Pettigrew, R. G. (2013). Epistemic Utility and Norms for Credence. *Philosophy Compass*, 8(10), 897-908.
<https://doi.org/10.1111/phc3.12079>

Peer reviewed version

License (if available):
CC BY-NC

Link to published version (if available):
[10.1111/phc3.12079](https://doi.org/10.1111/phc3.12079)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the peer reviewed version of the following article: Pettigrew, R. (2013), Epistemic Utility and Norms for Credences. *Philosophy Compass*, 8: 897–908. , which has been published in final form at doi: 10.1111/phc3.12079. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Epistemic Utility and Norms for Credences (forthcoming in *Philosophy Compass*)

Richard Pettigrew

July 2, 2013

Like the rest of us, Paul's beliefs come in degrees. Some are stronger than others. In particular, Paul believes that Linda is a bank teller and a political activist more strongly than he believes that she is a bank teller. That is, his credence in the former proposition is greater than his credence in the latter. Surely, Paul is irrational. But why?¹ In this survey, I describe a new strategy for answering such questions. It is a strategy that was first introduced by Jim Joyce [Joyce, 1998].

The traditional strategy—the strategy that Joyce sought to replace or, at least, supplement—is to show that such credences will lead the agent who has them to make decisions that are guaranteed to have a bad outcome. These are the well-known Dutch Book arguments.² For instance, Paul's credences will lead him to buy a book of bets on the two propositions concerning Linda that is guaranteed to lose him money. This, it is claimed, makes him irrational. Now, the validity of this argument has been the subject of much debate. However, even if it works, it only identifies one way in which Paul's credences are irrational: they are poor guides to action; from a pragmatic point of view, they are irrational. But, intuitively, there is something irrational about these credences from a purely epistemic point of view; they seem to exhibit a purely epistemic flaw. Even for an agent incapable of acting on her credences—and therefore incapable of making the bets that lead to the guaranteed loss—Paul's credences would be irrational. We will be concerned with identifying why that is so. That is, Joyce's strategy, which we describe here, provides a purely epistemic route to the norms that govern credences; this route does not rely on any connection between credence and action.

¹The example is taken from [Kahneman and Tversky, 1983, 297], where Kahneman and Tversky showed that the vast majority of subjects, when presented with information about Linda's enthusiasm for student politics, will form credences like Paul's. They dub this phenomenon the *conjunction fallacy*.

²These were given originally in [Ramsey, 1931] and [de Finetti, 1931]. For recent surveys of the literature that has grown up around these arguments, see [Hájek, 2008] and [Vineberg, 2011].

We will begin by showing how the strategy works in the case of Paul. Then we will show how to extend it to establish Probabilism, which is the norm that Joyce considers in his original paper. Probabilism is one of the core tenets of so-called Bayesian epistemology. Our next target is the other core tenet of that view, namely, Conditionalization. We will give an argument for that norm that uses Joyce's strategy as well: it is due to Hilary Greaves and David Wallace [Greaves and Wallace, 2006]. After considering how we might strengthen these arguments by weakening the assumptions they make, we conclude by describing possible avenues for future research.

1 Why Paul is irrational

First, let:

- $A = \text{Linda is a bank teller and a political activist}$
- $B = \text{Linda is a bank teller}$

Clearly, A entails B . Next, represent Paul's epistemic state by his *credence function* c_{Paul} . This is the function that takes each of the propositions about which he has an opinion and returns the real number in the closed unit interval $[0, 1]$ that gives his credence in that proposition.³ Suppose Paul has opinions only about A and B . Thus, c_{Paul} is defined only on A and B . Moreover, we have $c_{\text{Paul}}(B) < c_{\text{Paul}}(A)$, since Paul believes A more strongly than he believes B . We wish to show that c_{Paul} is irrational.

