AN OBJECTION OF VARYING IMPORTANCE TO EPISTEMIC UTILITY THEORY

Abstract. Some propositions are more epistemically important than others. Further, how important a proposition is is often a contingent matter—some propositions count more in some worlds than in others. Epistemic Utility Theory cannot accommodate this fact, at least not in any standard way. For EUT to be successful, legitimate measures of epistemic utility must be proper, i.e., every probability function must assign itself maximum expected utility. Once we vary the importance of propositions across worlds, however, normal measures of epistemic utility become improper. I argue there isn't any good way out for EUT.

0. Introduction

Bob writes an article for Wikipedia on this history of his front lawn. There's nothing particularly special about his lawn, but Bob has kept rather close track on the average height of the grass, the varying color of the blades, and the different types of mowers he's used over the years. Wikipedia promptly rejects the article. "We're not an indiscriminate collector of information," they write, "but instead a collector of *important* information." Bob's article, while containing many true and verified claims, is not worth publishing.

Wikipedia, of course, does not collect merely practically useful information. It has articles on medieval heresies, alchemy, and quasars. These articles do not, in general, come in handy when the average reader needs to make a decision. Instead, such articles are of epistemic importance. Even if they aren't directly useful, they make the cut—unlike facts about Bob's lawn—because they contain content that the editors judge worth knowing. Other information, while of practical use, is not of much epistemic value: the outcome of a particular coin-flip that you've wagered on, the solution to the next Bitcoin hash, or the length of time your brother will hog the shower.

Epistemic Utility Theory aims to quantify how epistemically well-off an agent is. Generally, epistemic utility theorists identify epistemic well-being with accuracy. The higher your credences in truths and the lower your credences in falsehoods, the better off you are all epistemic things considered. However, as we just noted, some propositions seem more epistemically important than others. It's epistemically better to have a high credence in, say, a true claim about fundamental physics than in a true claim about a random episode of your favorite '90s sitcom.

The EUT literature has not paid all that much attention to the differing levels of epistemic importance of various propositions. But where attention has been paid, epistemic utility theorists note simply that one can assign different weights to different propositions, and that this causes no problem for EUT.

In this paper, I will argue that the varying level of epistemic importance is a major problem for epistemic utility theory. In fact, as we'll see, it appears impossible for common measures of epistemic utility both to capture the differing levels of epistemic importance in an adequate way and to do the work epistemic utility theorists require of them.

The fundamental problem is that how important a proposition is can be a contingent matter. In some worlds, facts about the behavior of carbon atoms are very important, but in other worlds, they aren't. So, the weight given to a proposition should vary from world to world. Unfortunately, EUT cannot handle varying weights, at least not in any remotely orthodox way. For EUT to succeed, permissible measures of epistemic utility must be *proper*. That is, each probability function must expect itself to achieve more epistemic utility than any other function. However, once the weights vary across worlds, the standard measures of epistemic utility become improper, which undermines the entire epistemic utility project.

Here's the plan. In Section 1, I go over the basics of Epistemic Utility Theory. In Section 2, I motivate the view that the epistemic importance of propositions is at least sometimes contingent. Section 3 explains why Epistemic Utility Theory cannot handle contingent weightings of propositions. Finally, Section 4 looks at some potential ways out and concludes that none of them are especially attractive.

1. Epistemic Utility Theory

Epistemic Utility Theory aims to do two things. First, it establishes a method of quantifying the epistemic value of a credence function at a world. Second, it uses this notion of value to justify various norms (such as probabilism, conditionalization, and so on) as binding on rational agents.[1]

A core component of this view is that some properly specified notion of *accuracy* is epistemic utility. That is, accuracy is what is ultimately of epistemic value. Accuracy is alethic: the higher your credence in truths and the lower your credence in falsehoods, the more accurate you are. As we'll see soon, though, there's a lot more to the notion that just that. The kind of accuracy worth having will require a lot more spelling out and specification.

