

Downwards propriety in epistemic utility theory

Alejandro Pérez Carballo

University of Massachusetts, Amherst

To a first approximation, *epistemic utility theory* is an application of standard decision theoretic tools to the study of epistemic rationality. The strategy consists in identifying a particular class of decision problems—*epistemic* decision problems—and using the recommendations that our decision theory makes for them in order to motivate principles of epistemic rationality.

The resulting principles will of course be a function of, among other things, what we take epistemic decision problems to be and of what specific brand of decision theory we rely on.¹ But regardless of the details, epistemic utility theory inherits from the decision theoretic framework a distinction between axiological notions—of *epistemic value* or *epistemic utility*—and deontological notions—like *epistemic rationality* or *epistemic permissibility*.

From a purely formal point of view, there is no need to take a stand on which, if any, of the two families of notions is prior to the other. But proponents of epistemic utility theory typically adopt the further commitment that the axiological is prior to the deontological—that the epistemic good is prior to the epistemic right—where ‘priority’ is *justificatory* priority. Indeed, many proponents and critics alike seem to agree that the project of applying decision-theoretic tools to the study of epistemic rationality loses its point if it turns out that an axiology-first approach to epistemology cannot be made to work.²

I think epistemic utility theory, as a formal framework for clarifying and assessing our epistemic commitments, has much to recommend it. And I think this is so even if it turns out that the epistemic right is prior to the good, or if neither of the two is prior to the other. But my goal here is not to argue for that. Rather, I want to argue from within the framework of epistemic utility theory against an axiology-first approach to epistemology. Whether this casts doubt on the merits of the epistemic utility framework is a question I must leave for some other time.

¹ For a clear and careful discussion of some of the options, see [Greaves 2013](#).

² See e.g. [Caie 2013](#), [Greaves 2013](#), [Konek & Levinstein 2017](#), [Pettigrew 2016](#). Of course, not all agree that this is an essential component of epistemic utility theory: see e.g. [Horowitz 2018](#), [Meacham 2018](#), [Stalnaker 2002](#) and, on at least one reading, [Joyce 2013](#).

My argument will proceed in two steps. I will first argue that the success of this axiology-first approach to epistemology depends (in part) on the availability of a purely axiological justification of a non-trivial assumption about epistemic value, viz. that epistemic utility functions are what I will call *downwards proper*. Next, I will make a case that no such justification is forthcoming. More specifically, I will argue that, given some of the more widely shared presuppositions among proponents of the axiology-first approach, the assumption of downwards propriety cannot be motivated on purely axiological grounds. Thus, unless we abandon those presuppositions, an axiology-first approach to epistemology is unlikely to succeed.³

1 The framework

To keep things simple, let us stipulate that we are working with a fixed finite set W of possible worlds. A proposition, for our purposes, will just be a subset of W . Given a *partition* \mathcal{S} of W —a set of non-empty propositions that are pairwise disjoint whose union is the entire set W —we say that a *credence function over* \mathcal{S} is any function that assigns a real number in $[0, 1]$ to each member of \mathcal{S} . For a credence function c , I will sometimes use \mathcal{S}_c to denote the partition c is defined over. I will refer to \mathcal{S}_c as the *state space* of c .

(Note that our definition of a credence function is importantly different from the more familiar definition of a credence function as an assignment of numerical values to all members of a given *algebra* of propositions. But here I follow [Joyce 2009](#), and much of the literature, in identifying credence functions instead with assignments of numerical values over a partition.⁴)

An *epistemic utility function* is a function that assigns a real number to each pair consisting of a credence function and a possible world.⁵ I will assume that

- 3 There has been much critical discussion on whether axiology-first epistemology can provide justification of norms other than Probabilism ([Easwaran & Fitelson 2012](#), [Meacham 2018](#)). Most of it has focused on the particular version of the view often called ‘accuracy-first epistemology’—a combination of axiology-first epistemology together with the claim that accuracy is the sole fundamental source of epistemic value. But some of the most influential critical discussion ([Caie 2013](#), [Greaves 2013](#)) has targeted the viability of the axiology-first approach altogether (see also [Carr 2017](#)). But my arguments here will apply, though, even if we grant (as [Konek & Levinstein 2017](#) and [Joyce 2018](#) have argued) that the right formulation of epistemic utility theory ignores, *contra* Greaves and Caie, any dependence relation between epistemic ‘acts’ and states of the world.
- 4 Unless we assume that all credence functions are probabilistically coherent, an assignment of numerical values to a partition does not determine a unique assignment of numerical values to the smallest algebra containing all the members of the partition. But here I’m following most of the literature in assuming that epistemic utility functions are only sensitive to which numerical values a credence function assigns to the *atomic* propositions in the algebra (this assumption is made, as far as I can tell, for reasons of mathematical tractability—cf. [Leitgeb & Pettigrew 2010](#), p. 221f).
- 5 Thus, I’m ruling out at the outset functions, like the so-called log score, which take values in the extended real line $\mathbb{R} \cup \{\infty, -\infty\}$. Nothing in what I will say, however, hinges on this.

epistemic utility functions are *nice* in the following sense: for each credence function c , the function $u(c, \cdot)$ is constant throughout each member of \mathcal{S}_c . In other words, the utility of c at a world depends only on the truth-value of propositions in c 's state-space. If u is nice, then for any partition \mathcal{S} of W we can define a function $u_{\mathcal{S}}$ that assigns a real number to each pair consisting of a credence function over \mathcal{S} and a member of \mathcal{S} , by picking an arbitrary element w_S of each $S \in \mathcal{S}$ and letting

$$u_{\mathcal{S}}(c, S) = u(c, w_S).$$

Niceness ensures that this definition does not depend on our choice of w_S . Slightly abusing notation, I will write $u(c, S)$ rather than $u_{\mathcal{S}}(c, S)$ when the choice of partition is clear from context and c is a credence function over \mathcal{S} .

If u is an epistemic utility function, c a credence function over \mathcal{S} and $S \in \mathcal{S}$, we call $u(c, S) = u_{\mathcal{S}}(c, S)$ the *epistemic utility of c at S* . I will stipulate that epistemic utility functions are continuous⁶ and that they satisfy the following minimal constraint, often called ‘Truth-directedness’: whenever c and c' are defined over the same partition, if c 's assignments are at least as close and sometimes strictly closer to their truth-values if S obtains than those of c' , then the epistemic utility of c at S is strictly greater than that of c' at S .

An example of an epistemic utility function is, of course, the *Brier score*, defined as:

$$\mathbf{b}(c, w) := - \sum_{S \in \mathcal{S}_c} (c(S) - \mathbb{1}\{w \in S\})^2,$$

where we let $\mathbb{1}\{w \in S\}$ equal 1 if $w \in S$ and 0 otherwise. Note that \mathbf{b} is nice, and further that for each partition \mathcal{S} , each c defined over \mathcal{S} , and each $S \in \mathcal{S}$,

$$\mathbf{b}(c, S) = - \sum_{T \in \mathcal{S}} (c(S) - \mathbb{1}\{S = T\})^2,$$

where $\mathbb{1}\{S = T\}$ equals 1 if $S = T$ and 0 otherwise. More generally, for each $\theta \in \mathbb{R}$, the function \mathbf{b}_{θ} , defined by

$$\mathbf{b}_{\theta}(c, w) := \mathbf{b}(c, w) + \theta = \theta - \sum_{S \in \mathcal{S}_c} (c(S) - \mathbb{1}\{w \in S\})^2,$$

will be a nice epistemic utility function satisfying all our assumptions thus far.

