialists will say that we ought to do the act endorsed by the *best rule to adopt*, or the *best rule to try to follow* (there are a number of variants of this which won't matter for our purposes), which is the rule that we maximize utility by adopting or trying to follow. Such a rule might tell us to generally keep our promises, because in the long run we'll do better by doing so than we would by trying to do what directly maximized utility.

Indirect epistemic consequentialists might try a similar move. *Epistemic rule consequentialism* would say that a credal state is rational when endorsed by the best rule to adopt or to try to follow. How might this help vindicate coherence norms? One could argue that, were we to try to follow rules which treat potentially boring credences differently from interesting ones, we would do worse than we would if we adopted a rule that tells us to treat credences on all topics as if they were interesting. To motivate this, one might point out that we can be deceived about which credences are boring, and so we can make costly mistakes by trying to discriminate between interesting and boring propositions. For example, when considering credal states like

[28] There are other forms of indirect consequentialism, but my arguments below should apply to them as well.

Inc and *Coh* (from section 3.1), we might be deceived into thinking one of the relevant credences is boring, and so incorrectly think the incoherent credal state *Inc* is better than the coherent state *Coh*. If we were wrong about boringness often enough, the costs of incoherent credences when we are wrong about boringness might outweigh the advantages of incoherent credences when we are right about boringness. If so, the best epistemic rule to adopt, or to try to follow, would tell us to treat all credences as if they were interesting. This, combined with dominance arguments, would allow us to vindicate probabilism. This is just one way to try to get universal coherence requirements out of rule consequentialism, but I'll use it as an example in what follows since it fits so well with currently standard arguments for probabilism. The issues I discuss, however, should be applicable to any attempt to vindicate universal coherence norms via epistemic rule consequentialism.

Which rule is the best to adopt is a contingent fact that will vary from agent to agent and environment to environment. This means that epistemic rule consequentialism cannot give us a universal vindication of probabilism, or of any rationality norms. To see why, let's consider the rule that tells us to treat all propositions as if they are interesting. If this were a good rule to adopt, it would be because we can make mistakes when we judge beliefs to be boring. Let's imagine that we are very poor detectors of interesting beliefs: if a proposition is interesting, we are 99% likely to think that it is boring. Even so, if interesting beliefs are extremely rare in our environment–let's say only 1 in 10,000 propositions we will consider is actually interesting–then whenever we identify a proposition as boring we are almost definitely correct (we are about 99% likely to be correct). In environments like this, when we judge that one of the credences in *Inc* and *Coh* is boring, we are almost always correct. This means that our judgments that incoherent *Inc* is superior to coherent *Coh* are potentially extremely reliable. So, we could do better adopting a rule allowing us to treat some credences as boring than we would adopting a rule telling us to

treat all propositions as interesting: were we to adopt the former rule, the costs of the cases in which we are mistaken about value could be outweighed by the benefits of forming incoherent credences in the large number of cases in which we are correct. Agents with different capacities would fare differently in this environment, and agents with these same capacities would do poorly in other environments. An agent might know enough about her capacities and environment to know how well she is likely to do adopting one rule versus another. It would be inappropriate to evaluate her rationality by applying rules which she knows are not the best for her to adopt in her environment. Thus, there are some agents who should not adopt rules telling them to treat all propositions as interesting. More generally, we cannot use rule consequentialism to universally rescue the dominance arguments for probabilism.

We might try to reduce the contingency of rule consequentialism by saying that we ought to evaluate agents in light of the rule that would be best for *ideal* agents to try to follow in the environment the actual agent is in. This will not get us a universal vindication of probabilism. We should expect ideal agents to be extremely good, in at least some environments, at determining which propositions are interesting or boring.[29] If the agent is in an environment that

[29] To see why I say this, imagine that, instead, ideal agents were never all that good at determining which propositions are interesting and which are boring. There is some rule *I* that would be best for the ideal agent to follow, given their limitations. There presumably can be a agent who *is* quite good at determining which propositions are interesting, in at least some environments, and can know they are good at it in these environments. Consider such an agent who will do epistemically better following a different rule than the ideal agent would do by following *I*. It's hard to see why this putatively non-ideal agent should follow *I*. So, if it is at all

is conducive to distinguishing interesting from boring propositions, then ideal agents will almost never be wrong about which propositions are interesting. But then ideal agents should not always follow a rule that treats all propositions as if they are interesting, and so we cannot use dominance arguments to show that ideal agents should always have probabilistic credences.

