



Pettigrew, R. (2017). Epistemic utility and the normativity of logic. *Logos and Episteme*, 8(4), 455-492. <http://logos-and-episteme.acadiasi.ro/volume-viii-issue-4-december-2017/>

Publisher's PDF, also known as Version of record

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

This is the final published version of the article (version of record). It first appeared online via Institutul European at <http://logos-and-episteme.acadiasi.ro/volume-viii-issue-4-december-2017/> . Please refer to any applicable terms of use of the publisher.

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 THE NORMATIVITY OF LOGIC

Richard PETTIGREW

ABSTRACT: How does logic relate to rational belief? Is logic normative for belief, as some say? What, if anything, do facts about logical consequence tell us about norms of doxastic rationality? In this paper, we consider a range of putative *logic-rationality bridge principles*. These purport to relate facts about logical consequence to norms that govern the rationality of our beliefs and credences. To investigate these principles, we deploy a novel approach, namely, epistemic utility theory. That is, we assume that doxastic attitudes have different epistemic value depending on how accurately they represent the world. We then use the principles of decision theory to determine which of the putative logic-rationality bridge principles we can derive from considerations of epistemic utility.

KEYWORDS: normativity of logic, epistemic value, accuracy-first epistemology

How does logic relate to rational belief? Is logic normative for belief, as some say? What, if anything, do facts about logical consequence tell us about norms of doxastic rationality? Here are some putative norms that seek to connect logic and rational belief:

(BP1) If Priest's Logic of Paradox governs propositions A and B , and B is strictly stronger than A in that logic, then, if you believe A , then you ought to believe B .

(BP2) If classical logic governs A_1, \dots, A_n, B , and A_1, \dots, A_n together entail B in that logic, then you ought not to believe each of A_1, \dots, A_n while disbelieving B .

(BP3) If you know that strong Kleene logic governs A and B , and you know that A entails B in that logic, then you have reason to see to it that your credence in A is at most your credence in B .

These illustrate something of the variety of claims that we might make in this area. Following John MacFarlane, we call such claims *bridge principles*—in particular, they are *logic-rationality bridge principles*.¹ Below, I will extend MacFarlane's taxonomy of such bridge principles to bring some order to this variety. Having done that, I wish to explore a novel way of adjudicating between them. In the existing literature, the following sorts of reasons are used to justify rejecting a given proposal of this sort:

¹ John MacFarlane, "In What Sense (If Any) Is Logic Normative for Thought?" (unpublished manuscript).

Conflicts with intuition. For instance, we might reject (BP2) by appealing to our intuitive reaction to cases like Makinson's Preface Paradox.² Suppose A_1, \dots, A_n enumerate all of my beliefs about British birdlife. So, for each A_i , I believe it. But I also realise that I am fallible on this topic. And thus, I disbelieve B , the proposition that all of my beliefs are true—that is, I disbelieve $B = A_1 \& \dots \& A_n$. Nonetheless, A_1, \dots, A_n together entail B . So I violate (BP2). Yet intuitively, we judge that I am perfectly rational. For this reason, some argue, we should reject (BP2).

Conflicts with ought-can. It is often noted that principles like (BP1) are extremely demanding, partly because we are not in a position to discover *all* the logical consequences of our beliefs, but also because, even if we could, we would be unable to store beliefs in all of them.³ Suppose, for instance, that A is the conjunction of the second-order Dedekind-Peano axioms for arithmetic. Then presumably we cannot discover all of the consequences of A ; and even if we could, we could not store them.⁴ Thus, we might take (BP1) to fail on the grounds that it conflicts with an ought-can principle.

Also, in recent unpublished work, Claire Field and Bruno Jacinto have tried to justify bridge principles in the following way:⁵

Justification on the basis of norms. They consider various norms that govern our beliefs. They consider the Truth Norm of Belief and the Knowledge Norm of Belief. And they ask which bridge principles follow from those norms. When considering the consequence of the Knowledge Norm for Belief, they consider the effects of assuming different frame conditions on the accessibility relation in the epistemic logic.

I wish to explore an alternative approach: sometimes this approach supplies a reason for rejecting a putative logic-rationality bridge principle, and sometimes it supplies a justification for accepting such a principle.

Justification by appeal to epistemic utility. In recent years, a number of philosophers have appealed to considerations of epistemic utility in order to justify various epistemic norms. Kenny Easwaran, Branden Fitelson, and Kevin Dorst have sought to establish the Lockean thesis concerning the normative link between credences and full beliefs, while Ted Shear, Branden Fitelson, and

² David Makinson, "The Paradox of the Preface," *Analysis* 25, 6 (1965): 205-207.

³ Gilbert Harman, *Change in View* (Cambridge, MA: MIT Press, 1986)

⁴ And, even if we could store them, surely that would not be a good use of our storage facilities. Note, however, that this last point does not turn on a conflict with an ought-can principle, but rather a conflict with a plausible principle governing how we should sensibly use our limited storage capacities.

⁵ Field and Jacinto presented this work at a conference, *The Normativity of Logic*, held at the University of Bergen, 14-16 June 2017.

Jonathan Weisberg have offered justifications of some of the principles of belief revision.⁶ On the credal side of epistemology, Jim Joyce and I have offered very closely related epistemic utility arguments for Probabilism,⁷ Together, Hilary Greaves and David Wallace have argued for Conditionalization on this basis, and R. A. Briggs and I have recently offered an alternative justification of that updating rule;⁸ Jason Konek and I have both sought to justify the Principal Principle;⁹ I have provided a rationalisation of the Principle of Indifference;¹⁰ Sarah Moss and Ben Levinstein have both sought norms that govern peer disagreement situations;¹¹ and Miriam Schoenfield has appealed to accuracy considerations to motivate a particular solution to the problem of higher-order evidence.¹²

We will spell out the idea behind these arguments in detail below, but roughly it is this. Our actions have different pragmatic value given different ways the world might be. We call this their *utility*. For instance, my action of betting that Labour will win the next UK General Election has high utility in worlds where they win and low utility in worlds where they lose. Similarly, our doxastic states—either our full beliefs, disbeliefs and suspensions of judgements, or our

⁶ Kevin Dorst, “Lockeans Maximize Expected Accuracy,” *Mind* (forthcoming), Kenny Easwaran “Dr Truthlove, Or: How I Learned to Stop Worrying and Love Bayesian Probabilities,” *Noûs* 50, 4 (2016): 816–853, Kenny Easwaran and Branden Fitelson “Accuracy, Coherence, and Evidence,” in *Oxford Studies in Epistemology*, volume 5, eds. Tamar Szabó Gendler and John Hawthorne (Oxford: Oxford University Press, 2015), Ted Shear, Branden Fitelson, and Jonathan Weisberg, “Two Approaches to Belief Revision” (unpublished manuscript).

⁷ James M. Joyce, “A Nonpragmatic Vindication of Probabilism,” *Philosophy of Science* 65, 4 (1998): 575–603, James M. Joyce. “Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief” in Franz Huber, & Christoph Schmidt-Petri (eds.) *Degrees of Belief*. (Springer, 2009), Richard Pettigrew, *Accuracy and the Laws of Credence* (Oxford: Oxford University Press, 2016).

⁸ Hilary Greaves and David Wallace, “Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility,” *Mind* 115, 459 (2006): 607–632, R. A. Briggs and Richard Pettigrew, “An Accuracy-Dominance Argument for Conditionalization” (unpublished manuscript).

⁹ Jason Konek, “The Simplest Possible Accuracy Argument for the Principal Principle” (unpublished manuscript), Richard Pettigrew, “A New Epistemic Utility Argument for the Principal Principle,” *Episteme* 10, 1 (2013): 19–35.

¹⁰ Richard Pettigrew, “Accuracy, Risk, and the Principle of Indifference,” *Philosophy and Phenomenological Research* 92, 1 (2016): 35–59.

¹¹ Sarah Moss, “Scoring Rules and Epistemic Compromise,” *Mind* 120, 480 (2011): 1053–1069. Benjamin A. Levinstein, “With All Due Respect: The Macro-Epistemology of Disagreement,” *Philosophers’ Imprint* 15, 3 (2015): 1–20.