A credence function c is *probabilistically coherent* iff $\sum_S c(S) = 1$. If c is a probabilistically coherent credence function over \mathcal{S} and u is an epistemic utility

⁶ For a fixed \mathcal{S} and epistemic utility function u , we can think of $u_{\mathcal{S}}$ as a real-valued function whose domain is $\mathbb{R}^N \times \mathcal{S}$, with $N = |\mathcal{S}|$. Our continuity assumption is essentially the assumption that for each \mathcal{S} and each $S \in \mathcal{S}$, $u_{\mathcal{S}}(\cdot, S)$ is a continuous function from \mathbb{R}^N to \mathbb{R} .

function, then for any credence function c' over \mathcal{S} we define the *expected u -value* of c' relative to c , which we denote with $\mathbb{E}_c[u(c')]$ as follows:

$$\mathbb{E}_c[u(c')] := \sum_{S \in \mathcal{S}_c} c(S)u(c', S).$$

An *epistemic decision problem* over \mathcal{S} is a triple $\mathcal{D} = (c, O, u)$, where c is a credence function over \mathcal{S} , O is a set of credence functions over \mathcal{S} —the set of *available options*—, and u is an epistemic utility. (I will omit the qualifications ‘for \mathcal{S} ’, ‘over \mathcal{S} ’, etc. when it’s clear from context which partition we’re talking about.)

For any two credence functions c_1 and c_2 and any epistemic utility function u , we say that c_1 (*weakly*) *u -dominates* c_2 iff for any $w \in W$, $u(c_1, w) \geq u(c_2, w)$. We say that c_1 *strongly u -dominates* c_2 iff c_1 weakly u -dominates c_2 and for some $w \in W$, $u(c_1, w) > u(c_2, w)$. For a given epistemic decision problem $\mathcal{D} = (c, O, u)$ and $c_1 \in O$, we say that c_1 is *strongly* (resp. *weakly*) *dominated* in \mathcal{D} iff there is some $c_2 \in O$ such that c_2 strongly (resp. weakly) u -dominates c_1 . Finally, if C is coherent, then for any decision problem $\mathcal{D} = (c, O, u)$ and any $c^* \in O$, we say that c^* *maximizes expected value* in \mathcal{D} iff for any $c' \neq c^*$, with $c' \in O$,

$$\mathbb{E}_c[u(c^*)] > \mathbb{E}_c[u(c')].$$

2 Applying the framework

The framework of epistemic utility theory allows us to derive claims about epistemic rationality from claims about epistemic value, at least given some *bridge principles* telling us how axiological and deontological notions relate to one another. For instance, much like in practical decision theory, we could say⁷

DOMINANCE: A credence function c is rationally permissible relative to a decision problem \mathcal{D} only if it is not dominated in \mathcal{D} by a credence function that is itself not dominated in \mathcal{D} .⁸

⁷ It is worth noting that Dominance is bound to be rejected by those who think we should allow for the possibility of *state-act dependence* in the context of epistemic utility theory. Here, though, I will set that possibility aside, and trust I do not thereby beg any questions, since this ultimately only makes things easier for the target of my arguments. See [fn. 3](#) for further discussion.

⁸ Contrast Dominance with the strictly stronger bridge principle that says that a credence function c is permissible relative to \mathcal{D} only if it is not dominated in \mathcal{D} . This latter principle has the unfortunate consequence that when every available option is dominated in \mathcal{D} , no option is rationally permissible. Cf. [Pettigrew 2016](#), § 2.1 for an argument for why this latter principle is too strong, and why we should at least replace it with Dominance (a principle he calls ‘Undominated Dominance’).

By making assumptions about what epistemic utility functions are like, we can derive from principles like this claims about which credence functions are permissible relative to a given decision problem.

Typically, however, the framework is not put to use to establish claims about which credence functions are epistemically rational relative to which decision problem. Rather, it is put to use to establish claims about which credence functions can be epistemically rational for an agent at a time. For instance, one of the main selling points of epistemic utility theory is that it offers a way of vindicating *Probabilism*—the claim that all rationally permissible credence functions (for an agent at a time) are probabilistically coherent—from assumptions about epistemic value.⁹ And the notion of rationality relevant to Probabilism is not obviously relativized to a particular decision problem. So we need some story about how to construct, for a given agent and time, the relevant decision problem—what I will call the *canonical* decision problem for that agent at that time. Only then can we use bridge principles like Dominance to establish claims about which credence functions are rationally permissible for a given agent at a time.

In principle, one could have different views about which is the canonical decision problem for a given agent at a time. As far as I can tell, though, most agree (without explicitly stating) that in assessing the rationality of an agent at a time, it is *the agent's credence function* at that time, and the set of *all credence functions with the same domain*, that figure as the first two elements in the canonical decision problem. Thus, there is broad agreement in the literature on something like the following principle:

FIXED DOMAIN: A credence function is rationally permissible (for a given agent at a time) only if it is permissible relative to some decision problem of the form $\mathcal{D} = (c, O[c], u)$, where c is the agent's credence function at that time, $O[c]$ is the set of all credence functions defined over \mathcal{S}_c and u is an *admissible epistemic utility function*.

Disagreement is largely focused on what counts as an admissible epistemic utility function and on how to define permissibility relative to a decision problem.¹⁰

To get a better sense of how the framework can be put to use, it helps to go through an example, one we will revisit in due course. Start by assuming the following thesis:

⁹ See e.g. [Joyce 1998](#).

¹⁰ Some think that the agent's *epistemic values*, at a time, play a role in determining the relevant decision problem—this seems to be the view implicit in [Greaves 2013](#) (see e.g. §3) and [Moss 2011](#) (see e.g. §1). Others seem to presuppose instead that (objective) facts about epistemic value play a role in determining the canonical decision problem—this seems to be the view implicit in e.g. [Joyce 2009](#) as well as [Leitgeb & Pettigrew 2010](#).

PROPRIETY: If u is an admissible epistemic utility function, then it is *proper* in the sense that for each probability function p and each credence function $c \neq p$ defined over the same partition, the expected u -value of p relative to p is greater or equal than that of c relative to p .

Using Dominance as a bridge principle, we can derive Probabilism by relying on the following mathematical result:¹¹

JOYCE'S THEOREM: Fix $\mathcal{D} = (c^*, O[c^*], u)$, with u proper, and let $c \in O[c^*]$. If c is not probabilistically coherent, then it is dominated in \mathcal{D} by a probabilistically coherent credence function. If c is probabilistically coherent, it is not weakly dominated in \mathcal{D} .

The argument is straightforward: assuming Propriety, Fixed Domain entails that a given credence function c^* will be rational for an agent at a time only if it is permissible relative to some $\mathcal{D} = (c, O[c], u)$, where c is the agent's credence function at that time and u is a proper epistemic utility function. But Joyce's Theorem entails that c^* will be dominated in any such \mathcal{D} by a probabilistically coherent credence function unless c^* itself is probabilistically coherent. Thus, if c^* is not probabilistically coherent, it will be dominated by a credence function that is not itself dominated. So, from Dominance we can conclude that c^* is rational (for an agent at a time) only if it is probabilistically coherent, viz. Probabilism.