If we think of accuracy-first epistemology as a direct form of epistemic consequentialism, it cannot vindicate universal coherence norms. If accuracy-first epistemology were instead an indirect form of epistemic consequentialism, it might be able to show that certain agents should always be coherent in certain environments. But indirect forms of epistemic consequentialism cannot show that coherence requirements apply to all credences of all agents in all situations.

4. Conclusion

If epistemic utility does not measure something good, then accuracy-first epistemology cannot vindicate any norms of rationality at all. The epistemic repugnant conclusion shows us that our account of epistemic utility only measures something good if it makes the interesting/boring distinction. This is because accounts of epistemic utility which make the interesting/boring distinction are independently plausible, avoid the ERC, and avoid the problems faced by other accounts of value which also avoid the ERC. However, accounts of epistemic utility which make the interesting/boring distinction are incompatible with the ways of vindicating universal coherence requirements of rationality that accuracy-first epistemologists currently know of (although we can still use dominance arguments to show that sets of credences which are necessarily interesting should be coherent). If accuracy-first epistemology cannot universally vindicate coherence requirements, this poses problems for attempted vindications of

---

plausible that everyone should follow the same rule as the ideal agent, then the ideal agent must be good at identifying which propositions are interesting and which are boring.

many other norms of rationality, as these other vindications start by assuming that credal states should be coherent. Accuracy-first epistemology must either largely give up on coherence, and thus on many other putative norms of rationality, or find a new route to its vindication.

Acknowledgments

This paper would not exist without Julia Staffel's advice and assistance. To the extent that I have explained my ideas well, this is also due to help from Jacob Ross. I benefitted from excellent feedback from anonymous reviewers, and from help from the following people: Seamus Bradley, Liam Bright, Kenny Easwaren, Colin Elliot, Branden Fitelson, Simon Goldstein, Chris Heathwood, Eric Hochstein, Ben Levinstein, Kathryn Lindeman, Kris McDaniel, Alastair Norcross, Tyke Nunez, Graham Oddie, Richard Pettigrew, Joe Salerno, and Vahé Shirikjian. Finally, I would also like to thank the students in my Epistemic Value seminar, and audiences at the Formal Epistemology Workshop, St. Louis University, the University of Colorado, Boulder, and Washington University in St. Louis.

References

Alston, W. (2005) Beyond Justification: Dimensions of Epistemic Evaluation. Ithaca: Cornell University Press.

Arrhenius, G. (2000) An impossibility theorem for welfarist axiology. *Economics and Philosophy*, 16 247-266.

Brown, C. (2011) Consequentialize this. Ethics, 121 (4). pp. 749-771.

Chang, R. (1997). Incommensurability, incomparability, and practical reason. Harvard University Press.

Carr, J. (2015) Epistemic expansions. Res Philosophica 92 (2):217-236 (2015)

Craig, E. (1990) Knowledge and the State of Nature. Oxford: Clarendon Press.

Dreier, J. (1993). Structures of Normative Theories. The Monist 76: 22-40.

Easwaran, K. (2015). Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities. Noûs.

Goldman, A. (1999) Knowledge in a Social World. New York: Oxford University Press.

Greaves, H., & Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. Mind, 115(459), 607-632.

Huemer, M. (2008). In defence of repugnance. Mind, 117(468), 899-933.

Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. Philosophy of Science, 575-603.

Kvanvig, J. (2008). Pointless truth. Midwest Studies in Philosophy, 32(1), 199-212.

Leitgeb, H., Pettigrew, R. (2010) An objective justification of Bayesianism I: Measuring inaccuracy. Philosophy of Science, 77, 201-235.

Levinstein, B.A. (2015) With all due respect: The macro-epistemology of disagreement. Philosopher's Imprint, 15.

