

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 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 some of 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 pairwise disjoint propositions 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 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)

- ³ There has been much critical discussion on whether axiology-first epistemology can provide a justification of norms other than Probabilism (Easwaran & Fitelson 2012, Meacham 2018). Although 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—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 unlike the latter, more general criticisms, my arguments here will apply 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—but for reasons spelled out further in fn. 9, I suspect nothing of substance hinges on this way of proceeding). See e.g. Caie 2013, Greaves & Wallace 2006, Joyce 2009, Moss 2011, Predd et al. 2009.

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 epistemic utility functions are *nice* in the following sense: the utility of c at a world depends only on the truth-value of propositions in c 's state space. In other words, if u is an epistemic utility function, I will assume that for each credence function c , the function $u(c, \cdot)$ is constant throughout each member of \mathcal{S}_c . When u is nice and c is defined over a partition \mathcal{S} of W , it makes sense to talk not just about the u -utility of c at a world, but also about the u -utility of c at any $s \in \mathcal{S}$. This is because, 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 continuous 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:

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

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

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

where $\mathbb{1}\{S = T\}$ equals 1 if $S = T$ and 0 otherwise. More generally, for each $\theta, \lambda \in \mathbb{R}$, with $\lambda > 0$, the function b_{θ}^{λ} , defined by

$$b_{\theta}^{\lambda}(c, w) := - \sum_{S \in \mathcal{S}_c} (\lambda \cdot (c(S) - \mathbb{1}\{w \in S\})^2 + \theta),$$

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.

is a nice epistemic utility function satisfying all our assumptions thus far.⁶

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

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

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 say:⁷

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

⁶ Note that b_θ^λ is *not* the same as the function obtained by multiplying b by λ and adding θ to it. In other words, and in general, $b_\theta^\lambda(c, w) \neq \lambda \cdot b(c, w) + \theta$.

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

We can then derive from principles like this and assumptions about epistemic utility functions 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*.

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').

⁹ See e.g. Joyce 1998. Note that what we are calling 'Probabilism' is strictly weaker than the main tenet of Bayesian epistemology. At most, our formulation of Probabilism only implies that a rational agent's assignments of credence over her *state space* must add up to 1. And this is compatible with the agent's credence function not being a probability function in the familiar sense—compatible, e.g. with there being disjoint X and Y such that her credence in $X \cup Y$ is not the sum of her credences in X and in Y . Fortunately, if we assume, as seems plausible, that rationality is preserved under domain restrictions, what we are calling Probabilism entails the stronger principle that constraints credal assignments over all Boolean combinations of elements of your state space. (Thanks to an anonymous referee for pressing me on this point.)

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

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

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 than or equal to 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.

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.

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.

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 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 on 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.

12 It is tempting to use 'Epistemic Consequentialism' or 'Epistemic Teleology' as a label for the view that the epistemic good is metaphysically prior to the epistemic right (cf. Berker 2013a,b), by analogy with the familiar characterization of consequentialism in ethics (cf. e.g. Moore 1912, Ross 1930; for general discussion, see Berker 2018). As I'm understanding the project, though, axiology-first epistemology aims to provide a *justification* or *vindication* of epistemic norms, and is silent on questions of metaphysical priority. (Though this is arguably the majority view in the literature, Pettigrew 2016 seems to understand the view he labels 'Veritism'—a term he uses almost interchangeably with 'accuracy-first epistemology'—as also maintaining that "there is a single fundamental source of value that is relevant to the epistemic evaluation of credences—it is accuracy" (p. 10).)

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.

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. (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.) And as I mentioned there, the implicit assumption in much of the literature is that the available options for an agent at a time are all and only those credence functions with the same domain as the agent's credence function at that time. In other words, the assumption is that something like Fixed Domain is the right view about when an agent's credence function is rationally permissible.

I think this assumption is mistaken. If we count all credence functions defined over the same domain as the agent's credence function (at a time) as available options for an agent (at that time), then we should also count 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 the 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 uncountably 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.

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.

But note that by the same token, so are all credence functions defined over some but not all of the available propositions—call them *available credence functions*. In whatever sense our agent is able to have any other credence function defined over her opinion set, she is able to have any available credence function. By the agent’s own lights, all functions whose domain consists solely of available propositions have an equal claim to representing an epistemic state she could be in—they have an equal claim to being available options in the canonical decision problem. In other words: insofar as all credence functions defined over an agent’s opinion set are relevant to the epistemic evaluation of her credence function, so are all available credence functions.

Note that the same cannot be said of all credence functions whose domain includes unavailable propositions. Suppose, for instance, we think of the agent’s opinion set as containing all and only 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 those credence functions as one she could have, since by assumption she is unable to entertain the relevant propositions. Or suppose instead we think of the agent’s opinion set as containing only those propositions the agent is currently entertaining. Perhaps she is able to entertain the proposition that there are nowhere differentiable, continuous, real-valued functions (never mind what those are), 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.