¹² Miriam Schoenfield, “An Accuracy-Based Approach to Higher-Order Evidence,” *Philosophy and Phenomenological Research* (forthcoming).

credences—have different epistemic value given different ways the world might be. We will call this their *epistemic utility*. For instance, we might say that a true belief is more valuable than a false one, and a high credence in a true proposition is more valuable than that same high credence in a false proposition. Just as we choose between our actions using the principles of decision theory, so we might pick between different doxastic states using those same principles. After all, these decision-theoretic principles are simply claims about how facts about rationality are determined by facts about value—they govern how epistemic utility determines epistemic rationality just as much as they govern how pragmatic utility determines pragmatic rationality. Thus, just as previous authors have tested norms like Probabilism, Conditionalization, etc. by asking whether they follow from the principles of decision theory together with a particular account of epistemic utility, so we might test principles like (BP1), (BP2), (BP3), and their ilk, which claim to connect logic and rationality, in the same way.

It's worth emphasising here that proceeding in this way seems natural—more natural, perhaps, than appealing to intuition or to an ought-can principle. Presumably a large part of the reason why we think that logic might be normative for belief is that we think that beliefs aim at the truth; that is, we think that beliefs are better when true and worse when false. And presumably we also recognise that logic is the study of the relationships between the truth values of different propositions. If that's right, you should expect logic to tell you something about how best to obtain the aim of belief. Epistemic utility theory allows us to explore exactly how this might work. That's not to say that this is the only framework in which to explore this: seeking out the consequences of the Truth Norm for Belief, as Field and Jacinto do, is an alternative approach. But I hope to convince you that it is a fruitful way to do so.

A Taxonomy of Bridge Principles

Each principle that purports to connect logic and rationality shares the same form. It is a conditional. Its antecedent is a proposition $A(T \& L)$. T is a claim about the logic that governs some set of propositions; L is a claim about the consequence relation of that logic; A is a propositional operator that acts on the conjunction, $T \& L$. The consequent of a bridge principle is a normative claim C concerning an agent's beliefs or credences. Thus, our logic-rationality bridge principles have the form $A(T \& L) \rightarrow C$.

In (BP1), T is the claim that the Logic of Paradox governs A and B , and L is the claim that A entails B in that logic, but B does not entail A . In (BP3), T is the claim that strong Kleene logic governs A and B , while L is the claim that A entails

B . In (BP1) and (BP2), $A(T \& L)$ is just $T \& L$, so A is the identity operator in this case, whereas in (BP3), $A(T \& L)$ is the proposition that you know $T \& L$, so that A is the knowledge operator in this case. In (BP1), C is the conditional: if you believe A , then you ought to believe B . That is, C is a *narrow scope norm*. In (BP2), C is a *wide scope norm*: it ought not to be that you believe each of A_1, \dots, A_n and you disbelieve B . In (BP3), the normative claim in the consequent is not stated in terms of ought at all; it is stated in terms of reasons, so it is weaker.

We now expand a little on John MacFarlane’s taxonomy for bridge principles. MacFarlane lists a number of dimensions along which bridge principles can differ, and he lists the ways in which they might differ along these different dimensions. I simply add a couple of further dimensions to his list.

Grain What sort of doxastic states does the norm govern?

Credences The norm governs credences or degrees of belief.

Full beliefs The norm governs full beliefs, full disbeliefs, and suspensions of judgment.

Normativity What sort of norm is $A(T \& L) \rightarrow C$?¹³

Evaluation It is used only to evaluate an agent’s doxastic state.

Appraisal It is used to apportion epistemic blame and fault to the agent.

Directive It is used to direct the agent’s doxastic life.

Governing logic Which logic governs the propositions in question, according to T ?

Classical We denote the consequence relation of this logic \models_{cl}

Strong Kleene logic We denote the consequence relation of this logic \models_{skl}

Logic of Paradox We denote the consequence relation of this logic \models_{lp}

and so on...

Strength of logical claim What is the strength of the claim L about logical consequence that occurs in the antecedent? Weak or strong?

Weak $A \models B$.

Strong $A \models B$ and $B \not\models A$.

Antecedent operator What is the operator A in the antecedent $A(T \& L)$? That is, under what conditions on $T \& L$ does the bridge principle get triggered?

Identity $A(T \& L) = T \& L$.

¹³ Cf. Florian Steinberger, “Three ways in which logic might be normative” (unpublished manuscript).

Richard Pettigrew

Obvious $A(T \& L) = T \& L$ is obvious.

Knowledge $A(T \& L) =$ You know $T \& L$.

Belief $A(T \& L) =$ You believe $T \& L$.

Number of premises L is a fact about logical consequence. That is, it is a proposition of the form $A_1, \dots, A_n \models B$ for some propositions A_1, \dots, A_n, B . What is n ?

Consequent operator What is the operator in the consequent C ? Is it an ought operator, a reasons operator, or a permission operator?

Ought C states a norm in terms of *ought*.

Reasons C states a norm in terms of *reasons*.

Permission C states a norm in terms of *permission*.

Consequent scope What is the scope of the operator found in the consequent C ? Does it apply to the consequent of the conditional only, both antecedent and consequent separately, or the whole conditional together?

Consequent C takes the form $X \rightarrow N(Y)$, where N is the normative operator identified in the previous condition. Thus, C is a narrow scope norm.

Whole C takes the form $N(X \rightarrow Y)$. Thus, C is a wide scope norm.

Both C takes the form $N(X) \rightarrow N(Y)$.

Polarity What is the strength of the claim in the consequent of the conditional in C ?

Positive The consequent of the conditional in C is a positive demand that the agent has a particular attitude.

Negative The consequent of the conditional in C is a negative demand that the agent does not have a particular attitude.

Picking a different answer for each of these gives a different putative normative claim about the connection between logic and rationality. Thus, for instance, (BP1) arises from the following choices: it governs full beliefs; the logic is Logic of Paradox; the logical claim is weak; the operator in the antecedent is the identity operator; the claim about logical consequent involves just a single premise; the operator in the consequent is the ought operator and that operator takes narrow scope in the consequent; and the polarity of the consequent is positive.

In what follows, we'll use epistemic utility to adjudicate between these different bridge principles. We'll divide our treatment into two parts: first, we'll treat full beliefs; second, we'll treat credences.

Bridge Principles for Full Beliefs

We begin by considering epistemic utility for full beliefs. I will present the now-standard veritist story for the classical case. This originates with Carl Hempel, but in its current form it is due to work by Kevin Dorst, Kenny Easwaran, and Branden Fitelson.¹⁴ After that, I will extend it to the non-classical case.

Suppose you entertain a particular proposition; it is there before your mind. Then there are three categorical doxastic attitudes that you might adopt towards it: you can believe it (**B**), disbelieve it (**D**), or suspend judgment on it (**S**). Suppose \mathcal{F} is the set of propositions that you entertain. We can represent your doxastic state by a function $b : \mathcal{F} \rightarrow \{\mathbf{B}, \mathbf{S}, \mathbf{D}\}$. We call this your *belief function*. Our first order of business is to describe an epistemic utility function for doxastic states represented in this way. An epistemic utility function takes a doxastic state and a possible world and returns a measure of how much epistemic utility that state has at that possible world. Here and throughout, we will assume a *veritist* account. That is, we will assume that the sole fundamental source of epistemic value for doxastic states is their accuracy; a doxastic state has greater epistemic value the more accurately it represents the world. One consequence of this is that the epistemic utility of your doxastic state at a possible world depends only on the truth values at that world of the propositions that you entertain. So, just as we can represent a doxastic state as a function from \mathcal{F} to the set of possible doxastic attitudes, so we can represent a possible world as a consistent valuation function from \mathcal{F} to the set of possible truth values. Since we are currently presenting the classical case, the set of truth values is $\{\mathbf{t}, \mathbf{f}\}$, and the consistency in question is classical consistency.