Louise, J. (2004). Relativity of Value and the Consequentialist Umbrella. Philosophical Quarterly 54: 518-536.

Moss, S. (2011). Scoring rules and epistemic compromise. Mind, 120(480), 1053-1069.

Oddie, G. (1997). Conditionalization, cogency, and cognitive value. British Journal of the Philosophy of Science, 48, 533–541.

Oddie, G., & Milne, P. (1991). Act and value: Expectation and the representability of moral theories. *Theoria*, *57*(1-2), 42-76.

Parfit, D. (1984) Reasons and Persons. Oxford University Press.

Pettigrew, R. (2016a) The population ethics of belief: In search of an epistemic Theory X. Nous.

Pettigrew, R. (2016b) Accuracy and the Laws of Credence. Oxford University Press.

Portmore, D. W. (2007). Consequentializing moral theories. Pacific Philosophical Quarterly, 88(1), 39-73.

Rawls, J. (2009). A theory of justice. Harvard university press.

Sosa, E. (2007). A Virtue Epistemology: Apt Belief and Reflective Knowledge. New York: Oxford University Press.

Staffel, J. (2014). Disagreement and Epistemic Utility-Based Compromise. Journal of Philosophical Logic, 1-14.

DeBona, G., Staffel, J. (forthcoming). Graded incoherence for accuracy-firsters, Philosophy of Science.

Tännsjö, Torbjörn 2002: 'Why We Ought to Accept the Repugnant Conclusion'. Utilitas, 14, 339-359.

Appendix

In this Appendix, I prove that the interesting/boring distinction causes problems for dominance arguments for probabilism. I will first show how, given an assumption about the interestingness of a pair of propositions, there is a counterexample to the dominance arguments for probabilism: there is an incoherent assignment of credences to this pair that cannot be dominated by any coherent assignment. I will then argue that fairly ordinary pairs of propositions satisfy this assumption, so counterexamples to the dominance arguments for probabilism are not limited to exotic cases.

My initial counterexample to the dominance arguments will be an incoherent credal state *Inc* that contains credences in just two propositions, A and B. A is some contingent atomic proposition. B is a contingent disjunction of A and some other proposition. Let's say that *Inc* assigns a credence of .7 to A and .4 to B. Let's assume the following: B is boring in some possible worlds where A is interesting and true, and also boring in some possible worlds where A is interesting and false. My argument to follow needs this assumption, and I'll show later that it will be satisfied in a range of cases. I'll briefly illustrate what these propositions might look like. A might be the proposition that there is intelligent life elsewhere in the universe. B might be the disjunction of this proposition and some claim about whether an arbitrary location in space-time has a particle in it. I suspect that these propositions intuitively fit the description I've given. It seems valuable to have accurate credences about A whether or not there is intelligent life. But, it does not seem valuable to have an accurate credence in the claim, "There is intelligent life or location such and such has a particle in it." For those not satisfied by this example, I will prove later that some examples satisfying my assumption must exist.

To show that *Inc* is a counterexample to the dominance arguments, we have to show that there is no coherent credal state *Coh* containing just credences in A and B that dominates *Inc*. To show this, I will divide all coherent credal states into those that have the same credence in A as *Inc* (or a credence in A that is only infinitesimally different from *Inc*'s), and those that have different credences in A from *Inc*, and address each separately.

A coherent credal state *Coh* cannot dominate *Inc* if it assigns a different credence to A than *Inc* does. Let's say *Coh* assigns a higher credence to A than does *Inc*, such as .9. We have assumed that there will be worlds in which B is boring and A is interesting and false. In such worlds, *Inc* is superior to *Coh*. This is because *Inc*'s credence in A is (non-infinitesimally) closer to the truth, and, since B is boring, the accuracy of *Coh*'s credence in B cannot make up for this;

even if boring credences have some utility, infinitesimal utilities cannot outweigh non-infinitesimal utilities. So *Coh* does not dominate *Inc* if *Coh* assigns a higher credence to A. Consider instead a *Coh* which assigns a lower credence to A than does *Inc*; imagine that it assigns .6 to A. We are assuming that there are worlds in which B boring and A is interesting and true. *Inc*'s credence in A is superior to *Coh* in such worlds, so *Coh* does not dominate *Inc*. So there is no coherent state that assigns a different credence to A and dominates *Inc*.