Pending some strong reason to accept Fixed Domain, then, we should reject it.¹⁸ 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

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).

¹⁸ In recent work, Richard Pettigrew has made what could be taken to be an indirect argument for Fixed Domain (Pettigrew 2018, § 3.4). But he at best establishes something much weaker than Fixed Domain. I discuss this issue further in fn. 31. (Thanks to an anonymous referee for pressing me to address this concern.)

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 $\mathcal{D} = (c, O^\downarrow[c], u)$, 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 b .¹⁹ As Carr (2015) points out, if we measure epistemic utility using 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 $\{W\}$.²⁰ 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 as irrational every coherent credence function that assigns a non-extreme value to some proposition—the only rational credence functions would be those that assign full credence to some element of their opinion set and zero to every other proposition in their domain—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

¹⁹ See §1 for the definition.

²⁰ 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, $b(c, w) < 0$. But if c_T is the unique probability function defined over $\{W\}$, we have that for all $w \in W$, $b(c_T, w) = 0$. Hence c_T strictly dominates c .

expected u -value of c' relative to c whenever c is probabilistically coherent and the domain of c' is a coarsening of c . For we can set

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

since $u(c', S)$ will be well-defined whenever S is a subset of an element of the state space of c' .²¹

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 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.

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 could 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 relies on two assumptions that are widely shared among proponents of axiology-first epistemology, which I will explain before presenting my argument.

²¹ 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}_c is a subset of an element of $\mathcal{S}_{c'}$, we can conclude that $u(c', S)$ is well-defined.

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 on 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 that 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 u_X that assigns real values to each pair consisting of a real number between 0 and 1 and a truth-value (strictly, u_X assigns real-numbers to each pair of the form (x, i) with $x \in [0, 1]$ and $i \in \{0, 1\}$). Intuitively, $u_X(x, i)$ 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: There is an *aggregation function* F such that u is admissible only if there are admissible local epistemic utility functions u_S ($S \subseteq W$) such that for any credence function c and any w ,

$$u(c, w) = F(\langle u_S(c(S), \mathbb{1}\{w \in S\}) : S \in \mathcal{S}_c \rangle).$$

Given Atomism, any constraint on epistemic utility functions must go via a constraint on local epistemic utility functions and some assumption about how to aggregate local epistemic utility functions.

For instance, suppose that all admissible epistemic utility functions are *additive* in the sense that they satisfy the following condition: there are admissible local epistemic utility functions u_X ($X \subseteq W$) such that for each c and w ,

$$u(c, w) = \sum_{S \in \mathcal{S}_c} u_S(c(S), \mathbb{1}\{w \in S\}).$$

(On this view, the epistemic utility of a credence function at a world is just the sum of the local epistemic utilities of the individual credence assignments to propositions in its state space at that world.) Then, in order to determine whether b is admissible, we can think of it as the sum of local epistemic utility functions of the form

$$b_X(x, i) = b(x, i) = -(x - i)^2,$$

²³ Note that this claim (what I call Atomism, below) is not entailed by the less controversial claim that the epistemic utility of a credence function at a world supervenes on its credence assignments to individual propositions.

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' .

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

Adopting Atomism thus requires that we reformulate admissibility conditions not on ‘global’ epistemic utility functions, but rather on local epistemic 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.

Again assuming that all admissible utility functions are additive, for example, we would replace Propriety with:

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

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

(You can think of u_X as an epistemic utility function defined only over credence functions whose state space consists of X and its negation; the definition above is nothing more than the familiar definition of propriety applied to this narrow class of utility functions.) 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 *positive affine transformation* of the local Brier score is an admissible epistemic utility function. In other words:²⁵

BRIER ADMISSIBILITY: For some $\theta, \lambda \in \mathbb{R}$, with $\lambda > 0$, b_θ^λ is an admissible local epistemic utility function.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 b_θ^λ is downwards proper if and only if $\theta \geq \lambda/2$, a proof of which is in the appendix (see [Corollary 7](#)). Equivalently (see [Corollary 8](#)), b_θ^λ is downwards proper iff b_θ^λ assigns positive epistemic utility to any assignment of credence to a true proposition greater than $1 - 1/\sqrt{2} \approx 0.293$. (Note that Atomism is playing a crucial role here: without it, we cannot turn Downwards Propriety into a constraint on the epistemic utility of individual

²⁵ Admittedly, there are some interesting arguments against Brier Admissibility in the literature—see esp. [Levinstein 2012](#). But my arguments do not essentially depend on Brier Admissibility—see §5.