Now, we wish to define a function EU such that, if $b : \mathcal{F} \rightarrow \{\mathbf{B}, \mathbf{S}, \mathbf{D}\}$ is a belief function on \mathcal{F} and $w : \mathcal{F} \rightarrow \{\mathbf{t}, \mathbf{f}\}$ is a classical valuation function on \mathcal{F} , then $\text{EU}(b, w)$ is the epistemic utility of the doxastic state represented by b at the possible world represented by w . First, we assume that EU is *additive*: that is, $\text{EU}(b, w)$ is the sum of the epistemic utilities at w of the different doxastic attitudes that b comprises. That is, there is a *local epistemic utility function* $\text{eu} : \{\mathbf{t}, \mathbf{f}\} \times \{\mathbf{B}, \mathbf{S}, \mathbf{D}\} \rightarrow [-\infty, \infty]$ such that

$$\text{EU}(b, w) = \sum_{X \in \mathcal{F}} \text{eu}(w(X), b(X))$$

¹⁴ Carl Hempel, "Deductive-Nomological vs. Statistical Explanation," in *Minnesota Studies in the Philosophy of Science*, vol. III, eds. Herbert Feigl and Grover Maxwell (Minneapolis: University of Minnesota Press, 1962), 98–169, Dorst, "Lockeans Maximize Expected Accuracy," Easwaran "Dr Truthlove," Easwaran and Fitelson, "Accuracy, Coherence, and Evidence."

Thus, $eu(\mathbf{t}, \mathbf{B})$ is the epistemic utility of believing a proposition when it is true, while $eu(\mathbf{t}, \mathbf{D})$ is the epistemic utility of disbelieving a proposition when it is true, and so on. And $eu(\mathbf{t}, \mathbf{B}) + eu(\mathbf{f}, \mathbf{B})$ is the epistemic utility of an agent with one belief in a truth and one belief in a falsehood and no other doxastic attitudes. Next, we identify our proposed local epistemic utility function. It is this:

$$eu(\mathbf{t}, \mathbf{B}) = eu(\mathbf{f}, \mathbf{D}) = R \text{ (for getting it Right)}$$

$$eu(\mathbf{t}, \mathbf{S}) = eu(\mathbf{f}, \mathbf{S}) = 0$$

$$eu(\mathbf{t}, \mathbf{D}) = eu(\mathbf{f}, \mathbf{B}) = -W \text{ (for getting it Wrong)}$$

where $R, W > 0$. Thus, true beliefs and false disbeliefs are equally valuable, with epistemic utility R ; and they are more valuable than suspensions, which are equally valuable whatever the outcome, with epistemic utility 0; and they, in turn, are more valuable than false beliefs and true disbeliefs, which are equally valuable, with epistemic utility $-W$. According to William James, two principles guide our epistemic life: *Believe truth! Shun error!*¹⁵ If you agree, you might take R and W to measure the strength of those two exhortations, respectively. The higher R , the more you care about getting things right; the higher W , the more you care about not getting things wrong. Thus, if $R > W$, you might call yourself an epistemic radical; if $R = W$, you are an epistemic centrist; and if $W > R$, you are an epistemic conservative.

Now, let's see what these different positions have to say about the logic-rationality bridge principles that we categorized at the beginning of the paper. Throughout, we will have cause to refer to five different belief functions defined on A and B . We define them here for ease of reference:

	A	B
b_1	B	D
b_2	B	S
b^*	S	S
b_1^\dagger	D	B
b_2^\dagger	S	B

Thus, for instance, we might take A to be *Labour will win* and B to be *Labour or the Greens will win*. Thus, b_1 believes that Labour will win, but disbelieves that Labour or the Greens will win, while b_1^\dagger switches those attitudes, disbelieving that Labour will win, but believing that Labour or the Greens will

¹⁵ William James, "The Will to Believe," in *The Will to Believe, and Other Essays in Popular Philosophy* (New York: Longmans Green, 1905).

win. Similarly, b_2 believes Labour will win, but suspends on Labour or the Greens winning, while b_2^\dagger switches those attitudes. And b^* suspends on both propositions.

Let's start with the epistemic conservative; for them, recall, $W > R$. Then we have an epistemic utility argument for the following logic-rationality bridge principle:¹⁶

(BP4) If \models_{cl} governs A and B , and $A \models_{cl} B$, then you ought not to believe A while disbelieving B .

That is, when A and B are classical propositions, and A classically entails B , you ought not to have the belief function b_1 . This is the single-premise version of the bridge principle that MacFarlane calls (Wo-).

Here's the argument for (BP4). Suppose A classically entails B . Then consider b_1 and b^* . While b_1 believes A and disbelieves B , b^* suspends judgment on both. Now consider the different ways the world might be and the epistemic utility of the two belief functions at those different worlds:

	A	B	$EU(b_1, w)$	$EU(b^*, w)$
w_1	t	t	$R - W$	$0 + 0$
w_2	t	f	$R + R$	$0 + 0$
w_3	f	t	$-W - W$	$0 + 0$
w_4	f	f	$-W + R$	$0 + 0$

Since $W > R$, it follows that $R - W < 0$. Thus, at all worlds except the one at which b_1 gets everything right—the world at which A is true and B is false—the epistemic utility of b_1 is negative; and the epistemic utility of b^* is always 0. However, given that A entails B , there is no world at which A is true and B is false—that is, w_2 is not a classically consistent valuation and thus does not represent a genuine possibility. So, at all logically possible worlds—that is, at w_1 , w_3 , and w_4 — b^* has greater epistemic utility than b_1 . That is, as a matter of logical necessity, it is epistemically better to have belief function b^* than b_1 . That is, $EU(b_1, w) < EU(b^*, w)$ for all logically possible worlds w . In such cases, we say that b^* *strictly logically dominates* b_1 *relative to* EU.

Now, in decision theory, strict logical dominance is often taken to be a sign of irrationality. That is, the following is taken to be a principle of rationality:

Strict Logical Dominance If option o^* has greater utility than option o at every logically possible world, then o is irrational.

¹⁶See Easwaran's "Dr Truthlove," for very closely related results in which the only categorical doxastic attitudes are belief and suspension.

Richard Pettigrew

Thus, we have our first epistemic utility argument for a logic-rationality bridge principle:

(EU1) Epistemic Conservatism + Strict Logical Dominance \Rightarrow (BP4).

Before we move on, it helps to see this argument in a particular case. Consider, then, the person who believes that Labour will win, and disbelieves that Labour or the Greens will win. Such a person would do better for sure if they were to suspend judgment on both propositions. If Labour do win, then their belief is true but their disbelief false, and that means that they have negative epistemic utility; if Labour don't win but the Greens do, then they do maximally badly, since both attitudes are wrong, so they have negative epistemic utility; and if neither Labour nor the Greens win, then they are in the same situation as when Labour wins, namely, that one attitude is right and the other wrong, and that means that they have negative epistemic utility. Thus, they are guaranteed to have negative epistemic utility. On the other hand, if they were to suspend on both propositions, they would be guaranteed to have a neutral epistemic utility of 0. Thus, suspending dominates.

Hopefully, this gives a taste of the sort of epistemic utility argument we will pursue in this paper. Each argument consists of the components: (i) an account of epistemic utility—in this case, we assumed Epistemic Conservatism; (ii) a decision-theoretic principle—in this case, Strict Logical Dominance; (iii) a mathematical fact that shows that, if you apply the decision-theoretic principle using the account of epistemic utility, you obtain the epistemic norm that you seek, such as (BP4)—in this case, we demonstrated the mathematical result using the truth table above.

There are a number of ways in which we might try to adapt this argument. We might ask what happens when we switch Epistemic Conservatism for Epistemic Centrism or Epistemic Radicalism; or when we include more than one premise in the fact about logical consequence in the antecedent; or when we replace classical logic with some non-classical alternative; or when we consider the possibility of rational ignorance of logical truths.

Epistemic Conservatism, Centrism, and Radicalism

First, let's see what happens when we move from Epistemic Conservatism to Epistemic Centrism or Epistemic Radicalism. Recall: according to Epistemic Centrism, $R = W$ —getting things right is exactly as good as getting things wrong is bad. Now, we can see from the table above that, for the Epistemic Centrist, b^* does not strictly logically dominate b_1 . After all, at worlds at which A and B are both

true or both false, the epistemic utility of b_1 (namely, $R - W$) is the same as the epistemic utility of b^* (namely, 0); it does not exceed it. And indeed it is straightforward to see that no alternative belief function strictly dominates b_1 .¹⁷ Nonetheless, note that b^* is at least as good, epistemically speaking, as b_1 at all logically possible worlds, and strictly better at some. In such cases, we say that b^* *weakly logically dominates* b_1 relative to EU. In decision theory, weak logical dominance is often taken to be a sign of rationality in the same way that strict logical dominance is. That is, the following is taken to be a principle:

Weak Logical Dominance If option o^* has at least as great utility as option o at every logically possible world, and greater utility at some, then o is irrational.