For the sake of concreteness, let's start with one common way to measure epistemic utility. We can measure how *in*accurate a credence function $c : \mathcal{F} \to [0,1]$ is at a world w with the:

Brier Score: $\mathrm{BS}(c, w) = \sum_{X \in \mathcal{F}} (c(X) - w(X))^2$

where $w(X) = 1$ if X is true at w and $= 0$ otherwise.

The Brier Score simply measures the squared divergence between c and the perfectly accurate credence function at w, i.e., the one that assigns credence 1 to all truths and 0 to all falsehoods. So, the higher c's credence in truths and the lower c's credence in falsehoods, the lower c's Brier Score.

We can then use this measure of inaccuracy to identify one candidate measure of epistemic utility.

Brier Utility: The Brier-Epistemic utility of a credence function c at a world w is $1 - \mathrm{BS}(c, w)$.

[1]See, for instance, Joyce (1998, 2009); Greaves and Wallace (2006); Pettigrew (2016).

If this is correct, then the total all epistemic things considered value of a credence function at a world—i.e., how epistemically well-off an agent with the credence function at that world is—is given by the credence function's Brier Utility at that world.

Of course, the Brier Score is simply one specific candidate measure, and we will have to tweak our understanding of epistemic value as we go on.

1.1. **Justifying Norms.** Once we have a measure of value in hand, we can exploit the tools of decision theory to determine which norms are binding on epistemically rational agents.

For the sake of illustration, we'll quickly rehearse the EUT justification of probabilism.[2] Suppose c is a credence function that fails to satisfy the axioms of finitely additive probability. Then there exists a probabilistically coherent c' that will obtain strictly more Brier Utility than c at every possible world. Furthermore, there is no c'' (coherent or not) that will obtain strictly more Brier Utility than c'.[3]

In other words, if a credence function is incoherent, then it is Brier Utility dominated by some coherent function. But no coherent credence function is dominated. So, by appealing to the norm of dominance avoidance from decision theory, EUT can justify probabilism.

Epistemic utility theory has been especially fruitful in producing justifications of other standard norms as well, such as conditionalization, the Principal Principle, and so on. Indeed, these accomplishments are what maintains philosophical interest in EUT.

1.2. **Legitimate Measures of Epistemic Utility.** The obvious weakness in the argument above for probabilism is that it relies on Brier Utility as the measure of epistemic value. It's not clear that rational agents must care only about their Brier Utility—what's so valuable about squared divergence? Why not, say, cubed divergence?

As it turns out, there are many other measures of epistemic utility that will do the work EUT requires. For instance, we could have used Additive Logarithmic Utility:

$$\textsc{Additive Log Utility:}\ \mathrm{ALU}(c, w) = 1 - \sum_{X \in \mathcal{F}} \log(|1 - w(X) - c(X)|$$

The dominance argument for probabilism would have worked just as well with ALU as with Brier Utility.

So, we should look more generally at which measures of epistemic utility might count as legitimate (i.e., as measures of actual epistemic value) and that can do the work EUT needs. As noted, a fundamental commitment of EUT is a weak kind of alethic monism. That is, in general, the more accurate a credence function is—the higher its credences in truths and the lower its credences in falsehoods—the more epistemic utility it obtains at a world. We can formalize this requirement with the following Pareto condition:

[2] See Joyce (1998, 2009) or Pettigrew (2016) for a more detailed exposition of this argument.

[3] In fact the result is slightly stronger. There is no c'' that will obtain at least as much Brier Utility as c' at every world and strictly more Brier Utility at some world.

Truth-Directedness: Suppose U is a legitimate measure of epistemic utility. Let c and c' be credence functions, and w a world. Then if:
(i) for all $X \in \mathcal{F}$, $|w(X) - c(X)| \leq |w(X) - c'(X)|$, and
(ii) for some $X \in \mathcal{F}$ $|w(X) - c(X)| < |w(X) - c'(X)|$
then $U(c, w) > U(c', w)$

This condition says simply that if c always assigns credences at least as high to truths as c' and at least as low to falsehoods as c', and if for at least one truth (falsehood) c assigns a strictly higher (lower) credence than c', then c must obtain more epistemic utility than c' at that world according to any legitimate measure. In this way, truth has to trump any other desideratum when we cash out epistemic utility as accuracy.