Our argument relies on an application of standard decision theory in epistemology. In standard decision theory, an agent is faced with a set of options; and decision theory provides principles that rule out some of those options as irrational. For instance, the agent might be faced with the following options: holiday in Delhi; holiday in Nairobi; holiday in Istanbul. A decision principle will rule out some (possibly empty) subset of this set of options as irrational for that agent. Different decision principles require different ingredients. Some require only the agent's *utility function*, which takes each of the options and each of the ways the world might be and assigns a numerical value that measures how much the agent would value the outcome of that option if the world were that way. Thus, the utility function of an agent deciding on her holiday destination would measure how much she would value going to Delhi if it is raining in Delhi; how much she would value going to Delhi if it is sunny in Delhi; and so on. And a decision rule that required only her utility function might rule out Delhi as irrational for an agent for whom the utility of going to Istanbul is greater than the utility of going to Delhi regardless of how the world turns out

³Recall: $[0, 1] = \{x \in \mathbb{R} : 0 \leq x \leq 1\}$.

to be.⁴ Other decision principles require the agent's utility function along with some facts about her epistemic state. For instance, such a principle might require her state of knowledge about the objective chances of certain events occurring; or it might require her credences in certain propositions. But, while the input of these decision principles may vary, the output is the same: each rules out a subset of the options as irrational.

The argument strategy that Joyce introduced, and which we will study in this paper, applies this simple version of decision theory in epistemology. The idea is this: Treat the set of possible epistemic states of a particular kind as the set of options; and apply the principles of decision theory to rule out some of those options as irrational. To do this, of course, we will need an *epistemic utility function*: that is, a function that takes each epistemic state of the particular kind in question, and each possible way the world might be, and returns a measure of the epistemic value that the epistemic state would have if the world were that way. In this section and the next, the kind of epistemic state we are interested in is the kind represented by an agent's credence function—we might call it her *credal state*. So our epistemic utility function takes an agent's credence function, together with a possible world, and returns a measure of the epistemic goodness of that credence function at that world. In Section 3, we will consider *updating rules*: these are rules that purport to say how a rational agent updates her credences in the light of new evidence. So, in that section, our epistemic utility function will take an agent's updating rule, together with a possible world, and return a measure of the epistemic goodness of adopting that rule at that world. Having defined our epistemic utility function and identified our decision principle, we prove a mathematical theorem that characterises the epistemic states that are ruled out as irrational by the decision principle when it is applied to the epistemic utility function we have defined. And we take this to establish the norm that says an agent shouldn't have such an epistemic state. Thus, the argument strategy has the following general structure:

- (I) Define an epistemic utility function.
- (II) Pick a principle of decision theory.
- (III) Prove a mathematical theorem.

In the remainder of this section, we present the epistemic utility argument against Paul, which has precisely this structure.

⁴Aficionados will note that this is a very simple version of decision theory. We will consider the behaviour of these arguments in the presence of a more complicated version of decision theory in Section 5.

1.1 Stage (I): the epistemic utility function

We begin by defining the most popular measure of epistemic utility for credence functions. It is called the *Brier score*. We will use this measure throughout the main body of the paper. At this stage, I do not offer any argument that the Brier score is the only legitimate way of measuring epistemic utility, nor even that it is one of the legitimate ways. That will be the concern of Section 4, where we will strengthen the argument given in the main body of the paper by showing that they can be run using any one of a range of different measures of epistemic utility.

The Brier score is introduced in two stages. The idea is that the epistemic utility of a credence function at a world ought to be given by its ‘proximity’ to the credence function that is ‘perfect’ or ‘vindicated’ at that world—that is, the credence function that perfectly matches whatever it is that credence functions strive to match. Thus, we need to identify, for each world, the ‘perfect’ or ‘vindicated’ credence function; and we need to identify the correct measure of distance from each of these credence functions to another credence function.

(Ia) Identify the vindicated credence functions.

In the first stage, we identify, for each possible world, the ‘vindicated’ credence function at that world. This is the credence function that should be awarded maximum epistemic utility at that world; it is the credence function that would be the best, epistemically speaking, for an agent to have at that world. Joyce identifies it as follows:

$$v_w(X) := \begin{cases} 0 & \text{if } X \text{ is false at } w \\ 1 & \text{if } X \text{ is true at } w \end{cases}$$

That is, an agent has greatest epistemic utility if she has maximal credence in truths and minimal credence in falsehoods. This, Joyce contends, is the correct analogue of the thesis from traditional epistemology that belief aims at the truth [Joyce, 1998, 577-9].

(Ib) Define the distance measure.