Of course, this argument will contribute little to the project of vindicating Probabilism unless Propriety can itself be justified as a constraint on admissible epistemic utility functions. After all, Propriety is no more self-evident than, and so is as much in need of justification as, Probabilism.

What could a justification of Propriety look like? If the goal is to use Propriety to establish Probabilism, it wouldn't be of much help to justify Propriety by in turn appealing to Probabilism. But in principle that leaves us with plenty of options for finding a justification of Propriety. Some of these, however, are incompatible with a view most commonly associated with the epistemic utility framework.

3 Axiology-First Epistemology and Variable Domains

Proponents of epistemic utility theory typically seek more than a justification of Probabilism. They seek a justification of Probabilism (and other epistemic

¹¹ The main theorem in [Joyce 2009](#) actually relies on something weaker than Propriety, viz. the claim that no probabilistic credence function is u -dominated by any credence function with the same domain. For reasons pointed out in [Pettigrew 2016](#), §2.2, though, we would be wise to rely on the stronger constraint. Related results include those in [de Finetti 1970](#), [Joyce 1998](#), [Leitgeb & Pettigrew 2010](#), [Predd et al. 2009](#).

norms) in purely axiological terms—a justification that appeals solely to facts about epistemic value, together perhaps with one or more bridge principles. Those engaged in the project of justifying norms of epistemic rationality in purely axiological terms—the project I will call *axiology first epistemology*—cannot thus rely on arguments for Propriety that start from assumptions about epistemic rationality.¹²

For example, proponents of axiology first epistemology cannot argue for Propriety by appealing to the claim that probabilistic coherence, even if not rationally required, is nonetheless rationally permissible.¹³ Instead, proponents of axiology first epistemology who want to make use of the argument from Propriety to Probabilism need a purely axiological justification of Propriety.

It is an open question whether such a justification is forthcoming.¹⁴ Here, though, I will grant that it is. In other words, I will grant for the sake of argument that Propriety can be justified in purely axiological grounds. For, as I will argue in this section, an axiology first justification of Probabilism requires an axiological justification for something much stronger than Propriety. Let me explain.

3.1 Variable domains

As I emphasized in §2, in order for the epistemic utility framework to be of any use in establishing claims like Probabilism, we need to specify what I called a *canonical decision problem* for a given agent at a time.

And there seems to be broad agreement that something like Fixed Domain is the right way to do so. In particular, there seems to be broad agreement that the available options in the canonical decision problem for an agent at a time are all and only those credence functions with the same domain as the agent's

12 On a familiar telling, consequentialist (or teleological) moral theories hold that the good is 'prior' to the right, where 'priority' is said in a metaphysical tone of voice (cf. e.g. Moore 1912, Ross 1930—for discussion, see Berker 2018). It is tempting thus to use 'Epistemic Consequentialism' or 'Epistemic Teleology' as a label for the view that the *epistemic* good is prior to the epistemic right (cf. Berker 2013a,b). As I'm understanding the project, though, axiology-first epistemology is silent on questions of metaphysical priority. In what is arguably the founding document of accuracy first epistemology (Joyce 1998), the framework of epistemic utility is introduced in order to offer a *justification* of Probabilism. And as far as I can tell, everyone engaged in the project of axiology-first epistemology takes its goal to be that of *vindicating* epistemic norms. (I should note here that Pettigrew (2016) describes the view he labels 'Veritism', a term he uses almost interchangeably with 'accuracy-first epistemology' (see e.g. p. 11), as maintaining that "there is a single fundamental source of value that is relevant to the epistemic evaluation of credences—it is accuracy" (p. 10), while *also* taking Veritism as essentially seeking to provide a "justification for Probabilism" (p. 22) and other epistemic norms (cf. pp. 133, 155) governing credences.)

13 See Joyce 2009, p. 279. For critical discussion, see Hájek 2009, Pettigrew 2016, Weisberg 2015.

14 Richard Pettigrew (2016, ch. 4) has offered what is perhaps the most sophisticated attempt at offering a purely axiological justification of Propriety, but there is some reason for thinking that the assumptions Pettigrew relies on are more controversial than he takes them to be—cf. Levinstein 2017.

credence function at that time. (For brevity, I will say that a credence function is an available option for an agent at a time just in case it is an available option in the canonical decision problem for that agent at that time.)

I think this is a mistake. If all credence functions defined over the same domain as the agent's credence function (at a time) are available options for an agent (at that time), then so are all credence functions whose domain is *more coarse-grained* than that of the agent's credence function at that time.¹⁵ Let me explain.

Fix a particular agent at a time. Following Pettigrew 2016, call the domain of her credence function at that time her *opinion set* (at that time).¹⁶ Say that a proposition is *available* to an agent (at a time) iff that proposition is definable, using standard Boolean operations, in terms of those propositions in the agent's opinion set.

According to Fixed Domain, in order to assess whether her credence function is rationally permissible, we need to consider a decision problem whose available options include all credence functions defined over the agent's opinion set. Why are all of those options relevant?

Presumably, it is because in *some* sense, they are available to the agent at the relevant time. Of course, the relevant sense of availability has little to do with what credence functions the agent is actively entertaining, for no agent remotely like us is able to actively entertain the uncountable many distinct credence functions defined over her opinion set. Similarly, it has little to do with whether the agent is able to *choose* to adopt that credence function as her own. Plausibly, we cannot simply choose to change our epistemic state.

All credence functions defined over her opinion set are relevant, I submit, because in some sense the agent is able to *have* one such credence function. By the agent's own lights, each such credence function has a claim to representing an epistemic state that the agent could be in.

But note that by the same token, so are all credence functions defined over available propositions—call them *available credences*. Suppose for instance our agent's opinion set is given by a partition $\{S_1, S_2, S_3\}$ of W and consider a credence function c whose domain is the partition $\{S_1 \cup S_2, S_3\}$. Why wouldn't that credence function be relevant for the evaluation of the agent's current credence function?

In whatever sense credence functions defined over her opinion state are available options, it seems—in whatever sense she is able to have any such credence function—so are all available credence function. By the agent's own

¹⁵ Exactly what I mean by 'more coarse-grained' here will become clear shortly.

¹⁶ Henceforth, I will stop explicitly relativizing credal attributions to a particular time, and assume the reader can just fill those in as needed.

lights, each available credence function has a claim to representing an epistemic state she could be in. They thus have an equal claim to being available options in the canonical decision problem as all credence functions defined over the agent's opinion set. In other words: to the extent we think all credence functions defined over the agent's opinion state are relevant to the epistemic evaluation of an agent's credence function, we should also think available credence functions are equally relevant.

Note that the same cannot be said of credence functions whose domain includes unavailable propositions. Suppose, for instance, we think of the agent's opinion set as defining all propositions she is able to entertain—perhaps those propositions she can in principle consider, because she has the relevant concepts. Credence functions defined over propositions the agent is unable to entertain are not, in the relevant sense, available to her. After all, she cannot consider that credence function as one she could have, for by assumption she is unable to entertain the relevant propositions. Or suppose instead we think of the agent's opinion set as defining those propositions the agent is currently entertaining. Perhaps she is able to entertain the proposition that there are nowhere differentiable, continuous, real-valued functions, but like most of us in everyday situations, that proposition is not part of her epistemic landscape.¹⁷ Plausibly, credence functions whose domain includes that proposition (or any proposition she is not currently entertaining) are also not available to her at that particular time.