Now, note that, in order to apply this to the choice of belief functions on A and B , we must ensure that it is genuinely logically possible that A is false and B true. That is, Weak Logical Dominance will tell us nothing when B entails A . In that situation, b_1 and b^* are equally good in every logically possible world, and there's nothing irrational about picking an option with that feature. This gives us an epistemic utility argument for a slightly weaker version of (BP4):

(BP5) If \models_{cl} governs A and B , and $A \models_{cl} B$, and $B \not\models_{cl} A$, then you ought not to believe A while disbelieving B .

Here's the argument:

(EU2) Epistemic Centrism + Weak Logical Dominance \Rightarrow (BP5).

Weak Logical Dominance also proves crucial when we move to Epistemic Radicalism—that is, the claim that $R > W$. It is clear from the table above that b^* neither strictly nor weakly logically dominates b_1 when $R > W$. After all, in this situation, b_1 outperforms b^* when A and B have the same truth value, since $R - W > 0$. As above, it is straightforward to see that there is no alternative belief function that strictly dominates b_1 for the Epistemic Radicalist. But there is an alternative that weakly dominates it, namely, b_1^\dagger from above. Recall: b_1^\dagger disbelieves A and believes B .

	A	B	$EU(b_1, w)$	$EU(b_1^\dagger, w)$
w_1	t	t	$R - W$	$-W + R$
w_2	t	f	$R + R$	$-W - W$
w_3	f	t	$-W - W$	$R + R$
w_4	f	f	$-W + R$	$R - W$

¹⁷ By checking cases, we can see that, if a belief function b strictly outperforms b_1 when A and B are both true, then b_1 strictly outperforms b when A and B are both false.

Thus, b_1 and b_1^\dagger have the same epistemic value when A and B are both true or both false, and b_1^\dagger is strictly better than b_1 when A is false and B is true. Thus, we have another epistemic utility argument for (BP5):

(EU3) Epistemic Radicalism + Weak Logical Dominance \Rightarrow (BP5).

Thus, in sum: for every point on the scale between Epistemic Conservatism and Epistemic Radicalism, there is an epistemic utility argument for (BP5). Indeed, for every point on that scale, b_1^\dagger weakly logically dominates b_1 . Thus, we have:

(EU4) Weak Logical Dominance \Rightarrow (BP5).

And, for Epistemic Conservatism, there is an epistemic utility argument for (BP4), namely, (EU1).

Moreover, note that a similar trick can be used to establish the following bridge principle:

(BP6) If \models_{cl} governs A and B , and $A \models_{cl} B$, and $B \not\models_{cl} A$, then you ought not to believe A while suspending judgment in B .

That is, if A is strictly stronger than B , then you ought not to have the belief function b_2 defined above. We can justify this by noting that b_2 is weakly logically dominated by b_2^\dagger , which we defined above. Recall: b_2^\dagger suspends on A and believes B .

	A	B	$EU(b_2, w)$	$EU(b_2^\dagger, w)$
w_1	t	t	$R - 0$	$0 + R$
w_2	t	f	$R + 0$	$0 - W$
w_3	f	t	$-W - 0$	$0 + R$
w_4	f	f	$-W + 0$	$0 - W$

Thus,

(EU5) Weak Logical Dominance \Rightarrow (BP6)

Thus, if \models_{cl} governs A and B , and $A \models_{cl} B$, and $B \not\models_{cl} A$, then we have an argument against believing A and disbelieving B , and an argument against believing A and suspending on B . Since belief, disbelief, and suspension are the only available categorical doxastic attitudes to a proposition, it seems at first sight that these two arguments then furnish us with a further argument that, if you believe A , and you adopt any categorical doxastic attitude towards B , then you

ought to believe B . But that's not quite right.¹⁸ The problem is that, while arguments (EU4) and (EU5) establish flaws in believing A and either suspending on B or disbelieving B , they do not rule out the possibility that believing A and believing B is also flawed in the same way. Now it turns out that, if A is a contradiction then that is indeed the case. If A is a contradiction, believing A and believing B is strictly logically dominated by disbelieving A and believing B . However, if A is not a contradiction, then believing A and believing B is not even weakly logically dominated.¹⁹ Thus, we have an argument for:

(BP7) If \models_{cl} governs A and B , and $A \models_{cl} B$, and $B \not\models_{cl} A$, and A is not a classical contradiction, then you ought to see to it that, if you believe A , and you entertain B , then you believe B .

The argument:

(EU6) Weak Logical Dominance \Rightarrow (BP7)

(BP7) doesn't say that you ought to see to it that you believe B if you believe A . Rather, it says that you ought to see to it that, *if B is a proposition that you entertain and to which you assign an attitude at all*, then you believe B if you believe A . It does this by showing that it would be irrational to assign either of the alternative attitudes to B in a way that it wouldn't be irrational to believe B . This is as close as we can get to the principle that MacFarlane calls (Wo+).

The upshot of this section is that, when the logic is classical and the entailment is a strict single premise entailment, epistemic utility considerations vindicate the logic-rationality bridge principles that seem most natural. They justify the wide scope versions of the norms, and they justify the versions with negative polarity. So they support bridge principles that are not vulnerable to some of Harman's main criticisms. They are not excessively demanding, since they do not demand that an agent have any attitude at all towards B ; rather, they only say what she should do if she does have an attitude towards B . And they posit wide scope norms, so they are not vulnerable to Harman's objection that narrow scope norms arbitrarily favour one way of resolving an inconsistency in your beliefs.

Multi-Premise Entailments

Next, let's see what happens when we include more premises. It turns out that the answer depends on the values of R and W . In this section, we'll have cause to refer to three different belief functions. Again, we define them now for ease of

¹⁸ Thanks to Anandi Hattiangadi on this point.

¹⁹ No alternative does as well when A and B are both true.

Richard Pettigrew

reference:

	A_1	...	A_n	B
b_3	B	...	B	D
b_4	B	...	B	S
b°	S	...	S	S

Suppose \models_{cl} governs A_1, \dots, A_n, B , and $A_1, \dots, A_n \models_{cl} B$. Now, we might say that an assignment of epistemic utility is extremely conservative if $nR < W$. Then, assuming Extreme Epistemic Conservatism, b_3 is strictly logically dominated by b° relative to EU. After all, b° has exactly the same epistemic utility at every world—namely, 0—while b_3 performs best when A_1, \dots, A_n are all true, but B is false, and in that situation b_3 has epistemic utility $nR - W$, which by hypothesis is less than 0. Thus, we have an epistemic utility argument for the following bridge principle:

(BP8) If \models_{cl} governs A_1, \dots, A_n, B , and $A_1, \dots, A_n \models_{cl} B$, then you ought not to believe each of A_1, \dots, A_n while disbelieving B .

This is the multi-premise analogue of (BP4) and MacFarlane’s (Wo-). Here’s the argument:

(EU7) Extreme Epistemic Conservatism + Strict Logical Dominance \Rightarrow (BP8)

(EU7) rules out b_3 as irrational. Interestingly, we cannot strengthen this argument to rule out b_4 as irrational as well. That is, we cannot give an argument for

(BP9) If \models_{cl} governs A_1, \dots, A_n, B , and $A_1, \dots, A_n \models_{cl} B$, then you ought not to believe each of A_1, \dots, A_n while suspending judgment in B .

Indeed, for any number of premises n , and any values R, W for the goodness of getting things right and the badness of getting things wrong, respectively, there are classical propositions A_1, \dots, A_n, B such that $A_1, \dots, A_n \models_{cl} B$, and such that the belief function b_4 is not dominated. Indeed, there is always a regular probability function p that expects the belief function b_4 to have the highest epistemic utility of all possible belief functions defined on A_1, \dots, A_n, B .²⁰ And this is sufficient to show that b_4 is not weakly logically dominated.²¹

²⁰A probability function is regular if it assigns strictly positive credence to every possibility.