In the second stage, we identify a measure of the distance from a vindicated credence function v_w to a credence function c . Here is one candidate amongst many that Joyce endorses:

$$d(v_w, c) := \sum_X |v_w(X) - c(X)|^2$$

where the sum ranges over all the propositions X to which v_w and c assign credences. Thus, the distance between a vindicated credence function and another credence function is the sum of the squared differences between the credences they assign. Throughout, we assume

that an agent has credences in only finitely many propositions. On this assumption, the distance function d is well-defined.

Putting (Ia) and (Ib) together with the thesis that the epistemic utility of a credence function at a possible world is its proximity to the credence function that is vindicated at that world, we have the Brier score:

$$B(c, w) := 1 - d(v_w, c) = 1 - \sum_X |v_w(X) - c(X)|^2$$

1.2 Stage (II): the principle of decision theory

So we have a measure of epistemic utility. Now we need a principle that allows us to appeal to this measure in order to rule out certain credence functions as irrational. We will appeal to the most basic such principle. To state it, we require some terminology: Suppose $\mathcal{A}$ is a set of options and U is a utility function. If a, a' in $\mathcal{A}$, then we say that a is *dominated by a' relative to U* if $U(a, w) > U(a', w)$, for all worlds w .

(Dominance) Suppose that:

- a is dominated by a' relative to U ;
- a' is not dominated by any a'' relative to U .

Then a is irrational for an agent with utility function U .

Notice that Dominance is not stated as a choice principle. We need not assume that an agent is making a free choice between the options in $\mathcal{A}$ in order to use Dominance to evaluate a particular option as irrational. Thus, even though it seems likely that we have no volitional control over the credences we in fact have, our credences may nonetheless be evaluated as irrational by appealing to Dominance.

1.3 Stage (III): the mathematical theorem

The final step in our argument brings together (I) and (II) to establish Paul's irrationality. We can prove that, if we apply the principle of decision theory introduced in (II) to the epistemic utility function introduced in (I), Paul's credence function will belong to the set that is ruled out as irrational. This is best seen by considering Figure 1.

This completes our account of Paul's irrationality. In sum: Suppose we measure epistemic utility as proximity to vindication (I). And suppose further that the vindicated credence function assigns maximal credence to all truths and minimal credence to all falsehoods (Ia); while the distance between a vindicated credence function and another credence function is given by sum of the squares of the differences between the credences they assign (Ib). Then there is a credence function c' that has greater epistemic

Figure 1: Given a credence function c that is defined only on propositions A and B , we represent c by the ordered pair $(c(A), c(B))$. We note that the greater the Euclidean distance between ordered pairs, the greater the sum of squared differences between the corresponding credence functions, and vice versa. Now, let w_1, w_2, w_3 be the three possible worlds: at w_1 , A and B are both false; at w_2 , A is false and B is true; at w_3 , both are true; there is no fourth possible world at which A is true and B is false, since A entails B . And let v_{w_i} be the credence function that is vindicated at w_i . Now, the shaded area contains precisely those credence functions c such that $c(A) \leq c(B)$. It is easy to see that, for any credence function c that lies outside that set, there is a credence function c' that lies in that set that lies closer to each of v_{w_1} , v_{w_2} , and v_{w_3} . For any such c , we choose c' to be the credence function in the shaded area that is closest to c . The dark blue dashed lines illustrate their respective distances from v_{w_3} , in particular, but analogous triangles can be drawn in the case of v_{w_1} and v_{w_2} . And, conversely, given a credence function c' that lies in the shaded area, there is no credence function c that is closer to each of v_{w_1} , v_{w_2} , and v_{w_3} . This shows that the credence functions that are not dominated are precisely those that lie in the shaded area. Thus, it is irrational to have a credence function c such that $c(A) > c(B)$. In particular, c_{Paul} is irrational.

utility than c_{Paul} at every possible world; but there is no credence function c'' that has greater epistemic utility than c' at every possible world. Therefore, by Dominance (II), Paul is irrational.

2 Probabilism

Together, (I) and (II) above establish the following norm that governs credences:

(No Drop) Suppose X entails Y . And suppose our agent's credence function c is defined on X and Y only. Then our agent's credence function is irrational if $c(Y) < c(X)$.