²⁶ Recall from §1 that for any $\theta, \lambda \in \mathbb{R}, \lambda > 0$

$$b_\theta^\lambda(c, w) = \sum_{S \in \mathcal{S}_c} b_\theta^\lambda(c(S), \mathbb{1}\{w \in S\}),$$

where $b_\theta^\lambda(x, i) = \theta + \lambda \cdot b(x, i)$, and $b(x, i) = -(x - i)^2$

credence assignments at a world. This is because, without Atomism, the epistemic utility of assigning some specific credence to a proposition may vary depending on what other propositions you assign credence to.)

Thus, in order to vindicate Downwards Propriety without giving up on Brier Admissibility, we need to justify an admissibility constraint that rules out b_{θ}^{λ} only if $\theta < \lambda/2$. 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 that are satisfied by a given utility function are also satisfied by any linear transformation of it. In contrast, Constraint distinguishes between epistemic utility functions and their translations.²⁷

To be sure, there may well be true facts about epistemic utility that rule out some epistemic utility functions as inadmissible without also ruling out all of their 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 just to stipulate some such fact. 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 intuitively 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 is 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 is required in order for b_{θ}^{λ} to be downwards proper:

²⁷ In other words, Constraint entails that epistemic utility forms a *ratio scale*, and not just an ordinal, or even an interval scale.

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 us 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.

Now, 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. But 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 set. 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 interpretation we adopt, Non-Negativity entails some highly implausible claims about epistemic value. For it entails that we are always 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 claim is best seen by way of examples.

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.²⁸

Or take instead the proposition that Scorpios are quiet.²⁹ If we accept that any non-extremal assignment of credence to this proposition has positive utility, then we would have to accept that we gain something, epistemically, by acquiring the ability to entertain that proposition. And this would be so even in a world in which beliefs about people's zodiac signs had played no role in human history. Again, it is hard to see why this would be.³⁰

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,

28 To be sure, there may be something valuable, epistemically, about being able to entertain *some* 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 that there is something valuable about being able to entertain *any* proposition about phlogiston, no matter how inaccurate our views on the relevant propositions are.

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

30 Cf. Carr 2015, p. 231f.

though, that there is no such benefit. This assumption may turn out to be false, but I would not bet on it.³¹

5 Taking stock

I have argued that the success of axiology-first epistemology depends on the availability of a purely axiological justification of Downwards Propriety.³² This constitutes a challenge to axiology-first epistemology that, to my mind, has not been adequately appreciated. I also argued that this challenge cannot be met, at least given two widely accepted assumptions—what I called Atomism and Brier Admissibility.

Essentially, my argument for this last claim was that given Atomism and Brier Admissibility, any axiological justification of Downwards Propriety would

³¹ The challenge I've presented above is related to but importantly different from one raised by Jennifer Carr and later generalized by Richard Pettigrew (Carr 2015, Pettigrew 2018). Pettigrew's main result builds on some powerful results in population ethics. His claim is roughly that, unless we reject some plausible assumptions about epistemic utility, we must learn to live with the fact that any coherent credence function is dominated by another credence function with a more fine-grained domain. Going over the technical details of Pettigrew's arguments would take us too far afield. But let me just note that they seem to depend on some controversial assumptions about epistemic value: first, Pettigrew's challenge, unlike mine, can be dismissed by rejecting Atomism; second, some of Pettigrew's results (esp. Theorem 4) depend on his idiosyncratic definition of local epistemic utility functions (Pettigrew stipulates without much discussion that a perfectly accurate assignment of credence must get positive utility and a perfectly inaccurate assignment of credence must get negative utility, and thus that it isn't merely a matter of convention which credence assignments get zero epistemic utility).

Still, if we grant Pettigrew's assumptions, his main result raises an important challenge. The main selling point of axiology first epistemology was supposed to be that it offered a justification of Probabilism. Such a justification relies on the fact that coherent credence functions, unlike incoherent ones, were not dominated by another one with the same domain. But if all credence functions—coherent or not—are dominated by some other credence function, even if the dominating one has a richer domain, the case for Probabilism is arguably undermined.

In response to this challenge, Pettigrew essentially endorses Fixed Domain (see esp. Pettigrew 2018, p. 365). But, as far as I can tell, the only reason he offers in its support is that it allows us to accept that all coherent functions are dominated without giving up on the argument for Probabilism. And this is not a particularly good reason, since something much weaker than Fixed Domain allows us to do that. All we need is to reject that credence functions whose domain is a refinement of an agent's opinion set are ever available options in a canonical decision problem. This suffices to address Pettigrew's challenge, but it does not suffice to provide an axiological justification of Downwards Propriety. And without an axiological justification of Downwards Propriety, the axiology-first case for Probabilism is undermined. (Thanks to an anonymous referee for pushing me to better clarify how my challenge differs from Carr's and Pettigrew's.)