We should reject Fixed Domain. What should we replace it with? In other words, how else should we think of the canonical decision problem for an agent at a time?

3.2 A New Challenge for Axiology-First Epistemology

Fix a partition \mathcal{S} of W . Say that a partition \mathcal{S}^- of W is a *coarsening* of \mathcal{S} iff for each $S^- \in \mathcal{S}^-$ there is $S \in \mathcal{S}$ such that $S^- \subseteq S$. Thus, \mathcal{S}^- is a coarsening of \mathcal{S} iff any member of \mathcal{S}^- is the union of elements of \mathcal{S} . (Equivalently, \mathcal{S}^- is a coarsening of \mathcal{S} iff all elements of \mathcal{S}^- are definable, using standard Boolean operations, in terms of elements of \mathcal{S} . So, a partition is a coarsening of an agent's opinion set iff all of its members are available to the agent.) Our discussion so far suggests the following alternative to Fixed Domain:

DOWNWARDS CLOSED: A credence function is rationally permissible (for a given agent at a time) only if it is permissible relative to a decision problem

¹⁷ This may be because the relevant concepts are available in principle even though in some sense they are not 'active' (on this distinction, see e.g. [Fodor 1975](#), p. 85 and, more recently, [Kemp et al. 2010](#), §11.3), or instead because she simply isn't attending to the relevant propositions (cf. the literature on (un)awareness and related discussion in the literature on epistemic modals, e.g. [Franke & de Jager 2011](#), [Swanson 2006](#), [Yalcin 2007](#)).

$\mathcal{D} = (c, O^\downarrow[c], u)$, where c is the agent's credence function at that time, $O^\downarrow[c]$ is the set of all credence functions whose domain is a coarsening of \mathcal{S}_c and u is an admissible epistemic utility function.

A surprising consequence of replacing Fixed Domain with Downwards Closed, however, is that the argument for Probabilism sketched in §2 breaks down. For consider a probabilistically incoherent credence function c defined over any non-trivial partition \mathcal{S} —that is, a partition containing more than one element. All that Propriety guarantees is that c will be dominated by a probabilistically coherent credence function defined over \mathcal{S} which in turn is not dominated by any credence function defined over \mathcal{S} . But it tells us nothing as to whether the dominating credence function is dominated by a credence function whose domain is a coarsening of \mathcal{S} .

Indeed, turn again to the familiar Brier score \mathbf{b} .¹⁸ As Carr (2015) points out, if we measure epistemic utility using \mathbf{b} , any probabilistic credence function with a non-trivial credence space that assigns non-extreme values to some proposition will be dominated by the unique probability function whose state space is the trivial partition $\{W\}$.¹⁹ And while this might, strictly speaking, allow the argument for Probabilism to go through—since every credence function will be dominated by a non-dominated, coherent credence function—the cost would be too high. For Dominance would also rule out every other coherent credence function from being rational—the sole rational credence function would be the one that assigns full credence to the trivial proposition and is undefined over every other non-empty proposition—which surely would mean the principle is too strong to be of any use.

In order to get an argument for Probabilism that relies on Downwards Closed, then, we need a constraint on epistemic utility functions stronger than Propriety. To see what that constraint has to look like, note first that our definition of expected u -value (§1) can be generalized so that it makes sense to talk of the expected u -value of c' relative to c whenever c is probabilistically coherent and the domain of c' is a coarsening of c . For recall that

$$\mathbb{E}_c[u(c')] := \sum_{S \in \mathcal{S}_c} c(S)u(c', S),$$

¹⁸ See §1 for the definition.

¹⁹ Fix $S \in \mathcal{S}_c$ such that $c(S)$ is strictly between 0 and 1. Note that for any $x \in \{0, 1\}$ and $r \in (0, 1)$, $(r - x)^2 > 0$. Thus, $\mathbf{b}(c, w) < 0$. But if c_\top is the unique probability function defined over $\{W\}$, we have that for all $w \in W$, $\mathbf{b}(c_\top, w) = 0$. Hence c_\top strictly dominates c .

and $u(c', S)$ will be well-defined whenever S is a subset of an element of the state space of c' .²⁰

We can now replace Propriety with the following, stronger assumption:

DOWNWARDS PROPRIETY: If u is an admissible epistemic utility function, then it is *downwards proper* in the sense that for each probability function p and each credence function $c \neq p$ defined over a coarsening of \mathcal{S}_p , the expected u -value of p relative to p is greater or equal than that of c relative to p .

Using Downwards Propriety, Dominance, and Joyce's Theorem, we can now derive Probabilism.²¹

Of course, for this new argument to help justify Probabilism, we need a way to justify Downwards Propriety. So in order for this argument to provide an axiology-first justification of Probabilism, what we need then is a purely axiological justification of Downwards Propriety. It is not enough to justify Propriety, even if that can be done on purely axiological grounds.

4 Against an axiology-first justification of Downwards Propriety

So far, we have a new challenge for axiology-first epistemology—to provide an axiological justification of Downwards Propriety. In this section, I want to argue that this challenge cannot be met.²² My argument will rely on two assumptions widely shared among proponents of axiology-first epistemology.

- 20 Since by assumption u is nice, $u(c', w)$ will be constant throughout any element of $\mathcal{S}_{c'}$, but also throughout any subset of an element of $\mathcal{S}_{c'}$. Since any member of \mathcal{S} is a subset of an element of $\mathcal{S}_{c'}$, we can conclude that $u(c', S)$ is well-defined.
- 21 Strictly speaking, once we assume Dominance we can rely on something slightly weaker than Downwards Propriety. All we need is that for any probability function p and each credence function $c \neq p$ defined over a coarsening of \mathcal{S}_p , p is not dominated by c in any decision problem $\mathcal{D} = (p, O^+[p], u)$, where (as above) $O^+[p]$ is the set of all credence functions whose domain is a coarsening of \mathcal{S}_p and u is an admissible epistemic utility function. But as Pettigrew 2016, §2.2 points out, Dominance is slightly too strong. Relative to a particular decision problem, a credence function that is dominated by an undominated probability function may nonetheless be permissible if the dominating probability function is defective in some other way. And one way for a probability function to be defective, relative to a decision problem, is for it to assign greater epistemic value to some other one of the available options (cf. fn. 11). Once we replace Dominance with the better principle, what we need is precisely Downwards Propriety—nothing weaker will do.
- 22 Above (fn. 14), I briefly alluded to an argument for Propriety due to Richard Pettigrew (2016) that arguably relies on purely axiological assumptions. I cannot here get into the subtle details of Pettigrew's argument. Suffice it to say that, according to Pettigrew himself, the conditions he imposes epistemic utility functions (or rather, on accuracy measures), from which he derives Propriety, are all satisfied by **b**. (Indeed, the conditions Pettigrew imposes cannot distinguish between two epistemic utility functions that are linear transformations of one another.) And this in turn entails that Pettigrew's requirements are insufficient to motivate Downwards Propriety as a constraint on admissible epistemic utility functions, since **b** is not downwards proper.