However, there are many intuitive measures of accuracy that satisfy Truth-Directedness but must nonetheless be illegitimate according to Epistemic Utility Theory. Consider, for instance, the absolute value score and corresponding absolute value utility:

Abs Score: $\text{abs}(c, w) = \sum_{X \in \mathcal{F}} |c(X) - w(X)|$

Abs Utility: $\text{AbsU}(c, w) = 1 - \text{abs}(c, w)$

The Abs Score is very intuitive. It simply takes the absolute value of the difference between your credences and the corresponding truth-values and adds it all up. Clearly, Abs Utility satisfies Truth-Directedness. Unfortunately, however, Abs Utility will not vindicate the accomplishments of Epistemic Utility Theory.

To see why, suppose an urn has a Red ball, a Green ball, and a Blue ball, one of which will be drawn at random (i.e., each with a $1/3$ chance). An agent with credence 0 in each of the three propositions Red, Green, and Blue is guaranteed to obtain more Abs Utility than an agent with a credence of $1/3$ in each. The agent with a credence of 0 in all three will receive a total score of 0, whereas an agent with credence $1/3$ in all three will receive a score of $-1/3$.

So, according to EUT, Abs Utility must count as an illegitimate measure of epistemic utility despite the fact that it appears to be one measure of a credence function's accuracy.

1.2.1. *Propriety.* Explaining why Abs Utility and other seemingly reasonable measures of accuracy are not legitimate measures of epistemic utility is an especially vexed issue for EUT. For our purposes though, we can simply characterize abstractly what constraints must be satisfied in addition to Truth-Directedness to narrow down the field of epistemic utility measures to those that will work to vindicate EUT's justification of norms like probabilism.

The most important constraint, and the one we will be returning to below, is:

Propriety: Let U be a measure of epistemic utility. U is *proper* if for any distinct probability functions c and c', $\text{E}_c(U(c)) > \text{E}_c(U(c'))$.

That is, a measure of epistemic utility is proper if every probability function assigns itself highest expected utility.

Propriety is a necessary but insufficient additional constraint on epistemic utility measures to vindicate arguments like the one above for probabilism. So, Brier Utility is proper, but Abs Utility is not.

Normally, the challenging issue for EUT is justifying constraints on legitimate measures of epistemic utility that end up entailing PROPRIETY. While proper measures such as Brier Utility look perfectly reasonable, improper measures such as Abs Utility do too. So the epistemic utility theorist must find a way to exclude the latter from the realm of legitimacy. As we'll see later, however, once we account for relative importance of propositions, it's the proper measures that will look illegitimate.

1.2.2. *Additivity.* It's not entirely understood what exactly is required in addition to TRUTH-DIRECTEDNESS and PROPRIETY to vindicate the accomplishments of EUT. However, we do know a few seemingly reasonable constraints that are jointly sufficient to do the work.

The first is a simple differentiability property that I will take no issue with. That is, we require the score $U(c, w)$ to be differentiable in its first argument. The second is a bit more tendentious.

One intuitive way to generate a measure of *in*accuracy of an entire credence function is to add up the inaccuracy of each credence it assigns. For instance, the Brier Score of a function c at a world w is simply the sum of $(c(X) - w(X))^2$ for each X.

More generally, we can assign each credence a score of either $s^X(c(X), w)$.[4] We then say:

> **ADDITIVITY:** A measure of inaccuracy $\mathcal{I}(c, w)$ is *additive* if $\mathcal{I}(c, w) = \sum_{X \in \mathcal{F}} s^X(c(X), w)$. The corresponding measure of epistemic utility $U = 1 - \mathcal{I}$ is additive if $\mathcal{I}$ is additive.

Additivity is a nice property. It says simply that a credence function's value is nothing over and above the value of each of the credences it assigns. The value of $c(X)$ does not depend on the value of $c(X')$. Although some may dispute this constraint, it is generally assumed, holds of all the common epistemic utility functions used by proponents of EUT, and is argued for explicitly by some proponents such as Pettigrew (2016).