³² It is worth noting that Downwards Propriety is independently plausible, at least if we accept Extensionality (below): rejecting it would then amount to the claim that one can be epistemically rational while thinking one might do better, epistemically, by abandoning some of one's credences. And this, it seems, flies in the face of the widely held assumption that epistemic rationality is *immodest*, in the sense that an epistemically rational agent should take herself to be doing as well as she can, epistemically and by her own lights. (Thanks here to an anonymous referee.)

amount to an axiological justification of Non-Negativity. Since, I argued, no axiological justification of Non-Negativity is forthcoming, Atomism and Brier Admissibility entail that there is no axiological justification of Downwards Propriety.

This argument could be strengthened. For example, a similar argument to the one in the paper can be used to show that if some version of the *spherical score* is admissible, then justifying Downwards Propriety would also require justifying Non-Negativity. (See [Corollary 9](#) and [Corollary 10](#), in the appendix.) But it is an open question what is the strongest version of this argument. It follows from [Corollary 6](#) (also in the appendix) that if we accept Extensionality (below), then justifying Downwards Propriety requires justifying a privileged zero point of epistemic utility. More specifically, [Corollary 6](#) entails that if u is an extensional, downwards proper epistemic utility function, there will be some translation of it that is not downwards proper. It may be that axiological considerations can be used to rule out as inadmissible some translations of admissible epistemic utility functions, but whether this is so is a question for some other time.³³

It is worth revisiting, though, whether there is 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 (cf. the principle [Joyce \(2009\)](#) calls ‘Minimal Coherence’):

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

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-values of the relevant propositions at that world and the value the function assigns to those propositions—a principle often known as:³⁴

33 Recall that we have ruled out of considerations utility functions like the *additive log score* l , which is the additive epistemic utility function generated by the following local epistemic utility function l :

$$l(x, i) = \ln(|1 - i - x|),$$

since the range of the log score extends beyond the real line (since $\ln(0) = \infty$). It is not difficult, however, to show that the log score is not downwards proper, since the l -value of the trivial credence function will always be greater than or equal to that of any other credence function. And with suitable care when dealing with arithmetic calculations involving $+\infty$ and $-\infty$ —see e.g. [Rockafellar 1970](#), §4—one can adapt the proofs of [Corollary 7](#) and [Corollary 8](#) in the appendix to show that l_θ^λ is downwards proper if and only if $\theta \geq 2\lambda \ln 2$, which will obtain if and only if for all $x \geq 1/4$, $l_\theta^\lambda(x, 1) > 0$.)

34 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.

EXTENSIONALITY: Suppose c and c' are defined over \mathcal{S}_c and $\mathcal{S}_{c'}$, respectively, and suppose there is a bijection $g : \mathcal{S}_c \rightarrow \mathcal{S}_{c'}$ such that $c(S) = c'(g(S))$ 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 $w' \in g(S)$. Then $u(c, w) = u(c', w)$.³⁵

In other words, Extensionality ensures that the epistemic utility of a credence function at a world does not depend on the *content* of the propositions over which the function is defined.³⁶

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 better alternative, perhaps because I don't find either Atomism or Extensionality particularly attractive. Still, given the role

³⁵ Cf. Joyce 2009, p. 273 and Pettigrew 2016, p. 42.

³⁶ Of course, there are some risks that come with rejecting Extensionality. In an interesting recent paper, Brian Talbot argues that the Repugnant consequences of the epistemic utility framework are unacceptable, even if they do not suffice to undermine the argument for Probabilism (Talbot 2019). To drive his point home, he asks us (p. 542) to consider “an attractive credal state which contains only extremely high credences in all the wisdom that humanity will ever acquire.” and compare it 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).” Talbot then claims 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*. A proponent of Extensionality could simply reject the claim that facts about wisdom are relevant to epistemic evaluation—it is not as clear what to say in response to Talbot's concern if we reject Extensionality. My sense though is that before we can take this concern too seriously, we need a better sense of 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.

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.³⁷

Appendix

The purpose of this appendix is to provide a proof of my claims that b_θ^λ is downwards proper iff for all $x \geq 1 - 1/\sqrt{2}$, $b_\theta^\lambda(x, 1) \geq 0$ iff $\theta \geq \lambda/2$ (Corollary 7 and Corollary 8). Along the way, I prove two more general characterization results: one for the class of additive, downwards proper epistemic utility functions (Proposition 2) and one for the class of additive and extensional downwards proper epistemic utility functions (Proposition 5). I then use this last result to characterize the class of positive affine transformations of the spherical score that are downwards proper (Corollary 9 and Corollary 10).