The first assumption, is that the epistemic utility of a credence function at a world supervenes on the epistemic utility of credence assignments to individual propositions at a world.²³ Let me spell this out in more detail.

For a given proposition X , say that a *local* epistemic utility function for X is a function l_X that assigns real values to each pair consisting of a real number between 0 and 1 and a truth-value (strictly, l_X assigns real-numbers to each pair of the form (x, i) with $x \in [0, 1]$ and $i \in \{0, 1\}$). Intuitively, $l_X(x, i)$ measures the epistemic utility of assigning credence x to X in a world where X has i as its truth-value.²⁴ With this bit of jargon in place, we can now spell out our supervenience claim as follows:

ATOMISM: The fundamental bearers of epistemic value are credence assignments to individual propositions. In other words: For each utility function u there are local epistemic utility functions l_S^u ($S \subseteq W$) and a function F_u such that (a) for each partition $\mathcal{S} = \{S_1, \dots, S_n\}$ of W , all w , and any credence function c with state space \mathcal{S} ,

$$u(c, w) = F(l_{S_1}^u(c(S_1), \mathbb{1}\{w \in S_1\}), \dots, l_{S_n}^u(c(S_n), \mathbb{1}\{w \in S_n\})),$$

and (b) u is admissible only if each l_S^u is.²⁵

Given Atomism, any constraint on epistemic utility functions must go via a constraint on local epistemic utility functions. For instance, in order to determine whether \mathbf{b} is admissible, we can think of it as the sum of local epistemic utility functions of the form

$$b_X(x, i) = \mathbf{b}(x, i) = \theta - (x - i)^2,$$

and then ask whether each b_X —which is to say, \mathbf{b} —is admissible. If \mathbf{b} is admissible, then that must be in part because \mathbf{b} is.

Adopting Atomism thus requires that we reformulate admissibility conditions not on (global) epistemic utility functions, but rather on local epistemic

²³ Note that this is distinct from the claim that the epistemic utility of a credence function at a world supervenes on its credence assignments to individual propositions. Given some widely accepted assumptions, this latter claim is strictly weaker than Atomism, below.

²⁴ We could have defined a local epistemic utility function for X as a function taking as arguments pairs consisting of a real number in $[0, 1]$ and a possible world. But then it would have made sense to impose a niceness constraint to the effect that the epistemic utility of assigning x to X in w cannot differ from that of assigning x to X in w' unless the truth-value of X is different from its truth-value in w' .

²⁵ Note that Atomism is compatible with there being different admissible ways of aggregating admissible local epistemic utility functions. It may be that the best way of understanding our supervenience claim requires a single way of aggregating local utility functions for all epistemic utility functions. But for our purposes, this weaker principle will do.

utility functions. This will be more or less straightforward depending on what method of *aggregation* we use to define global epistemic utility functions in terms of local epistemic utility functions. For instance, suppose we require that each global utility function u be *additive* in the sense that it satisfies the following condition:

ADDITIVITY: If u is an admissible epistemic utility function, then u is *additive* in the sense that there are local epistemic utility functions l_p^u ($p \subseteq W$) such that for each C and w ,

$$u(C, w) = \sum_{S \in \mathcal{S}_C} l_S^u(C(S), \mathbb{1}\{w \in S\}).$$

Then we could replace Propriety with:

PROPRIETY (LOCAL): If l_p is an admissible local epistemic utility function, then l_p is *proper* in the sense that for any $r \neq r' \in \mathbb{R}$,

$$r \times l_p(r, 1) + (1 - r) \times l_p(r, 0) \geq r \times l_p(r', 1) + (1 - r) \times l_p(r', 0).$$

It is easy to check that an additive global epistemic utility function will be proper iff it is the sum of local epistemic utility functions that are proper.

The second assumption is that some version of the Brier score is an admissible epistemic utility function. More specifically, I will make the following assumption:

BRIER ADMISSIBILITY: For some $\theta \in \mathbb{R}$, \mathbf{b}_θ is an admissible epistemic utility function.

In other words, I will be assuming that some *translation* of the Brier score is admissible (recall from §1 that for any $\theta \in \mathbb{R}$, $\mathbf{b}_\theta = \theta + \mathbf{b}$).²⁶

With these two assumptions in place, we can now formulate an argument to the effect that there can be no axiological justification of Downwards Propriety. The argument begins with the observation that \mathbf{b}_θ is downwards proper if and

²⁶ Admittedly, there are some interesting arguments against Brier Admissibility in the literature—see esp. [Levinstein 2012](#). So perhaps my arguments are ultimately best understood as giving yet another reason for questioning Brier Admissibility. I do not think this is the right conclusion to draw from my arguments in this paper, but I do not intend to argue for this here. Instead, I will simply take on Brier Admissibility as a working assumption. Given the prominent role that versions of the Brier score have played in the literature so far, I trust my taking Brier Admissibility for granted is fair game.

only if $\theta \geq 1/2$, a proof of which follows in a footnote.²⁷ Equivalently, \mathbf{b}_θ is downwards proper iff \mathbf{b}_θ assigns positive epistemic utility to any assignment of credence to a true proposition greater than $1 - 1/\sqrt{2} \approx 0.293$.²⁸

Thus, in order to vindicate Downwards Propriety without giving up on Brier Admissibility, we need to justify an admissibility constraint that rules out \mathbf{b}_θ only if $\mathbf{b}_\theta \theta < 1/2$. And doing so in purely axiological terms requires, in light of Atomism, providing a purely axiological justification for the following claim:

CONSTRAINT: Any admissible local epistemic utility function assigns positive utility to an assignment of credence to true propositions greater than $1 - 1/\sqrt{2}$.

It is worth emphasizing how different Constraint is from other plausible assumptions about epistemic value that have been made in the literature, like Truth-directedness (see §1) or even Propriety itself. Indeed, unlike any of the possible constraints on epistemic utility discussed in Joyce (2009), for example, CONSTRAINT depends on where the zero point in the epistemic utility scale is located. More generally, most constraints on epistemic utility that have been defended in the literature rely on that are satisfied by a given utility function are satisfied by any linear transformation of it. In contrast, Constraint distinguishes between epistemic utility functions and their translations.²⁹

To be sure, there may well be true facts about epistemic utility that rule out some epistemic utility functions as inadmissible without making all of their

²⁷ *Proof:* Start by noting that for each coherent c and c' defined over a coarsening of \mathcal{S}_C that is compatible with c' (in the sense that $c(S) = c'(S')$ whenever $S \in \mathcal{S}_c \cap \mathcal{S}_{c'}$),

$$\mathbb{E}_c[\mathbf{b}_\theta(c')] = \sum_{S \in \mathcal{S}_{c'}} c(S) \times (\theta - c'(S) + c'(S)^2).$$

We can then show that \mathbf{b}_θ is downwards proper iff for each $x, y \geq 0$ such that $x + y = r \leq 1$,

$$r^2 - r \leq \theta - x + x^2 - y + y^2 = \theta - (x + y) + (x + y)^2 - 2xy,$$

or equivalently

$$\theta \geq 2xy.$$

(The left to right direction follows immediately from the definitions, and the fact that if u is proper and $\mathbb{E}_c[u(c)] \geq \mathbb{E}_c[u(c')]$ for all c' defined over a coarsening of \mathcal{S}_c that is compatible with c , then u is downwards proper. For the right to left direction, a simple induction on the size of the difference between \mathcal{S}_c and $\mathcal{S}_{c'}$ will do.)