²¹Let’s see why this is so. Suppose that one option o^* strictly logically dominates another o . Then o^* has strictly greater expected utility than o by the lights of *any* probability function. After all, the utility of o^* is greater than the utility of o at every world. So any weighted sum of the utilities of o^* will be greater than the corresponding weighted sum of the utilities of o . And of

Given how useful it is to know that a belief function maximises expected epistemic utility relative to a probability function, let's spell out exactly how this works. The *expected epistemic utility of a belief function b by the lights of probability function p* is defined as follows:

$$\text{Exp}_{\text{EU}}(b \mid p) = \sum_w p(w) \text{EU}(b, w)$$

It is straightforward to see that:

$$\begin{aligned} \text{Exp}_{\text{EU}}(b \mid p) &= \sum_w p(w) \text{EU}(b, w) \\ &= \sum_w p(w) \sum_{X \in \mathcal{F}} \text{eu}(w(X), b(X)) \\ &= \sum_{X \in \mathcal{F}} \sum_w p(w) \text{eu}(w(X), b(X)) \\ &= \sum_{X \in \mathcal{F}} p(X) \text{eu}(\mathbf{t}, b(X)) + p(\bar{X}) \text{eu}(\mathbf{f}, b(X)) \end{aligned}$$

That is, the expected utility of b is the sum of the expected utilities of the individual attitudes it assigns. Thus, b has maximal expected epistemic utility by the lights of p iff each attitude that b assigns has maximal expected epistemic utility by the lights of p .

Now, note the following fact:

Theorem 1 (Hempel-Easwaran-Dorst)

If $W \geq R$, then

- i. belief in X has maximal expected EU by the lights of p if $1 \geq p(X) \geq \frac{W}{W+R}$

course the expected utilities of o and o^* are just such weighted sums. Thus, if o maximises expected utility by the lights of some probability function, there is no option that strictly logically dominates o , for such an option would have strictly greater expected utility by the lights of that probability function. Next, suppose that o^* weakly logically dominates o . Then the utility of o^* is at least the utility of o at every world and strictly greater at some. So any weighted sum of the utilities of o^* that assigns strictly positive weight to each will be greater than the corresponding weighted sum of the utilities of o . And again the expected utilities of o and o^* by the lights of a regular probability function are just such weighted sums. Thus, if o maximises expected utility by the lights of some regular probability function, there is no option that weakly logically dominates o , for such an option would have strictly greater expected utility by the lights of that probability function.

- ii. suspension in X has maximal expected EU by the lights of p if $\frac{W}{W+R} \geq p(X) \geq \frac{R}{W+R}$
- iii. disbelief in X has maximal expected EU by the lights of p if $\frac{R}{W+R} \geq p(X) \geq 0$.

If $W < R$, then

- i. belief in X has maximal expected EU by the lights of p if $1 \geq p(X) \geq \frac{1}{2}$
- ii. suspension in X never maximises expected EU by the lights of p
- iii. disbelief in X has maximal expected EU by the lights of p if $\frac{1}{2} \geq p(X) \geq 0$.

It is in this sense that epistemic utility theory vindicates a normative reading of the Lockean thesis. Suppose $W \geq R$ — that is, you are an epistemic conservative or centrist. Then there is a threshold $t = \frac{W}{W+R}$ such that you are rationally required to believe a proposition if your credence in that proposition exceeds t , you are rationally required to suspend on that proposition if your credence lies strictly between $1 - t$ and t , and you are rationally required to disbelief it if you credence lies below $1 - t$. If your credence is exactly t , then believing and suspending both maximise expected EU; if your credence is exactly $1 - t$, then disbelieving and suspending both maximise expected EU. Next, suppose $W < R$ — that is, you are an epistemic radical. Then there is a threshold $t = \frac{1}{2} = 1 - t$ such that you are rationally required to believe a proposition if your credence in that proposition exceeds t , and you are rationally required to disbelief it if you credence lies below t . If your credence is exactly t , then believing and disbelieving both maximise expected EU.

Thus, given $W \geq R$, and n , in order to find propositions A_1, \dots, A_n, B such that $A_1, \dots, A_n \models_{\text{cl}} B$ and a probability function by the lights of which $b_4(A_1) = \dots = b_4(A_n) = \mathbf{B}$ and $b_4(B) = \mathbf{S}$ has maximal expected epistemic utility, we need only find p such that $p(A_1) = \dots = p(A_n) \geq \frac{W}{W+R}$ and $p(B) \leq \frac{W}{W+R}$. And that is straightforward to do, since conjunctions typically have lower probability than their conjuncts.²² If $W < R$, the situation is a little more complicated, since there is no probability function by the lights of which b_4 has maximal expected utility, since there is no probability function for which suspending judgment has maximal

²² Let X and Y be logically independent propositions. Let $A_1 = X$ and $A_2 = \dots = A_n = Y$, and $B = X \& Y$. So $A_1, \dots, A_n \models_{\text{cl}} B$. Then let $p(X), p(Y) = \frac{W}{W+R}$ and $p(\overline{X}Y \vee X\overline{Y}) > 0$. So $\frac{R}{W+R} < p(X \& Y) < \frac{W}{W+R}$. Thus, $p(A_1) = \dots = p(A_n) = t$, while $1 - t < p(B) < t$.

expected utility. But there are regular probability functions such that the only belief function that has higher expected epistemic utility than b_4 assigns belief to each A_i and assigns belief or disbelief to B ; and it is easy to see that neither of those strictly or weakly dominates b_4 ; so nothing does. So, unlike in the single-premise case, we cannot give a dominance argument for (BP9).

Similar reasoning shows that, if we do not assume Extreme Epistemic Conservatism, then there is no guarantee even that b_3 is weakly or strictly logically dominated. That is, we can find propositions A_1, \dots, A_n, B , and a probability function p such that b_3 has maximal expected epistemic utility by the lights of p . Here's an example. Suppose there is a fair lottery with $n + 1$ tickets. Let A_i be proposition *Ticket i does not win*, for $1 \leq i \leq n$, and let B be the proposition *Ticket $n + 1$ wins*. Then A_1, \dots, A_n entail B . However, if we suppose that each ticket has the same chance of winning, then $p(A_i) = \frac{n}{n+1} \geq \frac{W}{W+R}$, for each $1 \leq i \leq n$, while $p(B) = \frac{1}{n+1} \leq \frac{R}{W+R}$.²³ Thus, believing that each of the first n tickets does not win whilst disbelieving that the final ticket will win is not strictly or weakly logically dominated. Indeed, not only is it not dominated, it is in fact the belief assignment recommended by the objective chance function in this context.

The upshot of this section is that epistemic utility considerations vindicate intuitions such as the Preface Paradox, which entail that logical consistency is not a rational requirement on beliefs. If we care so much more about avoiding error than about believing truths, then we can recover the bridge principle that prohibits believing each of a set of propositions whilst disbelieving one of their classical logic consequences. But if we do not, we cannot. There will be situations, such as lottery or preface cases, in which such doxastic attitudes will be rationally required by the natural probability function that governs them.

Non-Classical Logics

Next, let's look at what happens when we move from classical logic to a non-classical alternative. We'll focus on two particular non-classical logics: Kleene's strong logic of indeterminacy (skl) and Priest's Logic of Paradox (lp). Both have three truth values: $\{\mathbf{t}, \mathbf{u}, \mathbf{f}\}$. And both specify the same truth-functional definitions for the connectives, namely,

²³Note: $\frac{n}{n+1} \geq \frac{W}{W+R}$ iff $nR \leq W$ iff $\frac{1}{n+1} \leq \frac{R}{W+R}$.

X	\bar{X}
t	f
u	u
f	t

$\&$	t	u	f
t	t	u	f
u	u	u	f
f	f	f	f

\vee	t	u	f
t	t	t	t
u	t	u	u
f	t	u	f

They differ in the interpretation of the third truth value **u**. In strong Kleene logic it is taken to mean *neither true nor false*, while in Logic of Paradox, it is taken to mean *both true and false*. And they differ in the role that those truth values play in the definition of logical consequence for the two logics. In both logics, A_1, \dots, A_n entails B iff whenever each of A_1, \dots, A_n takes one of the designated truth values, B does as well. But they differ in the specification of the designated truth values. For Kleene’s logic, **t** is the only designated truth value. For the Logic of Paradox, **t** and **u** are both designated. Thus, $A \vee \bar{A}$ is not a tautology in strong Kleene logic, since it has truth value **u** when A does, and **u** is not designated; but it is a tautology in Logic of Paradox, since it has value **t** or **u** regardless of the truth value of A , and both are designated.