What about coherent states that assign the same credence to A as *Inc*, or assign credences only infinitesimally different? If the accuracy of boring credences has infinitesimal utility, it is easy to show that such a *Coh* cannot dominate *Inc*. Since *Coh* assigns the same credence to A as does *Inc*, it must assign a different credence to B. So, in some possible world, *Inc* assigns a more accurate credence to B than *Coh*, and is thus superior whether or not B is boring in those worlds, since accurate boring credences still have some utility. So, if accurate boring credences have some utility, no coherent credal state *Coh* can dominate *Inc*.

Let's consider how the argument goes if accurate boring credences have no utility. As long as there are some possible worlds where B is true and interesting, and some possible worlds where B is false and interesting, we can give a similar argument to the one I just gave, since B's accuracy has utility in these worlds. If this is so, then (as above) there will be some world where accurate credences in B have epistemic value, and *Inc* assigns a more accurate credence to B than does *Coh*, and *Inc* and *Coh* assign the same credence to A. But, if accurate boring credences have no utility, we have to consider three more options: accurate credences in B have no possible epistemic value (B is boring in all possible worlds), B is only interesting in worlds where it is true, or B is only interesting in worlds where it is false. If any of these is true, it is not guaranteed that there will be worlds where *Inc* assigns a more accurate, and valuable, credence to B than *Coh*. But these options also break the dominance arguments for probabilism. If accurate credences in B

have no possible epistemic value, there is no possible world in which *Coh* is better then *Inc*, since they assign the same credence to A and credences in B never have epistemic utility. If, on the other hand, accurate credences in B only have epistemic value when B is true, then some version of *Coh* could dominate *Inc*. This *Coh* would assign a higher credence to B than does *Inc*, so it would be better in worlds where B is true, and no worse in worlds where B is false. But this option raises a new sort of problem for the dominance argument for probabilism. It would mean that we are required to assign credence 1 to B. This would be because any *Coh* that did not assign 1 to B would be dominated by a credal state that assigned the same credence to A as *Coh* but a higher credence to B. This is untenable, so we can exclude this option from consideration. We exclude the option that B is only interesting when false for similar reasons. Thus, if accurate boring credences have no epistemic utility, no coherent state *Coh* can dominate *Inc*.

The arguments I have made assume that there are possible worlds in which A is interesting and true and B is boring, and possible worlds in which A is interesting and false and B is boring. Some A and B meet these conditions. To see why, let's start by seeing why there must be propositions like A which can be interesting whether or not they are true. Intuitively, there are propositions like this. And, if boring credences have zero value, then the dominance arguments for probabilism requires that there are propositions like this: if, for example, A was only interesting when true, then any coherent state that doesn't assign 1 to A would be dominated by an incoherent state that assigns the same credence to B and a higher credence to A. If boring credences have infinitesimal value, the argument that some propositions must be interesting whether or not true is a bit more complex. Assume that all possibly interesting propositions are only interesting when true. Consider a question that has multiple possible but exclusive answers, where at least two of these answers–propositions X and Y–are possibly

interesting.[30] Consider a credal state *Max* that assigns 1 to X and 1 to Y. We can show that *Max* cannot be utility dominated, even though it assigns 1 to two propositions that cannot both be true. All alternative credal states must assign less than 1 to X or less than 1 to Y. If an alternative assigns less than 1 to X, *Max* is superior to it in worlds where X is true and Y false: *Max* is non-infinitesimally better with regards to X, and any benefit the alternative gets from superior accuracy with regards to Y is infinitesimal. If an alternative assigns less than 1 to Y, it will be worse than *Max* for similar reasons in worlds where Y is true and X false. So *Max* cannot be dominated. This is a bad result. To avoid it, we can either deny that any questions have more than one possibly interesting answer, or accept that credences about some propositions are interesting whether or not the propositions are true. Some questions do have more than possibly interesting answer. So, there are propositions like A which can be interesting when true and when false.