1.3. **Taking Stock.** It's worth pausing briefly to take stock of where we are. According to EUT, there is some space of potentially legitimate epistemic utility functions. Two non-negotiable requirements on these measures are that they are proper and truth-directed. Additionally, epistemic utility theorists maintain that on any legitimate measure, any non-probabilistically coherent function is utility-dominated by a probabilistically coherent function.

Note that for the argument for probabilism to have teeth, these utility functions must plausibly capture the overall epistemic value of a credence function at a world. Otherwise, it may be worth sacrificing some supposed epistemic 'utility' for some alternative epistemic good. Compare: in the practical case, I may show that one option is money-dominated by another. Taking one job will guarantee you less total money than another job. However, even if you like money, that argument is not enough to show you should not take the first job. You may care about other things such as leisure time or social status that could come with first job and not the second. Likewise, the dominance argument

[4] Notice here the superscript X will allow us to use different measures for different propositions.

for probabilism only works if the measures proposed by the epistemic utility theorist can capture all that is of final epistemic value.

Let's now look more in depth at how the value of accuracy can vary between propositions. From there, we'll see why a conflict emerges with EUT.

2. Epistemic Importance of Propositions

It's frequently claimed that some propositions are more 'worth knowing' than others (Talbot, 2017; Goldman, 1999; Alston, 2005). Having high accuracy in propositions concerning fundamental laws of nature is better than having high accuracy in the claim that either the nearest raven has eaten three worms today, the vase in the living-room is grue, or that I wore wool socks on January 8th, 2004.

In other words, an epistemic agent may value accuracy in an important proposition more than she values accuracy in unimportant propositions. If given the chance to have highly accurate credences in claims about my previous sartorial choices or in claims about fundamental physics, then you'd be epistemically better off if you go with the latter.

Indeed, the practices of fundamental research vindicate this view of epistemic importance. Some arcane areas of physics and mathematics (such as research into large cardinals) are not primarily driven by any practical concern. Rather, we study these things out of epistemic interest. However, there are many other curiosities of no practical import that researchers don't study, such as 'boring' mathematical theorems or the number of neutrons in the average cat.

2.1. EUT and Constant Importance.

Because EUT aims to capture features of overall epistemic value, it should be able to accommodate the view that epistemic importance can vary. Indeed, there is a straightforward and natural way to do so.

The Brier Score counts every proposition equally. That is, it just takes the squared divergence between the credence assigned to each proposition and its truth value, squares that quantity, and then adds everything up. But we can generalize the Brier Score to a family called the Weighted Quadratic Score:

Weighted Quadratic Score: $$\mathrm{WQS}_\lambda(c, w) = \sum_{X \in \mathcal{F}} \lambda(X)(c(X) - w(X))^2$$

where $\lambda(X) > 0$ encodes the relative importance of the accuracy of c's credence in X. In turn, the epistemic utility of c at w is given by $1 - \mathrm{WQS}(c, w, \lambda)$.

So far, so good. WQS is truth-directed, proper, additive, and so on. So, it meets all the requirements we've laid down so far for epistemic utility theory. Because it allows for some propositions to count more toward epistemic value than others, it is in fact a better candidate for a reasonable epistemic utility function than the Brier Score is.

2.2. Contingent Importance.

Unfortunately, the WQS assigns *constant* weight to each proposition. If X is important in one world, then it's just as important in every other world.

However, once we admit that importance can vary between propositions, it seems that such importance is often a contingent matter. How much X matters in one world might be different, I claim, from how much it matters in another world.

There are two possible ways importance could be contingent. The first is if the proposition X takes on different levels of importance in worlds where it's true from worlds where it's false. For example, it may be that your credence that there are exactly eight planets in the solar system is important in worlds where there are in fact eight planets, but relatively unimportant in worlds where there are 10,000 planets.[5] As it turns out, Epistemic Utility Theory can handle this kind of case, though in a rather complicated way.[6]

The second way is a little more complicated, but it is what will cause the basic problem for EUT. Sometimes the importance of one proposition depends on the truth-value of another proposition. Consider, for instance, the following cases:[7]

Fundamental Laws: In different worlds, there are different fundamental laws of nature. In one world, the speed of light plays an especially important theoretical role. In another, it plays no fundamental role whatever. So, having accurate credences about the speed of light matters more in the former types of worlds than in the latter types.