Moreover, we can see that, if X entails Y , then no stronger norm governing an agent with credences only in X and Y can be established by these means. Any stronger norm would have to rule out further credence functions; but we saw above that Dominance does not sanction this—no credence function that satisfies No Drop is dominated. Suppose we now turn from an agent with credences in only two propositions, one of which entails the other, to an agent with credences in all propositions in some full finite algebra.⁵ What, then, is the strongest norm that can be established by appealing to the Brier score and Dominance? The answer is Probabilism. First, a definition: A credence function c is a *probability function* if

(Normalization)

- If X is a tautology, then $c(X) = 1$;
- If X is a contradiction, then $c(X) = 0$.

(Additivity) If X and Y are mutually exclusive, then $c(X \vee Y) = c(X) + c(Y)$.

Now we can state Probabilism:

(Probabilism) An agent's credence function is irrational if it is not a probability function.

As the following theorem shows, Probabilism can be established using exactly the same strategy as No Drop:

Theorem 1 (de Finetti)⁶

(I) If c violates Probabilism, then there is c' that satisfies Probabilism such that $B(c, w) < B(c', w)$ for all worlds w .

⁵That is, the agent has credences in finitely many propositions, and whenever she has credences in two propositions X and Y , she also has credences in $\neg X$, $\neg Y$, and $X \wedge Y$.

⁶This theorem is due to Bruno de Finetti [de Finetti, 1974, 87-91]; its philosophical gloss is due to Joyce.

(II) If c satisfies Probabilism, then there is no c' such that $B(c, w) < B(c', w)$ for all worlds w .

Thus, if we measure epistemic utility by the Brier score, any agent who violates Probabilism is dominated, while any agent who satisfies it is not. This provides an epistemic utility argument for Probabilism. It is not yet Joyce's argument for Probabilism, since Joyce permits other measures of epistemic utility besides the Brier score. We will see how to strengthen the argument in this way in Section 4. In the meantime, we turn to another important norm that can be justified using Joyce's strategy.

3 Conditionalization

Probabilism tells us how rational credences at a particular time relate to one another: a rational agent's credence in an exclusive disjunction will be the sum of her credences in the disjuncts, etc. Our next norm tells us (roughly) how a rational agent's credences at one time relate to her credences at a later time. More precisely, it identifies the unique rational *rule* for updating our credences in the light of new evidence obtained between the earlier and later times. It is called *Conditionalization*.

Just as there is a pragmatic argument for Probabilism—the Dutch Book argument—so is there one for Conditionalization—the diachronic Dutch Book or Dutch Strategy argument [Lewis, 1999]. Again, of course, whether or not this argument is valid, it does not help us to identify the epistemic failing of an agent who plans to update other than by this rule. The argument we will consider, by contrast, does identify that failing.

Before we state this norm, we need to say what we mean by an *updating rule*. An updating rule tells an agent how she should respond to evidence that she might acquire. Formally, it is a function $\mathbf{R}$ that takes a credence function c , a partition $\mathcal{E}$ (i.e., a set of exhaustive and mutually exclusive propositions), and a proposition E in that partition $\mathcal{E}$, and returns a credence function $\mathbf{R}(c, \mathcal{E}, E)$. Suppose our agent's current credence function is c ; and her next piece of evidence will be given by some proposition in the partition $\mathcal{E}$. Then $\mathbf{R}(c, \mathcal{E}, E)$ is the credence function that the rule demands the agent adopt if the proposition from $\mathcal{E}$ that she learns is E .

For instance, suppose Paul and Linda are playing a game. Linda thinks of a famous person and, by asking 'yes/no' questions, Paul tries to guess who it is. On the basis of her answers to his previous questions, Paul is sure Linda is thinking of Lillian Ngoyi, Steve Biko, or Mahatma Gandhi. His next question is: Was the person born in South Africa? Thus, the partition from which he knows his evidence will come is $\{Ngoyi \vee Biko, Gandhi\}$. Thus, an updating rule must assign to each of the two members of that partition the credence function it requires Paul to adopt if he learns that proposition.

We define the standard Bayesian updating rule as follows: Let **Cond** be the updating rule:

$$\mathbf{Cond}(c, \mathcal{E}, E) = c(\cdot | E) := \frac{c(\cdot \wedge E)}{c(E)}$$

And we state the norm as follows:

(Conditionalization) All updating rules other than **Cond** are irrational.

That is, the only updating rule that Conditionalization does not rule out as irrational is **Cond**, which requires that an agent take her new credences to be her old credences conditional on the evidence she has acquired in the meantime.