Why must there be propositions like B, whose boringness is to some extent independent of their truth or the truth of A? The disjunction of an interesting proposition and a boring proposition cannot always be interesting. Why? To avoid the epistemic repugnant conclusion, we have to say that a credal state *Rep*, containing vast numbers of minimally accurate boring credences in true propositions, has no (or infinitesimal) value. Assume for *reductio* that the disjunction of an interesting proposition and a boring one is always interesting. Consider a credal state *Rep**; for each proposition that *Rep* has a credence in, *Rep** contains an equally accurate credence in the disjunction of that proposition and some interesting proposition P (each

[30] Y cannot be just the negation of X, but rather must be a positive answer to the question. That's because we are positing that X is boring if false, and so the proposition ¬X is not possibly interesting.

disjunction in *Rep** is about the same interesting proposition *P*). The result would contain a vast number of minimally accurate credences whether or not P is true (by stipulation the credences in *Rep* previously were in truths, so the disjunctions are all true, and the credences in them are no lower than the credences in the initial, boring propositions *Rep* contained). If disjunctions of interesting and boring credences are interesting, then we'd get the epistemic repugnant conclusion. Let's illustrate. Say that *Rep* consists just of credences about whether random locations in space/time have particles in them, and *Rep** just has credences about disjunctions of each of these and the proposition that intelligent life exists elsewhere in the universe. If all the credences in *Rep** were minimally accurate, but this told us nothing about intelligent life, *Rep** would not be more valuable than, e.g. all human wisdom. So these disjunctions must be boring. The same sort of argument shows that the disjunction of an interesting proposition and a boring one must sometimes be boring when the disjunction is false. Imagine that *Rep* contained minimally accurate credences in boring falsehoods, and *Rep** contained a minimally accurate credences in the disjunction of some interesting falsehood Q and every proposition in *Rep* previously contained; again, the credences in *Rep** cannot all be interesting. So, to avoid the epistemic repugnant conclusion, we have to say that the disjunction of an interesting and boring proposition cannot always be interesting, and should usually be boring; this is independent of the whether the interesting proposition is true or false, and whether the disjunction is true or false.

So, there are at least some incoherent credal states like *Inc* that cannot be dominated by coherent ones; it seems that there must be as many such pairs as there are interesting propositions and disjunctions of them and boring propositions. Note that, while I picked specific credences for *Inc* to assign to A and B, nothing in my argument depended on these particular numbers. As long as my assumption about A and B's possible interestingness holds, no assignment of credences to A and B can be dominated.

While the example I've been discussing is one in which B is a disjunctive proposition, the problem for the dominance arguments is more general. My initial counterexample to the dominance argument really only requires the following: propositions A and B are contingent, credences in A and B can possibly be incoherent, B is boring in some possible worlds where A is interesting and true, and B is boring in some possible worlds where A is interesting and false. We could have, for example, made B a conjunction of A and another proposition. To see why, start with a repugnant credal state whose credences are all about whether arbitrary space/time coordinates are occupied by particles. Take some contingent proposition X, which is interesting whether or not true. Now construct a credal state containing nothing but minimally accurate credences in conjunctions of X with each of the propositions from the repugnant state. This state will also be repugnant–it is clearly not better than all human wisdom, no matter how large it is–whether or not these minimally accurate credences are about true or false conjunctions, and whether or not X is true. So, these conjunctions must be boring whether or not X is true.[31] And so we can make counterexamples to dominance arguments that involve conjunctions. More generally, for incoherent credences in any pair of propositions, if one can be interesting whether or not it is true, and the other can be boring whether or not it is true, and the boringness of the

[31] One might be tempted to say that these propositions are not boring because they have non-infinitesimal value. After all, they might tell us something about X; since X is interesting, information about X might seem to have non-infinitesimal value. If it does, this value is instrumental, and not the kind relevant to the interesting/boring distinction. That can easily be seen by imagining the credal state described here also contained a coherent credence in X; to see these conjunctive propositions as having non-infinitesimal value in this case would be inappropriate double-counting.

second is to some extent independent of the truth of the first, then we cannot use dominance arguments to show that credences in the two should be coherent.