Years of LBJ: Lyndon Johnson shaped the course of American history in the mid-twentieth century both with his presidency and with his transformation of the US Senate. It was therefore worthwhile for Robert Caro to write a detailed biography of LBJ totaling many thousands of pages and eventually five volumes. However, LBJ almost lost his second bid for the US Senate, which would have ended his political career. Had he lost, it would not be worth knowing five volumes' worth of material about life.

Effective Altruism: Dave wants to know all about the most important problem facing humanity. In one world, humans are at grave risk of extinction because of runaway super-intelligent AI. In another world, humans are teetering on catastrophe due to global warming. In yet another, famine is the biggest drivers of unhappiness. In each world, Dave would most want accurate credences in different types of propositions—facts about machine learning, the carbon cycle, and crop yield take on massively different levels of importance.

Pop Music: Bill is obsessed with popular music, but he's also terribly elitist. He wants to know everything about the lives of the singers who are actually the most talented musicians. As it turns out, in one world,

[5]Thanks to xxx for this example.

[6]See (Merkle and Steyvers, 2013) for the details. In brief, they use a theorem from (Schervish, 1989), according to which one can construct a scoring rule for a credence x in X as follows. Let $s(x,1) = \int_x^1 (1-t)\omega(t)\,\mathrm{d}t$ and $s(x,0) = \int_0^x t\omega(t)\,\mathrm{d}t$, where $\omega(t)$ is positive, finite, and continuous over $(0,1)$. Here, $s(x,i)$ is the score for a credence of x when X's truth-value is i. By setting $\omega(t) = t^{\alpha-1}(1-t)^{\beta-1}$ for $\alpha, \beta > -1$, we can construct asymmetric scoring rules by letting $\alpha \neq \beta$. For instance, if $\alpha = 1$ and $\beta = 3$, then $s(.25,1) \approx s(.25,0)$ despite the fact that .25 is normally considered more accurate in worlds in which the proposition is false.

[7]In the first two of these cases, the epistemic importance of some claims is taken to be objective. The importance of the proposition varies from world to world for all rational agents. In the third and fourth cases, the epistemic importance is subjective. Some propositions count more to Dave and Bill, even if they don't count more for everybody. But how much utility Dave and Bill receive for their credences will vary depending on the world.

Whether epistemic importance should always be objective or subjective or a mix does not matter for our purposes here. No matter the view, importance is often contingent.

Beyoncé meets the cut, but in a very distant one, she doesn't. In only some worlds, Bill places high importance on knowing where Beyoncé was born, the name of her high school, and the sales figures of her first album.

In each of these cases the importance of some propositions varies depending on what else is true. If, for instance, it's true that LBJ was president, then other facts about him matter a lot. Likewise, if AI is an actual threat to humanity, then other propositions about neural networks matter a lot more than in worlds where AI is no threat. Let's see now what goes wrong.

3. What Goes Wrong

Since the actual utility an agent will receive from an accurate credence in X will vary depending on which world she's in, we should account for this fact in our value function. Most naturally, we can change the Weighted Quadratic Score to the Contingent WQS as follows:

Contingent WQS: $$\mathrm{CWQS}_{\lambda}(c, w) = \sum_{X \in \mathcal{F}} \lambda(X, w)(c(X) - w(X))^2$$

The Contingent WQS is a lot like the WQS with one change. Instead of λ being a function just of the proposition, it's now a function of both the proposition and world. So, $\lambda(X, w)$ is high in worlds where accuracy in X matters a lot but low in worlds where it doesn't. For instance, in worlds where LBJ is president, facts about his boyhood relationship with his father that shaped his psychological makeup achieve high λ value in worlds where LBJ was president but low λ value in worlds where he never left the Texas Hill Country.

3.1. The Problem.

CWQS seems to be even better than the WQS as a candidate for quantifying actual epistemic value. It's also proper and additive, just as the Epistemic Utility Theorist wants. Unfortunately, it is not proper.