How might we justify Conditionalization? Again, we can appeal to considerations of epistemic utility. And again, our argument has two stages: in the first, we use our epistemic utility function over credence functions to define a function that measures the epistemic utility of an updating rule at a world; in the second, we identify a principle of decision theory. As before, we bring together these two stages by applying the decision-theoretic principle to the epistemic utility function for updating rules; we prove that this rules out all updating rules except **Cond** as irrational. This style of argument was first given explicitly by Hilary Greaves and David Wallace [Greaves and Wallace, 2006].

3.1 Stage (I): the epistemic utility function

Let us continue to assume, as we did in our argument for Probabilism, that the epistemic utility of a credence function c at world w is given by the Brier score $B(c, w)$ of c at w . Then we can use this to define the epistemic utility of an updating rule $\mathbf{R}$ at a world w , when $\mathbf{R}$ is applied to a credence function c and a partition $\mathcal{E}$: it is $B(\mathbf{R}(c, \mathcal{E}, \mathcal{E}_w), w)$, where $\mathcal{E}_w$ is the element of $\mathcal{E}$ that is true at world w . That is, given one's current credence function and a partition from which one's future evidence will come, the epistemic utility of an updating rule at a particular world is just the Brier score, at that world, of the new credence function that the updating rule will demand if one learns the proposition in the partition that is true at that world. Consider our example; and consider the world at which Linda is in fact thinking of Steve Biko. Then, for Paul, the epistemic utility of an updating rule at that world is just the Brier score of whichever credence function that updating rule requires of Paul when Linda answers 'yes' to his question and thus provides him with the evidence $Ngoyi \vee Biko$.

3.2 Stage (II): the principle of decision theory

Next, consider the following well known principle of decision theory. Again, we need a little terminology in order to state it: Suppose $\mathcal{A}$ is a set of options, U is a utility function, and c is a probability function. If a in $\mathcal{A}$, then we define the *expected utility of a relative to U and c* as follows:

$$\text{Exp}_{U,c}(a) = \sum_w c(\bar{w})U(a, w)$$

where $\bar{w}$ is the proposition that is true at world w and only at world w .⁷

(Maximize Expected Utility) Suppose that:

- $\text{Exp}_{U,c}(a) < \text{Exp}_{U,c}(a')$; and
- $\text{Exp}_{U,c}(a'') \leq \text{Exp}_{U,c}(a')$ for all a'' in $\mathcal{A}$.

Then a is irrational for an agent with utility function U and credence function c .

This says that, when there are options with maximal expected utility relative to one's credence function and utility function, it is irrational to choose any other options.

3.3 Stage (III): the mathematical theorem

As before, we now apply the decision-theoretic principle identified in (II) to the epistemic utility function identified in (I). Together, they entail Conditionalization, as the following theorem shows:

Theorem 2 (Greaves and Wallace) ⁸ Suppose $\mathcal{E}$ is a partition and c is a probabilistic credence function. Then, for any updating rule $\mathbf{R} \neq \mathbf{Cond}$,

$$\sum_w c(\bar{w})B(\mathbf{R}(c, \mathcal{E}, \mathcal{E}_w), w) < \sum_w c(\bar{w})B(\mathbf{Cond}(c, \mathcal{E}, \mathcal{E}_w), w)$$

Thus, if a probabilistic agent were to plan to update her credences by any rule other than **Cond**, she would not be maximising her expected epistemic utility relative to her current credence function and the Brier score for updating rules. That is, she would expect the outcome of her alternative updating rule to have less epistemic utility than she would expect the outcome of **Cond** to have. This is our epistemic utility argument for Conditionalization. Again, it is not quite Greaves and Wallace's, since it assumes that epistemic utility is measured by the Brier score. Again, in Section 4,

⁷One might think of $\bar{w}$ as the state description corresponding to w . It is guaranteed to exist by the fact that the algebra on which c is defined is full and finite.

⁸For the proof of this theorem, see [Greaves and Wallace, 2006, 624].

we will describe ways to weaken this assumption and thereby strengthen the argument.