Before we begin, let me introduce some additional terminology. For any probability function c and any coarsening \mathcal{S}' of \mathcal{S}_c , the *restriction of c to \mathcal{S}'* is the unique probability function c' defined over \mathcal{S}' such that, for each $S' \in \mathcal{S}'$:

$$c'(S') = \sum_{\substack{S \subseteq S' \\ S \in \mathcal{S}_c}} c(S).$$

I will say that c' is a *restriction of c to \mathcal{S}'* iff $\mathcal{S}_{c'}$ is a coarsening of \mathcal{S}_c and c' is the restriction of c to $\mathcal{S}_{c'}$.

Our first step now is to establish the following lemma, which will simplify the proofs to follow.

Lemma 1. *An epistemic utility function u is downwards proper iff it is proper and for each probability function c and each restriction c' of c ,*

$$\mathbb{E}_c[u(c)] \geq \mathbb{E}_{c'}[u(c')].$$

Proof. Start by noting that for each c and each c^* defined over a coarsening \mathcal{S}^* of \mathcal{S}_c ,

$$\mathbb{E}_c[u(c^*)] = \mathbb{E}_{c \uparrow \mathcal{S}^*}[u(c^*)],$$

³⁷ A much earlier version of this material was presented at the workshop *Population Ethics Meets Formal Epistemology* at the Institute for Future Studies in Stockholm. Thanks to H. Orri Stefánsson, Gustaf Arrhenius, and to the workshop participants, especially Hilary Greaves and Richard Pettigrew, for their helpful questions. Thanks also to Hilary Kornblith and Phil Bricker for some insightful comments that helped me see my way out of what I thought was a fatal objection to this paper. Finally, thanks to an Associate Editor and three anonymous referees for this journal for their careful reading of the paper and their probing questions, and especially to Sophie Horowitz, Chris Meacham, and Katia Vavova for detailed and extremely helpful comments on earlier versions of this paper.

where $c \upharpoonright \mathcal{S}^*$ is the restriction of c to \mathcal{S}^* . The left to right direction of [Lemma 1](#) now follows immediately from the definition, since if c' is a restriction of c to \mathcal{S}' , $c \upharpoonright \mathcal{S}' = c'$.

For the right to left direction, assume u is proper and that for each coherent c and each restriction c' of c ,

$$\mathbb{E}_c[u(c)] \geq \mathbb{E}_{c'}[u(c')].$$

Now let c^* be an arbitrary credence function defined over \mathcal{S}^* and let $c' = c \upharpoonright \mathcal{S}^*$. From our initial observation, we have that

$$\mathbb{E}_c[u(c^*)] = \mathbb{E}_{c'}[u(c^*)],$$

since $\mathcal{S}^* = \mathcal{S}_{c'}$ and $c' = c \upharpoonright \mathcal{S}_{c'}$. Since u is proper, we also have that

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

We can thus conclude that

$$\mathbb{E}_c[u(c)] \geq \mathbb{E}_c[u(c^*)],$$

as desired. □

Recall that a *local epistemic utility function* for X is a function u_X that assigns a real number to each pair (r, i) , where $r \in [0, 1]$ and $i \in \{0, 1\}$. Recall too that if u is an additive epistemic utility function (see [§4](#)), there is a family of local epistemic utility functions $\{u_S : S \subseteq W\}$ such that, for each c and w :

$$u(c, w) = \sum_{S \in \mathcal{S}_c} u_S(c(S), \mathbb{1}\{w \in S\}).$$

Slightly abusing notation, for $r, r' \in [0, 1]$, I will write $\mathbb{E}_r[u_X(r')]$ to denote the expected u_X -value of assigning r' to X relative to a probability function that assigns r to X . In other words:

$$\mathbb{E}_r[u_X(r')] := r \cdot u_X(r', 1) + (1 - r) \cdot u_X(r', 0).$$

Much like I did in the formulation of Propriety (Local) ([§4](#)), I will say that u_X is *proper* iff for each $r \neq r' \in [0, 1]$,

$$\mathbb{E}_r[u_X(r)] \geq \mathbb{E}_r[u_X(r')].$$

I will say that a *family* of local epistemic utility functions $\{u_X : X \subseteq W\}$ is *downwards proper* iff each u_X is proper and, for any non-empty $Y, X \subseteq W$, with $X \cap Y = \emptyset$, and each $r, r' \in [0, 1]$ with $r + r' \leq 1$,

$$\mathbb{E}_r[u_X(r)] + \mathbb{E}_{r'}[u_Y(r')] \geq \mathbb{E}_{r+r'}[u_{X \cup Y}(r + r')].$$

The following gives us a simple characterization of additive epistemic utility functions that are downwards proper.

Proposition 2. *An additive epistemic utility function u is downwards proper iff the family $\{u_X : X \subseteq W\}$ of local epistemic utility functions is downwards proper.*

Proof. Start by noting that if u is additive, then for each probability function c and each credence function c' defined over \mathcal{S}_c ,

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

(This follows straightforwardly from the additivity of expectation.)