²⁸ *Proof:* Let $F(x) = (1 - x)^2$, and note that $F(x) = \theta$ iff $\mathbf{b}_\theta(x, 1) = 0$. Since $F(x) = 1/2$ iff $x = 1 - 1/\sqrt{2}$, and F is strictly decreasing for $x \in [0, 1]$, we know that for $x \in [0, 1]$,

$$F(x) \geq 1/2 \Leftrightarrow x \leq 1 - \frac{1}{\sqrt{2}}.$$

²⁹ More precisely, Constraint entails that epistemic utility forms a *ratio scale*, and not just an ordinal, or even an interval scale.

linear transformations as inadmissible. There may well be facts about exactly *how much* epistemic utility a particular credence assignment to a given proposition has. But if the goal is to provide a *justification* of claims like Probabilism, it is not enough to just stipulate some such fact and leave it at that. In particular, in order to provide a justification of Downwards Propriety, we cannot just stipulate the truth of Constraint. (I am assuming without argument that there's nothing intuitive plausible about Constraint.)

I do not, of course, have an argument that no such justification is forthcoming. But it strikes me as highly implausible that there would be some reason for thinking that any admissible way of valuing credences, epistemically, must give special treatment to $1 - 1/\sqrt{2}$ —at least if such a reason must appeal only to claims about what is epistemically valuable. What is so special about $1 - 1/\sqrt{2}$, aside from the fact that we need to work around it in order to ensure Downwards Propriety?

To my mind, the one reasonable strategy here would be to aim for something strictly stronger than what's required in order for b_θ to be downwards proper: something that places the cut-off for admissibility in a more 'natural' place. And the one plausible way to do so would be to seek a purely axiological justification of:

NON-NEGATIVITY: Admissible local epistemic utility functions assign only non-negative values to credence assignments in true propositions.

On this view, *no* assignment of credence to a true proposition can get negative utility. The question then is whether this view can be justified in purely axiological grounds.

I can think of one strategy to justify Non-Negativity. Here's the rough idea. First, think of epistemic value of a credence assignment to a proposition at a world as derived from how 'close' one is to the *epistemically ideal credence assignment* to that proposition at that world. Next, argue that the epistemically ideal credence function at a world is the one that is *maximally accurate* with respect to that proposition at that world. Finally, argue that if all else is equal, any credence assignment to a proposition that is true in w is closer to the epistemically ideal credence assignment to that proposition at w than no credence assignment to that proposition.

There's something odd about this strategy, though. After all, it's not as if we have an intuitive notion of distance that allows to compare how close an assignment of credence that is undefined on a proposition is to an assignment of credence of .7 to that proposition. (Grant that it makes no sense to apply the notion of height to abstract entities like numbers. The strategy we're considering would thus be like saying that any person is closer in height to the tallest person in the world than number 27 is.)

Further, this strategy also cannot be motivated by appeal to some pre-theoretic notion of being ‘closer to getting it right’. You guess the coin will land heads. I decline to guess. As it happens, the coin lands tails. It would be odd, to say the least, to claim that you were closer to getting it right just because I didn’t make any guesses as to how the coin would land.

But perhaps there is some sense in which, when it comes to *credence* assignments, assigning no credence to the proposition that the coin will land heads (say) is further from the epistemically ideal credence assignment than one that assigns any value to that proposition. Still, even if there is such a sense, this would not suffice. In order to justify Non-Negativity it will not do to find some reason that applies only to the proposition that the second toss will land heads. We need a reason to think that for *any* true proposition, no matter how ‘gruesome’, failing to assign credence to that proposition is worse, epistemically, than assigning *any* value to that proposition (no matter how far from that proposition’s truth-value).

It is hard to see, though, what such a reason could be. Recall our discussion at the end of §3.1 about how to think of an agent’s opinion state. On one interpretation, an agent’s opinion state defines the range of propositions the agent is able to entertain. On another, it defines the range of propositions the agent is currently entertaining. Whichever way we go, to think that assigning no value to a true proposition is always at least as bad as assigning any value to it would involve thinking that we are epistemically better off entertaining a true proposition no matter how arbitrarily far our credence is from the proposition’s truth-value than we would be if we were unable to entertain that proposition. The implausibility of this position is best seen by way of examples.

There may be something valuable, epistemically, about being able to entertain propositions about phlogiston—it may help understand why certain ways of thinking about combustion are mistaken, which in turn might help better understand the development of the oxygen theory of combustion. But Non-Negativity entails something much stronger than this. Consider the proposition that phlogiston is not released during combustion—or, if you think such propositions lack a truth-value, the proposition that it is not true that phlogiston is released during combustion. Non-Negativity entails that having an (almost) maximally inaccurate credence in the proposition that phlogiston is not released during combustion—being almost certain that phlogiston is released during combustion—is better, epistemically, than simply failing to entertain that proposition. But it is hard to see why there would be something epistemically better about having (almost) maximally inaccurate views about phlogiston rather than merely not having any views on the matter. If Non-Negativity is right, though, then any non-extreme assignment of credence to the proposition that phlogiston

is not released during combustion (no matter how inaccurate) is better than no assignment at all.

Or take instead the proposition that Scorpios are quiet.³⁰ To think that any non-extremal assignment of credence to this proposition has positive utility would be to think that you gain something, epistemically, by acquiring the ability to entertain that propositions. And this would be so even in a world in which beliefs about people's zodiac signs had played any role in human history. Again, it is hard to see why this would be.

Granted, this is not a decisive argument against NON-NEGATIVITY. There may be, for all I know, some special kind of epistemic benefit you gain by merely coming to entertain the proposition that Scorpios are quiet (or its negation, as the case may be) as long as that credence isn't 0. I think it's safe to assume, though, that there is no such benefit. This assumption may turn out to be false, but I would not bet on it.

5 Taking stock

So where are we? I have argued that the success of axiology-first epistemology depends on the availability of a purely axiological justification of Downwards Propriety. This constitutes a challenge to axiology-first epistemology that, to my mind, has not been adequately appreciated.³¹ I also argued that this challenge cannot be met, at least given two widely accepted assumptions—what I called Atomism and Brier Admissibility.

One issue worth revisiting, though, is whether we can take a different route towards Probabilism. In arguing that the success of axiology-first epistemology depends on whether we can justify Downwards Propriety, I implicitly relied on the following assumption:³²

WEAK PROBABILISM: For each partition, there is some coherent assignment over that partition that is rationally permissible for some agent at some time.

³⁰ So I learn from <https://www.astroved.com/articles/libra-moon-sign-compatibility>.

³¹ To be fair, Pettigrew 2018a, building on Carr 2015, formulated a related challenge to accuracy-first epistemology. But Pettigrew's results seem to rely on his stipulation that Non-Negativity is false (see the principle he labels 'Truth-Directedness' on p. 5). Both Pettigrew and Carr seem to have suspected that for a reasonable epistemic utility function u , if no probability function is u -dominated by one of their restrictions, then some probability function is u -dominated by one of its extensions. And while Pettigrew recognizes that the case for Probabilism may survive even if probability functions are sometimes dominated by their extensions, he does not fully appreciate what is involved in justifying something like Downwards Propriety.