One might worry that Conditionalization is not exactly the updating norm endorsed by Bayesians. Conditionalization assumes that an updating rule is a function that takes as its input a credence function, a partition, and a proposition in that partition. And one might worry that it should take only a credence function and a proposition. After all, an agent often does not know the partition from which her evidence will come; and even after she has acquired that evidence, she often does not know which other propositions she might have learned. Nonetheless, Conditionalization does entail that, for every partition $\mathcal{E}$ to which the proposition E belongs, a rational agent whose evidence will come from $\mathcal{E}$ will plan to adopt the credence function $c(\cdot|E)$ if she learns E . Perhaps this is sufficiently close to the norm the Bayesians endorse.

4 Other epistemic utility functions

Our epistemic utility arguments for Probabilism and Conditionalization have relied on the assumption that there is only one way to measure epistemic utility, namely, the Brier score. Is this assumption justified? Can we strengthen our arguments by weakening this assumption?

Hannes Leitgeb and Richard Pettigrew argued that the Brier score is the only legitimate measure of the epistemic utility of a credence function [Leitgeb and Pettigrew, 2010]; and Reinhard Selten argued that it is the only correct measure at least in the case of probabilistic credence functions [Selten, 1998]. But their arguments rely on strong geometrical assumptions that have yet to be well motivated. Joyce, by contrast, is more permissive. He enumerates a set of properties and claims that a function must have at least these if it is to count as a legitimate measure of epistemic utility [Joyce, 1998, Section 4]. He then shows that Theorem 1 holds for any function that boasts the properties he lists. Thus, any legitimate epistemic utility function can be used to justify Probabilism. Greaves and Wallace take a similar approach [Greaves and Wallace, 2006, Section 3.1]. In fact, the properties they endorse are of greater interest, since they constitute the weakest known properties that suffice for analogues to Theorems 1 and 2, the mathematical results that underpin our arguments for Probabilism and Conditionalization. Let's consider them briefly.

Definition 1 *An epistemic utility function S is separable if, for each proposition X , there is a function $s_X : \{0, 1\} \times [0, 1] \rightarrow [0, \infty]$ such that*

$$S(c, w) := 1 - \sum_X s_X(v_w(X), c(X))$$

We might think of $s_X(1, x)$ as the epistemic utility of having credence x in X at a world at which X is true, while $s_X(0, x)$ is the epistemic utility of having credence x in X at a world at which X is false. Note that, if $s_X(1, x) = |1 - x|^2$ and $s_X(0, x) = |x|^2$, then S is the Brier score.

Definition 2 *An epistemic utility function S is proper if, for any probabilistic credence function c and credence function $c' \neq c$,*

$$\sum_w c(\overline{w})S(c', w) < \sum_w c(\overline{w})S(c, w)$$

That is, S is proper if any probabilistic credence function expects itself to be better than it expects any other credence function to be. Then:

- Predd, et al. show that Theorem 1 holds for any separable, proper epistemic utility function [Predd et al., 2009].
- Greaves and Wallace show that Theorem 2 holds for any proper epistemic utility function [Greaves and Wallace, 2006].

These are very strong results. They permit many different measures of epistemic utility: the Brier score, as well as the so-called logarithmic scoring rule, the spherical scoring rule, and many others.⁹ Do they strengthen our arguments for Probabilism and Conditionalization? In the case of Conditionalization, they certainly do; in the case of Probabilism, we have to be a bit careful. Let's see why.

Suppose our agent doesn't know what her epistemic utility function is; or suppose that she simply does not have sufficiently determinate attitudes to fix a unique such function. Nonetheless, she knows from some generalization of Theorem 1 that whichever legitimate epistemic utility function she has will make any non-probabilistic credence function dominated. Even so, it doesn't follow that any such credence function is irrational for her. After all, she may be in the following situation, which is not ruled out by the generalization of Theorem 1:

- S_1 and S_2 are legitimate epistemic utility functions by the lights of this generalization of Theorem 1.
- c is a credence function that violates Probabilism
- For every probabilistic c_1 that dominates c relative to S_1 , we have $S_2(c_1, w) \ll S_2(c, w)$ for some world w .
- For every probabilistic c_2 that dominates c relative to S_2 , we have $S_1(c_2, w) \ll S_1(c, w)$ for some world w .

⁹See [Joyce, 2009, Section 8] for a survey of separable, proper epistemic utility functions.