Now, suppose u is additive and downwards proper. Pick $X, Y \subseteq W$ and fix $r, r' \in [0, 1]$ with $r + r' \leq 1$. Let c be the unique probability function defined over the partition $\{X, Y, \overline{(X \cup Y)}\}$ with $c(X) = r$ and $c(Y) = r'$, and let c' be the restriction of c to $\{X \cup Y, \overline{(X \cup Y)}\}$ (where \bar{S} denotes the set-theoretic complement of S).

Letting $Z = X \cup Y$ and

$$\Phi(x, S) := \mathbb{E}_x[u_S(x)],$$

we know by definition that

$$\mathbb{E}_c[u(c)] = \Phi(r, X) + \Phi(r', Y) + \Phi(1 - (r + r'), \bar{Z})$$

and

$$\mathbb{E}_c[u(c')] = \Phi(r + r', Z) + \Phi(1 - (r + r'), \bar{Z}).$$

Since u is downwards proper and c' is defined over a coarsening of \mathcal{S}_c , we can infer

$$\Phi(r, X) + \Phi(r', Y) + \Phi(1 - (r + r'), \bar{Z}) \geq \Phi(r + r', Z) + \Phi(1 - (r + r'), \bar{Z}),$$

which in turn entails

$$\mathbb{E}_r[u_X(r)] + \mathbb{E}_{r'}[u_Y(r')] \geq \mathbb{E}_{r+r'}[u_{X \cup Y}(r + r')],$$

as desired.

Suppose now that u is an additive epistemic utility function and $\{u_S : S \subseteq W\}$ is downwards proper. Pick a coherent c and fix a restriction c' of c . Since $\{u_S : S \subseteq W\}$ is downwards proper, we know that for any two disjoint, non-empty $X, Y \subseteq W$ and any $r, r' \in [0, 1]$ such that $r + r' \in [0, 1]$,

$$\Phi(r, X) + \Phi(r', Y) \geq \Phi(r + r', X \cup Y).$$

A simple induction on N then allows to conclude that for any partition $\{X_1, \dots, X_N\}$ of a non-empty $X \subseteq W$ and non-negative real numbers $\{r_1, \dots, r_N\}$ whose sum is in $[0, 1]$,

$$\sum_N \Phi(r_i, X_i) \geq \Phi\left(\sum_N r_i, \bigcup_N X_i\right) = \Phi\left(\sum_N r_i, X\right).$$

In particular, for each $S' \in \mathcal{S}_{c'}$ (again, letting S range only over members of \mathcal{S}_c),

$$\sum_{S \in \mathcal{S}'} \Phi(c(S), S) \geq \Phi\left(\sum_{S \in \mathcal{S}'} c(S), S'\right) = \Phi(c'(S'), S').$$

And since

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

we can conclude that

$$\mathbb{E}_c[\mathbf{u}(c)] \geq \sum_{S'} \Phi(c'(S'), S') = \mathbb{E}_{c'}[\mathbf{u}(c')].$$

From [Lemma 1](#) we then conclude that \mathbf{u} is downwards proper. □

A function $\varphi : X \rightarrow \mathbb{R}$ with $X \subseteq \mathbb{R}$ is said to be *convex* iff for all $\lambda \in [0, 1]$, $x, y \in X$ with $x + y \in X$,

$$\varphi(\lambda x + (1 - \lambda)y) \leq \lambda\varphi(x) + (1 - \lambda)\varphi(y).$$

It is *subadditive* iff for all $x, y \in X$ with $x + y \in X$,

$$\varphi(x + y) \leq \varphi(x) + \varphi(y).$$

The following fact will come in handy in establishing our second main result.

Fact 3. *Let $\varphi : [0, 1] \rightarrow \mathbb{R}$ be convex. The following are all equivalent:*

- (i) φ is subadditive.
- (ii) For all $x \in [0, 1]$ and all $n > 0$ with $nx \in [0, 1]$, $\varphi(nx) \leq n\varphi(x)$.
- (iii) For all $x \in [0, 1]$ such that $2x \in [0, 1]$, $\varphi(2x) \leq 2\varphi(x)$.

Proof. It is straightforward to show, by induction on n , that (i) entails (ii); (ii) clearly entails (iii). All that is left to show is that (ii) entails (i). So fix a convex $\varphi : [0, 1] \rightarrow \mathbb{R}$ satisfying (iii). Take $x, y \in [0, 1]$ such that $x + y \in [0, 1]$. Clearly, $1/2(x + y) \in [0, 1]$, and thus from (iii) and the fact that φ is convex we can conclude

$$\varphi(x + y) = \varphi\left(2 \cdot \frac{1}{2}(x + y)\right) \leq 2\varphi\left(\frac{1}{2}(x + y)\right) \leq 2 \cdot \frac{1}{2}(\varphi(x) + \varphi(y)) = \varphi(x) + \varphi(y).$$

□

I will also rely on a well-known fact—an immediate consequence of Savage’s characterization of proper scoring rules ([Savage 1971](#)).³⁸

³⁸ For a proof, see Proposition 2 in [Predd et al. 2009](#).