Having introduced strong Kleene logic and Logic of Paradox, how might we define epistemic utility for belief functions when one of those logics governs the propositions that our agent entertains? It is easy to see what the possible worlds are in such a situation. They are the logically consistent valuation functions $w: \mathcal{F} \rightarrow \{\mathbf{t}, \mathbf{u}, \mathbf{f}\}$. And a local epistemic utility function is a function $eu_{\text{skl}}/eu_{\text{lp}}: \{\mathbf{B}, \mathbf{S}, \mathbf{D}\} \times \{\mathbf{t}, \mathbf{u}, \mathbf{f}\} \rightarrow [-\infty, \infty]$. We assume that the local epistemic utility functions in these situations extend the classical ones, which specify the epistemic utility for **B**, **S**, and **D** when the proposition towards which the attitude is directed takes truth value **t** or **f**. Thus, we need only specify the epistemic value of each of these attitudes when directed towards a proposition with truth value **u**.

Take strong Kleene logic first. Here, **u** is interpreted to mean *neither true nor false*. We must define $eu_{\text{skl}}(\mathbf{u}, \mathbf{B})$, $eu_{\text{skl}}(\mathbf{u}, \mathbf{S})$, and $eu_{\text{skl}}(\mathbf{u}, \mathbf{D})$. There are a number of views one might take on these values. These will depend in part on the use to which the logic is being put, but there will also be disagreements once we have fixed the use of the logic. I do not seek to adjudicate these disagreements here, but rather to spell out their consequences for the logic-rationality bridge principles.

For instance, Hartry Field, following Kripke, claims that strong Kleene logic is the logic that governs the liar sentence.²⁴ That is, the liar sentence takes truth value **u**; it is neither true nor false. What’s more, he takes the ideal attitude to

²⁴ Hartry Field, *Saving Truth from Paradox* (New York: Oxford University Press, 2008), Saul Kripke, “Outline of a Theory of Truth,” *Journal of Philosophy* 72, 19 (1975): 690–716.

propositions with truth value \mathbf{u} to be disbelief (or rejection). Indeed, Michael Caie notes that this is the consensus amongst those who take a paracomplete approach to semantic paradoxes.²⁵ Thus, for Field, the following is the natural assignment of epistemic value to the various categorical doxastic attitudes to a proposition with truth value \mathbf{u} :

$$\begin{aligned} \text{eu}(\mathbf{t}, \mathbf{B}) &= \text{eu}(\mathbf{f}, \mathbf{D}) = \text{eu}(\mathbf{u}, \mathbf{D}) = R \\ \text{eu}(\mathbf{t}, \mathbf{S}) &= \text{eu}(\mathbf{f}, \mathbf{S}) = \text{eu}(\mathbf{u}, \mathbf{S}) = 0 \\ \text{eu}(\mathbf{t}, \mathbf{D}) &= \text{eu}(\mathbf{f}, \mathbf{B}) = \text{eu}(\mathbf{u}, \mathbf{B}) = -W \end{aligned}$$

Given this, we have a strict dominance argument for the strong Kleene version of (BP4) and weak dominance arguments for the strong Kleene versions of (BP5) and (BP6), which are obtained from those principles by replacing \models_{cl} with \models_{skl} . In fact, this follows from a more general fact, which also covers the original, classical versions of (BP4-6):²⁶

Theorem 2

Suppose:

- i. The logical consequence relation for a many-valued logic k is defined in terms of the preservation of designated truth values.
- ii. That is, $A_1, \dots, A_n \models_k B$ iff, for all worlds w , if $w(A_i)$ is a designated truth value in k , for each $1 \leq i \leq n$, then $w(B)$ is a designated truth value in k .
- iii. $A \models_k B$
- iv. If \mathbf{i} is a designated truth value, then $\text{eu}(\mathbf{i}, \mathbf{B}) = R$, $\text{eu}(\mathbf{i}, \mathbf{S}) = 0$, $\text{eu}(\mathbf{i}, \mathbf{D}) = -W$.
- v. If \mathbf{i} is not a designated truth value, then $\text{eu}(\mathbf{i}, \mathbf{B}) = -W$, $\text{eu}(\mathbf{i}, \mathbf{S}) = 0$, $\text{eu}(\mathbf{i}, \mathbf{D}) = R$.

Then:

- a. If Epistemic Conservatism holds, then believing A and disbelieving B is strictly logically dominated by suspending on A and suspending on B .

²⁵ Michael Caie, "Belief and Indeterminacy," *Philosophical Review* 121, 1 (2012): 1–54.

²⁶ *Proof.* The proof is easily adapted from the classical case. In that case, there were three possibilities: worlds at which A and B both take value \mathbf{t} , worlds at which A takes \mathbf{f} while B takes \mathbf{t} , and worlds at which A and B both take value \mathbf{f} . In the present case, the worlds can also be divided into three groups: worlds at which A and B both take a designated value, worlds at which A takes an undesignated value and B takes a designated value, and worlds at which A and B both take an undesignated truth value. Because of (iii) and (iv), b_1 , b_1^\dagger , etc. have the same epistemic values at each of these three possibilities as they have at the corresponding possibility in the classical case. Thus, the reasoning in the classical case transfers to this case. QED.

Richard Pettigrew

- b. That is, b_1 is strictly logically dominated by b^*
- c. Believing A and disbelieving B is weakly logically dominated by disbelieving A and believing B . That is, b_1 is weakly logically dominated by b_1^\dagger .

Believing A and suspending on B is weakly logically dominated by disbelieving A and suspending on B .

That is, b_2 is weakly logically dominated by b_2^\dagger .

We can also apply this in the case of the Logic of Paradox, if we follow Priest's claim that belief (or acceptance) is the correct attitude to a proposition that is assigned truth value \mathbf{u} , which he interprets as *both true and false*.²⁷ In that case, the natural account of local epistemic utility is this:

$$\begin{aligned} eu(\mathbf{t}, \mathbf{B}) &= eu(\mathbf{f}, \mathbf{D}) = eu(\mathbf{u}, \mathbf{B}) = R \\ eu(\mathbf{t}, \mathbf{S}) &= eu(\mathbf{f}, \mathbf{S}) = eu(\mathbf{u}, \mathbf{S}) = 0 \\ eu(\mathbf{t}, \mathbf{D}) &= eu(\mathbf{f}, \mathbf{B}) = eu(\mathbf{u}, \mathbf{D}) = -W \end{aligned}$$

And, since \mathbf{u} is a designated truth value in Logic of Paradox, this account satisfies the hypotheses of Theorem 2. Thus, we have a strict dominance argument for the Logic of Paradox version of (BP4) and weak dominance arguments for the Logic of Paradox versions of (BP5) and (BP6), which are obtained from those principles by replacing \models_{cl} with \models_{lp} .

However, there are other ways in which strong Kleene logic and Logic of Paradox may be applied for which the local epistemic utility functions described so far are not appropriate.²⁸ Suppose, for instance, that strong Kleene logic governs propositions that involve vague predicates.²⁹ Then the appropriate doxastic attitude to a proposition with truth value \mathbf{u} is surely suspension, not disbelief. If the colour of my socks lies in the borderline region between determinately red and determinately orange, it seems better to suspend judgment on the proposition *My handkerchief is red* than to believe or disbelieve it. Thus, we might think that the local epistemic utilities are assigned as follows:

²⁷Graham Priest, *Doubt Truth to be a Liar* (Oxford: Oxford University Press, 2005)

²⁸ Thanks to Hartry Field, Patrick Greenough, and Ole Thomassen Hjortland for helpful discussion on this point.