In this situation, it would not be irrational for our agent to stick with credence function c . After all, while moving to a credence function c_1 that dominates c with respect to S_1 is guaranteed to serve her better by the lights of S_1 , it risks serving her very badly by the lights of S_2 . And recall that she doesn't know whether S_1 or S_2 reflects her epistemic values. Moreover, a similar argument shows that it is not mandatory to move to any credence function c_2 that dominates c with respect to S_2 .

This objection was first raised by Aaron Bronfman in an unpublished note on Joyce's original paper [Bronfman, ms]. To block it, we must assume that, while there are many legitimate epistemic utility functions, any particular agent has a single determinate such function. This response allows us to drop the implausible assumption that there is just one legitimate way of valuing credences, whilst avoiding Bronfman's objection.

Why does Bronfman's objection not affect Greaves and Wallace's argument for Conditionalization? The reason is that, in the argument for Conditionalization, there is a single option—namely, the updating rule **Cond**—that has higher expected epistemic utility than all others relative to any proper epistemic utility function.

5 Recent and future work

This completes our survey of the two central applications of the argument strategy that Joyce introduced in his 1998 paper. In this final section, I wish to describe three ways in which this style of argument might develop—indeed, in some cases, it is already developing in these ways.

5.1 Characterising epistemic utility functions

As I mentioned in Section 4, we now have very powerful mathematical theorems that generalise Theorems 1 and 2 by showing that they hold for any of a number of different ways of measuring epistemic utility. What we lack are really strong philosophical arguments that every legitimate way of measuring epistemic utility belongs to this number. In short, we need a philosophical argument that every legitimate epistemic utility function is both separable and proper. Joyce has proposed an argument for propriety [Joyce, 2009, 279]; but Alan Hájek has raised a powerful objection [Hájek, 2008, 814-5].

Throughout, we have agreed with Joyce that the vindicated credence function at a world is v_w . But consider an agent with credences in a set of propositions that are governed by a non-classical logic. Then v_w may not be well defined. For instance, if there is a world at which one of the propositions is neither true nor false, then v_w is not a total function; if there is a world at which one of the propositions is both true and false, then v_w is ill-defined. Robbie Williams has investigated how we might define the

vindicated credence function at such a world [Williams, 2012]. And he has drawn on mathematical results due to Jeff Paris to characterise the credence functions that are permitted by Dominance relative to an epistemic utility function that measures proximity to these vindicated credence functions [Paris, 2001].

5.2 Choosing decision principles

Our argument for Probabilism depends on Dominance; our argument for Conditionalization depends on Maximize Expected Utility. But there are other decision principles out there. What epistemic norms do these alternative principles entail when coupled with a particular measure of epistemic utility?

In [Pettigrew, 2013], Richard Pettigrew considers a decision principle that (roughly) rules out as irrational any option that is expected by every possible objective chance function to have lower utility than some other option. Combined with any separable proper scoring rule, this principle entails David Lewis' Principal Principle.

Other decision principles that have had some attention are the principles of evidential and causal decision theory. These principles apply when the ways the world might be are not taken to be independent of the options from which the agent chooses. Michael Caie and Hilary Greaves have explored the consequences of these norms [Caie, ta], [Greaves, ms]. As Caie shows, considerations of epistemic utility should lead us to reject Probabilism for credences in certain propositions that concern a particular agent's credence in that proposition. These are situations in which Dominance does not apply, because the assumption of act-state independence that we have made throughout does not hold.

5.3 Modelling epistemic states

Finally, while we have focussed on the epistemic utility of credences, the same argument strategy may be applied to other sorts of epistemic state, such as full beliefs, indeterminate (or mushy) credences, and comparative probabilities. In each case, we need only identify the legitimate measures of epistemic utility. And, in each case, that may amount to identifying the vindicated epistemic states of that sort as well as measures of distance between such states. In the case in which both credences and full beliefs are present, important work was done by Hempel, Levi, and Maher [Hempel, 1962], [Levi, 1974], [Maher, 1993]. More recently, Kenny Easwaran has considered the case in which there are only full beliefs [Easwaran, ms]. In the case of comparative probabilities, Branden Fitelson and David McCarthy have some initial results [Fitelson and McCarthy, ms].