Fact 4. A local epistemic utility function $u : [0, 1] \times \{0, 1\} \rightarrow \mathbb{R}$ is proper iff the function φ_u defined by

$$\varphi_u(x) := \mathbb{E}_x[u(x)]$$

is convex. □

Proposition 5. Let u be an additive and extensional epistemic utility with corresponding local epistemic utility function u , and for $x \in [0, 1]$ let $\varphi_u(x) = \mathbb{E}_x[u(x)]$. The following are equivalent:

- (i) u is downwards proper.
- (ii) φ_u is convex and subadditive.
- (iii) φ_u is convex and $\varphi_u(2x) \leq 2\varphi_u(x)$ for all $0 \leq x \leq 1/2$.
- (iv) φ_u is convex and for all $n \geq 0$, $x \in [0, 1]$, $nx \in [0, 1]$ entails $\varphi_u(nx) \leq n\varphi_u(x)$.

Proof. Fix an additive and extensional u . Extensionality entails that the family $\{u_X : X \subseteq W\}$ is constant with $u_X = u$. Hence, from [Proposition 2](#) we can conclude that u is downwards proper iff φ_u is proper and subadditive. [Fact 4](#) now ensures that u is downwards proper iff φ_u is convex and subadditive, and thus that (i) and (ii) are equivalent. That (ii)-(iv) are pairwise equivalent is a straightforward consequence of [Fact 3](#). □

Corollary 6. Suppose u is an additive, proper, and extensional epistemic utility function with corresponding local epistemic utility function u . Let u_θ^λ be the additive epistemic utility function generated by u_θ^λ , where

$$u_\theta^\lambda(x, i) := \lambda \cdot u(x, i) + \theta,$$

and let

$$\varphi_\theta^\lambda(x) := \mathbb{E}_x[u_\theta^\lambda(x)].$$

Then u_θ^λ is downwards proper if and only if for all $0 \leq x \leq 1/2$

$$\theta \geq \lambda(\varphi_\theta^1(2x) - 2\varphi_\theta^1(x)).$$

Proof. Since u is proper, so is u , so the linearity of expectation ensures that u_θ^λ is proper for all $\lambda, \theta \in \mathbb{R}$. From [Fact 4](#) we can infer that each φ_θ^λ is convex, and [Proposition 5](#) allows us to conclude that u_θ^λ is downwards proper if and only if for all $0 \leq x \leq 1/2$,

$$\varphi_\theta^\lambda(2x) \leq 2\varphi_\theta^\lambda(x),$$

which by definition obtains if and only if for all $0 \leq x \leq 1/2$

$$\theta \geq \lambda(\varphi_\theta^1(2x) - 2\varphi_\theta^1(x)),$$

as desired. □

Corollary 7. b_θ^λ is downwards proper if and only if $\theta \geq \lambda/2$.

Proof. Recall that b_θ^λ is an additive and extensional epistemic utility function with corresponding local epistemic utility function b_θ^λ . Of course, b is proper. So, given [Corollary 6](#) it suffices to establish that $\theta \geq \lambda/2$ iff for all $0 \leq x \leq 1/2$, $\theta \geq \lambda(\varphi_b(2x) - 2\varphi_b(x))$, where

$$\varphi_b(x) := \mathbb{E}_x[b(x)] = -x + x^2.$$

Now,

$$\varphi_b(2x) - 2\varphi_b(x) = -2x + 4x^2 + 2x - 2x^2 = 2x^2.$$

And since $2x^2$ is a strictly increasing function of x for $x \geq 0$, we conclude that for all $0 \leq x \leq 1/2$, $\theta \geq \lambda \cdot 2x^2$ iff

$$\theta \geq \lambda \cdot 2(1/2)^2 = \lambda/2,$$

as desired. □

Corollary 8. b_θ^λ is downwards proper if and only if for all $x \in [0, 1]$ with $x \geq 1 - 1/\sqrt{2}$, $b_\theta^\lambda(x, 1) \geq 0$.

Proof. Given [Corollary 7](#), it suffices to show that $\theta \geq \lambda/2$ iff for all $x \in [0, 1]$ with $x \geq 1 - 1/\sqrt{2}$, $b_\theta^\lambda(x, 1) \geq 0$.