³² Cf. the principle Joyce 2009 calls 'Minimal Coherence'.

This assumption can be motivated by appealing to a principle usually taken for granted by proponents of axiology-first epistemology, viz. that the epistemic utility of a credence function at a world is sensitive only to the truth-value of the relevant propositions at that world and the value the function assigns to those propositions, a principle often known as:³³

EXTENSIONALITY: Suppose C and C' are defined over \mathcal{S}_C and $\mathcal{S}_{C'}$, respectively, and suppose there is a bijection $\phi : \mathcal{S}_C \rightarrow \mathcal{S}_{C'}$ such that $C(s) = C'(\phi(s))$ for all $s \in \mathcal{S}_C$. Further suppose w and w' are such that for all $s \in \mathcal{S}_C$, $w \in s$ iff $w' \in \phi(s)$. Then $u(C, w) = u(C', w)$.³⁴

Now, perhaps Weak Probabilism is too strong. Perhaps there are collections of propositions such that no credence assignment over that collection of propositions, coherent or not, is ever epistemically rational. Say that a partition is *admissible* iff there is some credence function defined over that partition that is rationally permissible for some agent at some time. We could build a reasonable argument for Probabilism if instead of relying on Downwards Propriety, we relied instead on a principle that ensured that no coherent function *defined over an admissible partition* is dominated by a credence function defined over a coarsening of that partition. Doing so would not require vindicating Non-Negativity, but rather the principle that no assignment of credence to any proposition *that is a member of an admissible partition* could ever get negative utility. Whether such a principle can be motivated, though, is a question for some other day.

Perhaps, then, the best way of formulating the upshot of my arguments is as a sort of dilemma for those who seek a purely axiological vindication of principles of rationality: they must either abandon Extensionality or abandon Atomism. I do not have a view on which is the best way to go, perhaps because I don't find either Atomism or Extensionality particularly attractive. Still, given the role that these principles have played in the literature, explicitly or implicitly, it is worth highlighting the consequences they have on the project of axiology-first epistemology.

6 Coda: The population ethics of belief

Given Atomism, the project of devising a global measure of epistemic utility is structurally analogous to that of finding a way of aggregating individuals' well-being in a possible world—one of the central tasks of *population ethics*.

³³ Note that Weak Probabilism only follows from Extensionality if we assume that for each n there is some partition of size n and a probability function defined over that partition that is rationally permissible for some agent at some time.

³⁴ Cf. Joyce 2009, p. 273 and Pettigrew 2016, p. 42.

So one might reasonably expect that any of the measures in $\{\mathbf{b}_\theta : \theta \in \mathbb{R}\}$ —and more generally, any additive epistemic utility function—is vulnerable to analogous objections to those faced by Total Utilitarianism—the view on which a population’s well-being is the sum of the well-being levels of each member of that population.³⁵ In particular, one might suspect that the corresponding additive view of epistemic utility—the view presupposed, given Atomism, by using any member of our family $\{\mathbf{b}_\theta : \theta \in \mathbb{R}\}$ —might give rise to something close to Parfit’s *Repugnant Conclusion* (Parfit 1984).³⁶

In its original formulation, the Repugnant Conclusion is the claim that for any population made up of people with very high levels of well-being, there is an larger population made up of people whose lives are barely worth living whose existence would be better. As Parfit pointed out, if we measure the well-being of a population by looking at the net sum of the well-being of its members, we are led to the Repugnant Conclusion. Is there an epistemic analogue of the Repugnant Conclusion? Does aggregating the epistemic utility of a credence function by simply taking the net sum of the local epistemic utility of its credal assignments give rise to this Epistemic Repugnant Conclusion?

The most natural epistemic analogue of the Repugnant Conclusion could be formulated as follows:

NAÏVE REPUGNANCE: For any credence function c whose credences have very high epistemic utility relative to w , there is a credence function c' defined over a larger set of propositions whose credences have very low epistemic utility relative to w , such that the epistemic utility of c' relative to w is greater than that of c .

Is this a consequence of measuring epistemic utility with any member of our family $\{\mathbf{b}_\theta\}$?

Answering this question requires saying something about what ‘very high’ and ‘very low’ epistemic utility amounts to. A natural benchmark for ‘very high’ is the epistemic utility of a maximally accurate credence. Thus, we can say that $l_P(r, w)$ is high to the extent that it’s close to $l_P(\mathbb{1}\{w \in P\}, \mathbb{1}\{w \in P\})$. What though of ‘very low’ utility?

In discussions of population ethics, the usual benchmark is that of zero utility, or sometimes (instead) the utility level above which lives are considered to be

³⁵ A battery of increasingly powerful results suggests that it is impossible to find a way to aggregate the well-being of a population into a measure that is not incompatible with some extremely plausible judgment about which states of affairs are better than others. The literature here is vast, but much of the recent literature can be seen as responding to Parfit 1984. For an overview of the wide range of impossibility results see Arrhenius 2011.

³⁶ Cf. Carr 2015, Pettigrew 2018b, Talbot 2019.

‘worth living’. If we assume that local epistemic utility takes on both positive and negative values, we can use the zero point as our benchmark for ‘very low’ utility. Accordingly, Naive Repugnance could be reformulated as:

EPISTEMIC REPUGNANCE: For any credence function c whose utility is maximally accurate there is a *possibly non-probabilistic* credence function all of whose utilities have utility arbitrarily close to 0 whose total epistemic utility is greater than that of c .

And again, assuming we are working with additive epistemic utility functions, and assuming that local epistemic utility functions take on both negative and positive values, it is not hard to see that Epistemic Repugnance is true.³⁷

Now, it is far from clear what is the upshot of all this for axiology-first epistemology. We know that for any credence function all of whose credences have maximally high epistemic value, there will be another credence function all of whose credences have barely any (positive) epistemic value. But of course, the latter credence function will be defined over a much larger number of propositions. So the question is: is it better to be almost maximally accurate about a small number of issues than to have views about a much larger number of issues all of which have only a little epistemic value? Pre-theoretically, I do not have a view as to how to answer this question.³⁸

37 Indeed, an even stronger result can be shown with much weaker assumptions—it can be shown that, for any additive epistemic utility whose local epistemic utility functions all satisfy some extremely plausible conditions and take positive and negative values, and any credence function all of whose credences have high epistemic utility at a world, there is another (possibly non-probabilistic) credence function all of whose credences have very low utility at that world that is better, epistemically, than the initial credence function. This is a consequence of the much stronger Theorem 4 in Pettigrew 2018a.

38 In an interesting recent paper, Brian Talbot argues that the Repugnant consequences of the epistemic utility framework are unacceptable (Talbot 2019). To drive his point home, he writes (p. 542): “Consider an attractive credal state which contains only extremely high credences in all the wisdom that humanity will ever acquire. [...] 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).” The point is that, if we use any reasonable, additive epistemic utility function (whose component functions all assign both positive and negative values), we’ll have to think that the ‘attractive’ credal state is worse, epistemically, than at least one ‘repugnant’ state. And this, Talbot argues, just shows that our putative measures of epistemic utility are just not good measures of epistemic *value*. There’s certainly something plausible about Talbot’s observations, but it is hard to feel secure standing on them. For one thing, we need a story as to what ‘wisdom’ is—presumably, any proposition which counts as a bit of wisdom is something that can only be believed by someone who has a large enough body of beliefs to see how that proposition connects with others. And it may well be that any large enough such body of beliefs will be infinite. For another thing, many of the practitioners of epistemic utility theory take Extensionality (above) for granted, and it is not obvious they would be moved by claims about some proposition containing more wisdom than another.