²⁹ See, for instance: Michael Tye, "Sorites Paradoxes and the Semantics of Vagueness," in *Philosophical Perspectives: Logic and Language*, vol. 8, ed. James Tomberlin (Atascadero: Ridgeview Press, 2008), 189–208, Hartry Field, "No Fact of the Matter," *Australasian Journal of Philosophy* 81 (2003): 457–480.

$$\begin{aligned}
 \text{eu}(\mathbf{t}, \mathbf{B}) &= \text{eu}(\mathbf{f}, \mathbf{D}) = R \\
 &\quad \text{eu}(\mathbf{u}, \mathbf{S}) = N \\
 \text{eu}(\mathbf{t}, \mathbf{S}) &= \text{eu}(\mathbf{f}, \mathbf{S}) = 0 \\
 \text{eu}(\mathbf{u}, \mathbf{B}) &= \text{eu}(\mathbf{u}, \mathbf{D}) = -Z \\
 \text{eu}(\mathbf{t}, \mathbf{D}) &= \text{eu}(\mathbf{f}, \mathbf{B}) = -W
 \end{aligned}$$

where $-W \leq -Z < 0 < N \leq R$. In this case, under certain assumptions, we can again argue for strong Kleene versions of (BP4-6). After all, consider the truth table:

	A	B	EU(b_1, w)	EU(b^*, w)	EU(b_1^\dagger, w)	EU(b_2, w)	EU(b_2^\dagger, w)
w_1	t	t	$R - W$	$0 + 0$	$-W + R$	$R + 0$	$0 + R$
w_2	t	u	$R - Z$	$0 + N$	$-W - Z$	$R + N$	$0 - Z$
w_3	t	f	$R + R$	$0 + 0$	$-W - W$	$R + 0$	$0 - W$
w_4	u	t	$-Z - W$	$N + 0$	$-Z + R$	$-Z + 0$	$N + R$
w_5	u	u	$-Z - Z$	$N + N$	$-Z - Z$	$-Z + N$	$N - Z$
w_6	u	f	$-Z + R$	$N + 0$	$-Z - W$	$-Z + 0$	$N - W$
w_7	f	t	$-W - W$	$0 + 0$	$R + R$	$-W + 0$	$0 + R$
w_8	f	u	$-W - Z$	$0 + N$	$R - Z$	$-W + N$	$0 - Z$
w_9	f	f	$-W + R$	$0 + 0$	$R - W$	$-W + 0$	$0 - W$

If $A \models_{\text{skl}} B$, then w_2 and w_3 do not represent logical possibilities, but the rest do. Thus:

- If $N + Z > R$ and $W \geq R$, b_1 is weakly logically dominated by b^* .
- Even if $N + Z > R$ and $W \geq R$, b_1 is not even weakly logically dominated by b_1^\dagger .
- After all, b_1 outperforms b_1^\dagger when A has truth value **u** and B has truth value **f**. And indeed there are values of $W \geq Z$ and $N \leq R$ such that nothing even weakly logically dominates b_1 for those values.
- If $W - Z = N$, b_2 is weakly logically dominated by b_2^\dagger .

Thus, we have arguments for the strong Kleene versions of (BP4-6), which are obtained from those principles by replacing \models_{cl} with \models_{skl} . But those arguments are weaker than the corresponding arguments for the original, classical versions of (BP4-6), since they make stronger assumptions about local epistemic utilities:

$$(\text{EU8}) (N + Z > R) + (W \geq R) + \text{Weak Logical Dominance} \Rightarrow (\text{BP4}_{\text{skl}}) \text{ and } (\text{BP5}_{\text{skl}}).$$

$$(\text{EU9}) (W + Z = N) + \text{Weak Logical Dominance} \Rightarrow (\text{BP6}_{\text{skl}}).$$

Next, suppose that the Logic of Paradox governs future contingents.³⁰ Thus, the truth value of a future contingent X is: (i) **t** if X is true in all possible futures; (ii) **f** if X is false in all possible futures; and (iii) **u** if X is true in some futures and false in others. The idea is that, in the latter case, the proposition is both true at some point in the future and false at some point in the future, and thus both true and false now. This is a paraconsistent approach to the logic of future contingents. In this situation, we might think it natural to order the local epistemic utilities of the various categorical doxastic attitudes as follows:

$$\begin{aligned}
 \text{eu}(\mathbf{t}, \mathbf{B}) &= \text{eu}(\mathbf{f}, \mathbf{D}) = R \\
 \text{eu}(\mathbf{u}, \mathbf{B}) &= \text{eu}(\mathbf{u}, \mathbf{D}) = N \\
 \text{eu}(\mathbf{t}, \mathbf{S}) &= \text{eu}(\mathbf{f}, \mathbf{S}) = 0 \\
 &\quad \text{eu}(\mathbf{u}, \mathbf{S}) = -Z \\
 \text{eu}(\mathbf{t}, \mathbf{D}) &= \text{eu}(\mathbf{f}, \mathbf{B}) = -W
 \end{aligned}$$

where $-W \leq -Z < 0 < N \leq R$. If we do this, here's the truth table:

	A	B	$\text{EU}(b_1, w)$	$\text{EU}(b^*, w)$	$\text{EU}(b_1^+, w)$	$\text{EU}(b_2, w)$	$\text{EU}(b_2^+, w)$
w_1	t	t	$R - W$	$0 + 0$	$-W + R$	$R + 0$	$0 + R$
w_2	t	u	$R + N$	$0 - Z$	$-W + N$	$R - Z$	$0 + N$
w_3	t	f	$R + R$	$0 + 0$	$-W - W$	$R + 0$	$0 - W$
w_4	u	t	$N - W$	$-Z + 0$	$N + R$	$N + 0$	$-Z + R$
w_5	u	u	$N + N$	$-Z - Z$	$N + N$	$N - Z$	$-Z + N$
w_6	u	f	$N + R$	$-Z + 0$	$N - W$	$N + 0$	$-Z - W$
w_7	f	t	$-W - W$	$0 + 0$	$R + R$	$-W + 0$	$0 + R$
w_8	f	u	$-W + N$	$0 - Z$	$R + N$	$-W - Z$	$0 + N$
w_9	f	f	$-W + R$	$0 + 0$	$R - W$	$-W + 0$	$0 - W$

If $A \models_{\text{lp}} B$, then worlds w_3 and w_6 fail to represent logical possibilities. But given this, we can see that no alternative weakly or strictly dominates b_1 . The reason is that no alternative belief function performs as well as b_1 at world w_2 . Similarly, no alternative either weakly or strictly dominates b_2 . The only alternatives that perform as well as b_2 at w_2 assign **B** to A and either **B** or **D** to B ; but these perform worse than b_2 at worlds w_9 and w_7 , respectively. So we do not have an Logic of Paradox versions of (BP4-6) in this case.

Indeed, this all follows from a more general fact:³¹

³⁰ Roberto Ciuni and Carlo Proietti, "The Abundance of the Future: A Paraconsistent Approach to Future Contingents," *Logic and Logical Philosophy* 22, 1 (2013): 21–43.

³¹ *Proof.* The truth value **i** from (4) is either designated or undesignated. Suppose first that it is designated. Then the only alternative belief function that performs at least as well as b_1 at the logical possibility where A takes **t** and B takes **i** is the belief function that assigns belief to A and belief to B . But that performs worse than b_1 when A and B both take **f**. Next, suppose that it is

Theorem 3

Suppose:

- i. The logical consequence relation for a many-valued logic k is defined in terms of the preservation of designated truth values.
- ii. $A \vDash_k B$
- iii. eu_k extends eu_{cl} .
- iv. There is a truth value \mathbf{i} such that
- v. $-W \leq eu_k(\mathbf{i}, \mathbf{S}) < eu_k(\mathbf{i}, \mathbf{B}) = eu_k(\mathbf{i}, \mathbf{D}) \leq R$

Then:

- a. No alternative even weakly logically dominates believing A and disbelieving B .

That is, nothing weakly dominates b_1 .

The upshot of this section is that the fate of logic-rationality bridge principles is sensitive to the logic that governs the propositions in question, the interpretation of the truth values in that logic, and the resulting assignments of epistemic value to beliefs, disbeliefs, and suspensions in propositions that take truth-values other than \mathbf{t} or \mathbf{f} . This explains why we have been careful throughout to specify in the antecedent of those principles which logic governs the propositions in question. There are bridge principles that hold when the logic is classical that do not hold for alternative logics.

Logical, Doxastic, and Epistemic Possibilities

In the preceding sections, we have offered epistemic utility arguments in favour of certain logic-rationality bridge principles, and we have given epistemic utility-based reasons for doubting that others can be justified. For each of the bridge principles we have considered, its antecedent is a plain fact about logical consequence—something of the form $\vDash_k \textit{governs } A_1, \dots, A_n, B, \textit{ and } A_1, \dots, A_n \vDash_k B$ for some logic k and some $n \geq 1$. As a result, the principles are quite demanding. They require that you manage your beliefs in line with a logical fact that you might not know or even believe. We have been able to justify these demanding principles only because we've assumed similarly demanding principles of decision theory. For instance, Strict Logical Dominance says that an option is irrational if

undesigned. Then the only belief function that performs at least as well as b_1 at the logical possibility where A takes \mathbf{i} and B takes \mathbf{f} is the belief function that assigns disbelief to A and disbelief to B . But that performs worse than b_1 when A and B both take \mathbf{t} . QED.

there is an alternative that is better at all *logically possible worlds*; Weak Logical Dominance says that an option is irrational if there is an alternative that is at least as good at all *logically possible worlds*, and better at some. And you might think that these are too strong, even in the practical case. For instance, Strict and Weak Logical Dominance render it irrational for my nine year old niece to pay any positive amount for a bet against Fermat's Last Theorem, even if she has never heard of it until I describe it to her, and even if I tell her nothing about its proof status. If we weaken these decision-theoretic principles, we obtain epistemic utility arguments for the correspondingly weakened logic-rational bridge principles.