6 Conclusion

In his 1998 paper, Joyce introduced an entirely novel style of argument for epistemic norms: characterise the legitimate measures of the epistemic utility of an epistemic state at a world; then apply decision principles to this function to derive epistemic norms that govern these states. So far, we have seen justifications for some central norms given in this way. But there are many more to go.

References

- [Bronfman, ms] Bronfman, A. (ms.). A Gap in Joyce's Argument for Probabilism.
- [Caie, ta] Caie, M. (ta). Rational Probabilistic Incoherence. *Philosophical Review*.
- [de Finetti, 1931] de Finetti, B. (1931). Sul significato soggettivo della probabilit. *Fundamenta Mathematicae*, 17:298–329.
- [de Finetti, 1974] de Finetti, B. (1974). *Theory of Probability*, volume 1. Wiley, New York.
- [Easwaran, ms] Easwaran, K. (ms). Dr. Truthlove, or How I Learned to Stop Worrying and Love Bayesian Probabilities.
- [Fitelson and McCarthy, ms] Fitelson, B. and McCarthy, D. (ms). Steps towards a new foundation for subjective probability.
- [Greaves, ms] Greaves, H. (ms). Epistemic Decision Theory.
- [Greaves and Wallace, 2006] Greaves, H. and Wallace, D. (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459):607–632.
- [Hjek, 2008] Hjek, A. (2008). Arguments For—Or Against—Probabilism? *The British Journal for the Philosophy of Science*, 59(4):793–819.
- [Hjek, 2008] Hjek, A. (2008). Dutch Book Arguments. In Anand, P., Pattanaik, P., and Puppe, C., editors, *The Oxford Handbook of Corporate Social Responsibility*. Oxford University Press, Oxford.
- [Hempel, 1962] Hempel, C. G. (1962). Deductive-Nomological vs Statistical Explanation. In Feigl, H. and Maxwell, G., editors, *Minnesota Studies in the Philosophy of Science*, volume III, pages 98–169. University of Minnesota Press, Minneapolis.
- [Joyce, 1998] Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4):575–603.
- [Joyce, 2009] Joyce, J. M. (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*. Springer.
- [Kahneman and Tversky, 1983] Kahneman, D. and Tversky, A. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90:293–315.

- [Leitgeb and Pettigrew, 2010] Leitgeb, H. and Pettigrew, R. (2010). An Objective Justification of Bayesianism I: Measuring Inaccuracy. *Philosophy of Science*, 77:201–235.
- [Levi, 1974] Levi, I. (1974). *Gambling with Truth*. MIT Press, Cambridge, Mass.
- [Lewis, 1999] Lewis, D. (1999). Why Conditionalize? In *Papers in Metaphysics and Epistemology*. Cambridge University Press, Cambridge, UK.
- [Maher, 1993] Maher, P. (1993). *Betting on Theories*. Cambridge Studies in Probability, Induction, and Decision Theory. Cambridge University Press, Cambridge.
- [Paris, 2001] Paris, J. B. (2001). A Note on the Dutch Book method. In de Cooman, G., Fine, T., and Seidenfeld, T., editors, *ISIPTA '01, Proceedings of the Second International Symposium on Imprecise Probabilities and their Applications*, pages 301–306, Ithaca, NY. Shaker.
- [Pettigrew, 2013] Pettigrew, R. (2013). A New Epistemic Utility Argument for the Principal Principle. *Episteme*, 10(1):19–35.
- [Predd et al., 2009] Predd, J., Seiringer, R., Lieb, E. H., Osherson, D., Poor, V., and Kulkarni, S. (2009). Probabilistic Coherence and Proper Scoring Rules. *IEEE Transactions on Information Theory*, 55(10):4786–4792.
- [Ramsey, 1931] Ramsey, F. P. (1931). Truth and Probability. *The Foundations of Mathematics and Other Logical Essays*, pages 156–198.
- [Selten, 1998] Selten, R. (1998). Axiomatic Characterization of the Quadratic Scoring Rule. *Experimental Economics*, 1(1):43–61.
- [Vineberg, 2011] Vineberg, S. (2011). Dutch Book Arguments. In Zalta, E. N., editor, *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/dutch-book/>.
- [Williams, 2012] Williams, J. R. G. (2012). Gradational Accuracy and Non-Classical Semantics. *Review of Symbolic Logic*, 5(4):513–537.