But note how differently things turn out to be if we end up accepting Non-Negativity. If our local epistemic utility functions must take only non-negative values, we can show that any credence function is worse than another credence function *all* of whose credences have arbitrarily low epistemic utility. More precisely, the following is a straightforward consequence of Non-Negativity on the assumption of Atomism:

EXTREME REPUGNANCE: For any $\varepsilon > 0$ and any credence function c and world w such that all of c 's credences are maximally accurate at w , there is a credence function all of whose credences have epistemic utility at most ε greater than the utility of a maximally inaccurate credences at w , whose total epistemic utility is greater, at w , than that of c .

To the extent that we find Naive Repugnance worrying, we should find Extreme Repugnance all the more worrying. I do not know what to make of Extreme Repugnance. But it is something that proponents of epistemic utility theory must have to accept in the event that, contrary to what I have argued, they find some way of justifying Non-Negativity.

References

- Ahlstrom-Vij, Kristoffer & Jeff Dunn (eds.). 2018. *Epistemic Consequentialism*. Oxford: Oxford University Press.
- Arrhenius, Gustaf. 2011. The Impossibility of a Satisfactory Population Ethics. In Hans Colonius & Ehtibar N. Dzhafarov (eds.), *Descriptive and Normative Approaches to Human Behavior, Advanced Series on Mathematical Psychology*.
- Berker, Selim. 2013a. Epistemic Teleology and the Separateness of Propositions. *Philosophical Review* 122(3). 337–393.
- Berker, Selim. 2013b. The Rejection of Epistemic Consequentialism. *Philosophical Issues* 23. 363–387.
- Berker, Selim. 2018. The Unity of Grounding. *Mind* 127(507). 729–777.
- Caie, Michael. 2013. Rational Probabilistic Incoherence. *Philosophical Review* 122(4). 527–575.
- Carr, Jennifer Rose. 2015. Epistemic Expansions. *Res Philosophica* 92(2). 217–236.
- Carr, Jennifer Rose. 2017. Epistemic Utility Theory and the Aim of Belief. *Philosophy and Phenomenological Research* 95(3). 511–534.
- Easwaran, Kenny & Branden Fitelson. 2012. An 'Evidentialist' Worry About Joyce's Argument for Probabilism. *Dialectica* 66(3). 425–433.
- de Finetti, Bruno. 1970. *Teoria della Probabilità*. Vol. I. Turin: Einaudi. English translation: Bruno de Finetti. 1974. *Theory of Probability*. Vol. 1. New York: John Wiley & Sons.
- Fodor, Jerry A. 1975. *The Language of Thought*. Cambridge, Mass.: Harvard Univ Press.

- Franke, Michael & Tikitou de Jager. 2011. Now That You Mention It: Awareness Dynamics in Discourse and Decisions. In Anton Benz, Christian Ebert, Gerhard Jäger & Robert van Rooij (eds.), *Language, Games, and Evolution* (LNAI 6207), 60–91. Heidelberg: Springer.
- Greaves, Hilary. 2013. Epistemic Decision Theory. *Mind* 122(488). 915–952.
- Hájek, Alan. 2009. Arguments for—or against—Probabilism? In Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of belief*, vol. 342 (Synthese Library), 229–251. Dordrecht: Springer.
- Horowitz, Sophie. 2018. Epistemic Value and the Jamesian Goals. In Kristoffer Ahlstrom-Vij & Jeff Dunn (eds.), *Epistemic Consequentialism*, 269–289. Oxford: Oxford University Press.
- Huber, Franz & Christoph Schmidt-Petri (eds.). 2009. *Degrees of Belief*. Vol. 342 (Synthese Library). Dordrecht: Springer Netherlands.
- Joyce, James M. 1998. A Nonpragmatic Vindication of Probabilism. *Philosophy of Science* 65(4). 575–603.
- Joyce, James M. 2009. Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of Belief*, vol. 342 (Synthese Library), chap. 10, 263–297. Dordrecht: Springer Netherlands.
- Joyce, James M. 2013. Why Evidentialists Need Not Worry about the Accuracy Argument for Probabilism. Unpublished manuscript, University of Michigan.
- Joyce, James M. 2018. Accuracy, Ratification, and the Scope of Epistemic Consequentialism. In Kristoffer Ahlstrom-Vij & Jeff Dunn (eds.), *Epistemic Consequentialism*, 240–266. Oxford: Oxford University Press.
- Kemp, Charles, Joshua B. Tenenbaum, Sourabh Niyogi & Thomas L. Griffiths. 2010. A Probabilistic Model of Theory Formation. *Cognition* 114(2). 165–196.
- Konek, Jason & Benjamin Anders Levinstein. 2017. The Foundations of Epistemic Decision Theory. *Mind* 128(509). 69–107.
- Leitgeb, Hannes & Richard Pettigrew. 2010. An Objective Justification of Bayesianism I: Measuring Inaccuracy. *Philosophy of Science* 77(2). 201–235.
- Levinstein, Benjamin Anders. 2012. Leitgeb and Pettigrew on Accuracy and Updating. *Philosophy of Science* 79(3). 413–424.
- Levinstein, Benjamin Anders. 2017. Accuracy Uncomposed: Against Calibrationism. *Episteme* 14(1). 59–69.
- Meacham, Christopher J. G. 2018. Can All-Accuracy Accounts Justify Evidential Norms? In Kristoffer Ahlstrom-Vij & Jeff Dunn (eds.), *Epistemic Consequentialism*, 149–181. Oxford: Oxford University Press.
- Moore, George Edward. 1912. *Ethics*. London: Williams & Norgate.
- Moss, Sarah. 2011. Scoring Rules and Epistemic Compromise. *Mind* 120(480). 1053–1069.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- Pettigrew, Richard. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pettigrew, Richard. 2018a. The Population Ethics of Belief: In Search of an Epistemic Theory X. *Noûs* 52(2). 336–372. Early View, 2016.

- Pettigrew, Richard. 2018b. The Population Ethics of Belief: In Search of an Epistemic Theory X. *Noûs* 52(2). 336–372.
- Predd, J. B., R. Seiringer, E. H. Lieb, D. N. Osherson, H. V. Poor & S. R. Kulkarni. 2009. Probabilistic Coherence and Proper Scoring Rules. *IEEE Transactions on Information Theory* 55(10). 4786–4792.
- Ross, William David. 1930. *The Right and the Good*. Oxford: Clarendon Press.
- Stalnaker, Robert C. 2002. Epistemic Consequentialism. *Aristotelian Society Supplementary Volume* 76(1). 153–168.
- Swanson, Eric. 2006. *Interactions With Context*. PhD dissertation, Massachusetts Institute of Technology.
- Talbot, Brian. 2019. Repugnant Accuracy. *Noûs* 53(3). 540–563.
- Weisberg, Jonathan. 2015. You’ve Come a Long Way, Bayesians. *Journal of Philosophical Logic* 44(6). 817–834.
- Yalcin, Seth. 2007. Epistemic Modals. *Mind* 116(464). 983–1026.