Note now that from the definition of b_θ^λ we have that for each $\lambda, \theta, \theta' \in \mathbb{R}$ with $\lambda > 0$ and each $x \in [0, 1]$,

$$b_\theta^\lambda(x, 1) > b_{\theta'}^\lambda(x, 1) \text{ iff } \theta > \theta'.$$

Note too that, since for each $\lambda, \theta \in \mathbb{R}$, with $\lambda > 0$, b_θ^λ is a strictly increasing function of x over $[0, 1]$, we know that for each $\lambda, \theta \in \mathbb{R}$ with $\lambda > 0$, each $x \in [0, 1]$, and each $r \in \mathbb{R}$, $b_\theta^\lambda(x, 1) \geq r$ iff for all $x' \geq x$, $b_\theta^\lambda(x', 1) \geq r$. The key observation now is that for each $\lambda > 0$, $b_{\lambda/2}^\lambda(1 - 1/\sqrt{2}, 1) = b_{1/2}^1(1 - 1/\sqrt{2}, 1) = 0$, since that allows us to conclude that

$$\begin{aligned} \theta \geq \lambda/2 &\Leftrightarrow b_\theta^\lambda(1 - 1/\sqrt{2}, 1) \geq b_{\lambda/2}^\lambda(1 - 1/\sqrt{2}, 1) \\ &\Leftrightarrow \text{for all } x \in [0, 1] \text{ with } x \geq 1 - 1/\sqrt{2}, b_\theta^\lambda(x, 1) \geq b_{\lambda/2}^\lambda(1 - 1/\sqrt{2}, 1) = 0, \end{aligned}$$

as desired. □

The *spherical* score s is an additive and extensional utility function with corresponding local utility function $s : [0, 1] \times \{0, 1\} \rightarrow \mathbb{R}$ given by

$$s(x, i) := \frac{(x + i - 1)}{\sqrt{x^2 + (1 - x)^2}}.$$

For $\theta, \lambda \in \mathbb{R}$, $\lambda > 0$ let s_θ^λ be the additive extensional utility function generated by $s_\theta^\lambda := \lambda s + \theta$.

As the following observations make evident, the spherical score is downwards proper, but many of its translations are not.

Corollary 9. s_θ^λ is downwards proper iff $\theta \geq \lambda(1 - \sqrt{2})$.

Proof. Let $\varphi_s(x) := \mathbb{E}_x[s(x)]$, and note that

$$\varphi_s(x) = x \cdot \frac{x}{\sqrt{x^2 + (1-x)^2}} + (1-x) \cdot \frac{(1-x)}{\sqrt{x^2 + (1-x)^2}} = \sqrt{x^2 + (1-x)^2}.$$

Let

$$g(x) := \varphi_s(2x) - 2\varphi_s(x) = \sqrt{4x^2 + (1-2x)^2} - \sqrt{4x^2 + 4(1-x)^2},$$

and note that

$$g(x) = \sqrt{4x^2 + (1-2x)^2} - \sqrt{4x^2 + (1-2x)^2 + (3-4x)},$$

so that g is strictly increasing over $[0, 1/2]$. Thus, $\theta \geq \lambda g(x)$ for all $0 \leq x \leq \frac{1}{2}$ iff

$$\theta \geq \lambda g(1/2) = \lambda(1 - \sqrt{2}).$$

□

Corollary 10. s_θ is downwards proper iff for all $x \in [0, 1]$, $s_\theta(x, 1) \geq 0$ whenever

$$x \geq \frac{(\sqrt{8+8\sqrt{2}}-2)(-1+2\sqrt{2})}{14} \approx 0.07396.$$

Proof. Routine calculation shows that

$$s_{\lambda(1-\sqrt{2})}^\lambda \left(\left(\frac{(\sqrt{8+8\sqrt{2}}-2)(-1+2\sqrt{2})}{14}, 1 \right) \right) = s_{1-\sqrt{2}} \left(\left(\frac{(\sqrt{8+8\sqrt{2}}-2)(-1+2\sqrt{2})}{14}, 1 \right) \right) = 0.$$

We can now simply repeat, *mutatis mutandis*, the steps of the proof of [Corollary 8](#).

□

References

- Ahlstrom-Vij, Kristoffer & Jeff Dunn (eds.). 2018. *Epistemic Consequentialism*. Oxford: Oxford University Press.
- 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: **finetti1974**.
- Fodor, Jerry A. 1975. *The Language of Thought*. Cambridge, Mass.: Harvard Univ Press.
- Franke, Michael & Tikitu 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.
- Greaves, Hilary & David Wallace. 2006. Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind* 115(459). 607–632.
- 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.
- Pettigrew, Richard. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pettigrew, Richard. 2018. 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.
- Rockafellar, R. Tyrrell. 1970. *Convex Analysis*. Princeton: Princeton University Press.
- Ross, William David. 1930. *The Right and the Good*. Oxford: Clarendon Press.
- Savage, Leonard J. 1971. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association* 66(336). 783–801.
- 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.