More generally, there is no additive, differentiable, truth-directed, and proper epistemic utility function that handle contingent importance in general as is shown in the Appendix.[8]

So, we now see the problem. Epistemic Utility Theorists want measures of accuracy that capture the overall epistemic value of an agent's credences. Moreover, they want these measures to vindicate the main accomplishments of EUT, such as the standard arguments for probabilism, conditionalization, and so on. We now turn to some options for EUT.

4. Ways Out

Since Epistemic Utility Theory is predicated on alethic monism, truth-directedness is non-negotiable. Giving up on propriety is also not an option for a number of reasons, the most important of which is that improper measures won't vindicate probabilism. So, something else must give way if EUT is to be successful.

The first way out for the Epistemic Utility Theorist is to deny that propositions vary in importance across worlds. As we've seen, though, this is implausible. The reason it's worth knowing about the interaction of bosons and fermions in

[8]For a somewhat related problem for EUT and accompanying result in the context of truth-likeness, see (Oddie, 2017).

particle physics is that this is a world with bosons and fermions that interact. It would not be worth knowing if no such particles existed.

A second way out is to deny additivity. On this view, the value of your credence in X may depend also on how accurate your credence in some other proposition Y is. Such a view is not intrinsically implausible, but there are some problems. First, this move is that it is a major departure from orthodoxy. The favorite epistemic utility functions are all additive. So, giving up additivity does not seem like a move to make lightly.

Next, it's unknown whether there really are any non-additive options available. Although the result proved below holds just for additive rules, the Epistemic Utility Theorist must find a non-additive rule that both assigns contingent importance and vindicates the major accomplishments of EUT such as probabilism and so on.

Last, it's unclear whether this move is actually well-motivated. Suppose Fuka-Eri has credence 1 that LBJ was born in Texas. It seems that in the world in which he is president this simply counts more toward her epistemic well-being than in the world in which he is not president, regardless of what else Fuka-Eri thinks.

I tentatively conclude, then, that EUT cannot capture all facts about epistemic value in an attractive way while maintaining its commitment to propriety and truth-directedness. EUT thus fails in its aims to capture the entirety of the epistemic value of an agent's credence function.

Appendix A. Proof of Main Result

A scoring rule is proper if, for every probability, that probability function expects itself to be least inaccurate. For the proof of the main result, it will be useful to define a stronger property:

Definition A.1 (Strong Propriety). An additive scoring rule of the form $\mathcal{I}(c, w) = \sum s^X(c(X), w)$ is *strongly proper* if $\mathrm{E}_{\mathrm{Pr}}(s^X(x))$ is uniquely minimized at $x = \mathrm{Pr}(X)$ for all $X \in \mathcal{F}$ and all probability functions Pr.

That is, a scoring rule is strongly proper, if it expects each of its credences to be least inaccurate. In other words, a rule is strongly proper if probability functions think of themselves as doing locally best.

As we see, however, this doesn't make much of a difference in our case:

Lemma A.2. *If $\mathcal{I}$ is a proper, additive, differentiable, and truth-directed scoring rule, then $\mathcal{I}$ is strongly proper.*

Proof. Consider the gradients. □

To capture contingent importance in the cases in the main text, we noted that a scoring rule I must assign different weights to some propositions depending on the truth-values of others. We then make the following definition:

Definition A.3 (Recognizing contingent importance). Let $\mathcal{I} \sum \lambda(X, w) s^X(c(X), w)$ be an additive scoring rule. $\mathcal{I}$ *recognizes contingent importance* if there exist worlds w and w' such that $w(X) = w'(X)$ but $\lambda(X, w) \neq \lambda(X, w')$.

We can now prove:

Theorem A.4. *There does not exist a scoring rule $\mathcal{I}$ that is additive, strictly proper, truth-directed, and differentiable, such that $\mathcal{I}$ recognizes contingent importance.*

Proof. For the *reductio*, suppose $\mathcal{I}$ is additive, strictly proper, truth-directed, differentiable, and that $\mathcal{I}$ recognizes contingent importance. First, note that by Lemma A.2, $\mathcal{I}$ must be strongly proper.

Let $\mathcal{I}(c, w) = \sum \lambda(X, w)s(c(X), w)$. Fix a probability function Pr and proposition X such that $0 < \Pr(X) < 1$. This results in no loss of generality since we're assuming that $\mathcal{I}$ is proper. Consider the function:

$$f_{\Pr}(x) := \mathrm{E}_{\Pr}(\lambda(X)s(x))$$

That is, $f_{\Pr}$ is just the expected inaccuracy of a credence of x in X according to probability function Pr.

We then have:

$$\begin{aligned} f(x) &= \sum_{w \in \Omega} \Pr(w)\lambda(X, w)s(x, w) \\ &= \sum_{w(X)=1} \Pr(w)\lambda(X, w)s(x, 1) + \sum_{w(X)=0} \Pr(w)\lambda(X, w)s(x, 0) \end{aligned} \tag{1}$$

Using equation (1) to take the derivative of f, we get:

$$\frac{df_{\Pr}}{dx} = s'(x, 1) \sum_{w(X)=1} \lambda(X, w)\Pr(w) + s'(x, 0) \sum_{w(X)=0} \lambda(X, w)\Pr(w) \tag{2}$$

To find the minimum, we set equation (2) to 0. Doing some simple algebraic manipulations, we see the minimum is achieved at x_0 only if:

$$\frac{s'(x_0, 1)}{s'(x_0, 0)} = -\frac{\sum_{w(X)=0} \lambda(X, w)\Pr(w)}{\sum_{w(X)=1} \lambda(X, w)\Pr(w)} \tag{3}$$

By assumption there are worlds w_1 and w_2 such that $w_1(X) = w_2(X)$, but $\lambda(X, w_1) \neq \lambda(X, w_2)$. In this case, pick a probability function $\Pr'$ such that:

$$\Pr'(w) = \begin{cases} \Pr(w_1) + \varepsilon & w = w_1 \\ \Pr(w_2) - \varepsilon & w = w_2 \\ \Pr(w) & \text{Otherwise} \end{cases}$$

where $\varepsilon < \min(\Pr(w_1), \Pr(w_2), 1 - \Pr(w_1), 1 - \Pr(w_2))$ to guarantee that $\Pr'$ is a probability function.

Note that $\Pr'(X) = \Pr(X)$, so, if $\mathcal{I}$ is proper, $f_{\Pr}$ is minimized at the same point as $f_{\Pr'}$. So, from equation (3), we then have a minimum at x_0 only if:

$$\frac{s'(x_0, 1)}{s'(x_0, 0)} = -\frac{\sum_{w(X)=0} \lambda(X, w)\Pr(w)}{\sum_{w(X)=1} \lambda(X, w)\Pr(w)} = -\frac{\sum_{w(X)=0} \lambda(X, w)\Pr'(w)}{\sum_{w(X)=1} \lambda(X, w)\Pr'(w)} \tag{4}$$

However, this is impossible by the construction of $\Pr'$. Either the denominators of the two right hand terms are equal with unequal numerators, or the numerators are equal with unequal denominators. In either case, the two right hand terms are not equal. So, we have a contradiction. Therefore, for all worlds w_1 and w_2 such that $w_1(X) = w_2(X)$, we have $\lambda(X, w_1) = \lambda(X, w_2)$. □

References

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

Goldman, A. (1999). *Knowledge in a Social World*. Oxford University Press.

Greaves, H. and D. Wallace (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind 115*(632), 607–632.

Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science 65*, 575–603.

Joyce, J. M. (2009). Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In F. Huber and C. Schmidt-Petri (Eds.), *Degrees of Belief*, Volume 342, pp. 263–297. Springer.

Merkle, E. C. and M. Steyvers (2013). Choosing a strictly proper scoring rule. *Decision Analysis 10*(4), 292–304.

Oddie, G. (2017). What accuracy could not be. *The British Journal for the Philosophy of Science Online first*.

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

Schervish, M. (1989). A general method for comparing probability assessors. *The Annals of Statistics 17*, 1856–1879.

Talbot, B. (2017). Repugnant accuracy. *Nous 00*(00), 1–24.