Here are general versions of our dominance norms, where \mathcal{C} is a set of worlds.

Strict \mathcal{C} Dominance If option o^* has greater utility than option o at every world in \mathcal{C} , then o is irrational.

Weak \mathcal{C} Dominance If option o^* has at least as great utility as option o at every world in \mathcal{C} , and greater utility at some, then o is irrational.

We obtain Strict/Weak Logical Dominance if \mathcal{C} is the set of logically possible worlds. We obtain Strict/Weak Doxastic Dominance if \mathcal{C} is the set of doxastically possible worlds—that is, the worlds at which everything she believes is true and everything she disbelieves is false. And we obtain Strict/Weak Epistemic Dominance if \mathcal{C} is the set of epistemically possible worlds—that is, the worlds compatible with what the agent knows. And so on.

Now, suppose we replace Strict/Weak Logical Dominance with Strict/Weak Doxastic or Epistemic Dominance in our arguments for logic-rationality bridge principles. Then surely we obtain arguments for the corresponding bridge principles in which the antecedent is no longer just a proposition about logical consequence, but is rather the proposition that the agent believes or knows that proposition about logical consequence.³²

Thus, for instance, let's assume Weak Doxastic Dominance:

Weak Doxastic Dominance If option o^* has at least as great utility as option o at every doxastically possible world, and greater utility at some, then o is irrational.

Then, in order to adapt argument (EU2), we need to assume that the agent believes that classical logic is the correct logic and that A is strictly stronger than B in that logic. By believing that classical logic is the correct logic, our agent narrows down the set of doxastically possible worlds to the four— w_1 , w_2 , w_3 , w_4 —represented in the relevant table above; by also believing that A is strictly stronger

³² J. R. G. Williams, "Rational Illogicality" *Australasian Journal of Philosophy* (forthcoming).

than B , our agent narrows the field further by ruling out world w_2 at which A is true and B is false, but retains world w_3 at which A is false and B is true—that is, she narrows the field to w_1, w_3, w_4 . We thus obtain:

(BP10) If you believe that \models_{cl} governs A and B , and $A \models_{cl} B$, and $B \not\models_{cl} A$, then you ought not to believe A while disbelieving in B .

Or so it seems. The problem with using doxastic possibility in this context is that we are using facts about what we believe to delimit the set of worlds that feature in a dominance principle that we then use to choose our beliefs! Why might this be problematic? Initially, you might think that it could give rise to a sort of instability in the beliefs it is rational for you to have. You start with a set of beliefs. They determine the worlds that are doxastically possible for you. Determined in this way, it turns out that epistemic utility theory rules out your beliefs as irrational. So you pick another set of beliefs. They determine the worlds that are doxastically possible for you. Determined in this way, it turns out that epistemic utility theory rules out those beliefs as irrational. And so on. At first sight, this seems a possibility. But, in fact, it is the opposite that happens. Pick a set of consistent beliefs and disbeliefs. These then determine a set of doxastically possible worlds. At each of these worlds, each of the beliefs you picked is true and each of the disbeliefs you picked is false. Thus, you have maximal epistemic utility at each of these worlds. Any alternative assignment of beliefs, suspensions, and disbeliefs to the same propositions will be weakly doxastically dominated. Thus, any consistent set of beliefs renders itself the only rational option.

Here's another way in which it might be problematic to use beliefs to determine the doxastic possibilities, and then use those possibilities to pick the beliefs. Suppose I believe that A entails B , and I believe A , but I disbelieve B . Then there are no worlds at which all of my beliefs are true and all my disbeliefs false. Thus, there are no worlds that are doxastically possible for me. One consequence of that is that every belief function is strictly doxastically dominated by every other one — for every pair of belief functions b and b' on the same set of propositions, it is vacuously true that b is strictly better than b' at all doxastically possible worlds, for there are no doxastically possible worlds. Thus, unless we restrict Strict Doxastic Dominance, every belief function is irrational. In fact, we're best to restrict Strict Doxastic Dominance in this case, and say that dominance principles only apply when the relevant set of worlds is non-empty. But if we do this, nothing is ruled irrational for the agent who believes that A entails B , believes A , and disbelieves B . And that looks troubling too!

A final worry about moving to Strict/Weak Doxastic Dominance principles. The norms that result from the epistemic utility arguments that appeal to those

principles are narrow scope norms of the sort that we typically reject in this area. Consider the standard narrow scope norm in this area: if you believe *If A, then B*, and you believe that *A*, then you ought to believe that *B*. As Harman noted in *Change in View*, such a norm cannot be correct, since it is just as legitimate to respond by dropping your belief that *If A, then B* or by dropping your belief that *A* as it is to respond by keeping both of those beliefs and further adopting a belief that *B*. Similarly, surely the logic-rationality bridge principles that follow from the doxastic dominance arguments cannot be correct either. Surely it is just as legitimate to respond to your belief that *A and B are governed by* \models_{cl} and your belief that $A \models_{cl} B$ by dropping one or other or both of those beliefs as it is to retain both beliefs and then ensure that you do not believe *A* and disbelieve *B*.

I offer two different solutions to these problems. First, the *Two-Tier Solution*. We might save our new doxastic dominance arguments for the doxastic versions of the bridge principles if we take our beliefs about logical consequence to be of a rather different sort from our beliefs about other matters. We might take the beliefs about logical consequence to delimit the doxastically possible worlds, perhaps, and then use those in our dominance principles to assess the different possible sets of beliefs we might have towards other propositions. If you opt for this solution, you owe an account of why there are two sorts of beliefs, ones that get to delimit doxastic possibilities and ones that don't. And you have to say, in particular, why logical beliefs—beliefs about the logic that governs a class of propositions, and beliefs about the consequence relation of that logic—are of the former sort. You might appeal, for instance, to Quine's notion of a web of belief.³³ It is perhaps the propositions sufficiently close to the centre of our web of belief—that is, those least vulnerable to revision—that delimit the set of doxastic possibilities. It is then those further out—those more vulnerable to revision—that are governed by Strict/Weak Doxastic Dominance. This would provide a principled distinction between the two sorts of belief, and it would also explain why the narrow scope norm is appropriate. It explains why it is legitimate to demand that you respond to your logical belief that *A and B are governed by* \models_{cl} and your logical belief that $A \models_{cl} B$ by ensuring that you do not believe *A* and disbelieve *B*, rather than by dropping one or other or both of your logical beliefs. The explanation is that the logical beliefs lie closer to the centre of the web of belief—when something's got to give, it shouldn't be them.

Here's the second solution to the problems raised above for the doxastic dominance arguments for the doxastic versions of the logic-rationality bridge

³³ W. V. O. Quine, "Two Dogmas of Empiricism," *Philosophical Review* 60, 1 (1951): 20–43. W. V. O. Quine and Joe Ullian, *The Web of Belief* (New York: Random House, 1970).

principles we've been considering. We might call it the *Wide Scope solution*. It proceeds by analogy with the standard retreat from narrow to wide scope norms in the face of Harman's criticism. Thus, instead of trying to justify the narrow scope norm (BP10), we might try to justify the wide scope version of it:

(BP11) You ought to see to it that you don't believe that \models_{cl} governs A and B , believe $A \models_{cl} B$, believe $B \not\models_{cl} A$, believe A , but disbelieve B .

Can we offer an epistemic utility argument for (BP11)? We have five propositions in play: (i) \models_{cl} governs A and B ; (ii) $A \models_{cl} B$; (iii) $B \not\models_{cl} A$; (iv) A ; and (v) B . (BP11) says that you ought not to believe (i)-(iv) whilst disbelieving (v). But now notice that (i)-(iv) entail (v). So, we have a multi-premise entailment and we wish to justify a norm that prohibits believing the premises and disbelieving the conclusion. Thus, if we simply treat each of these propositions as a normal proposition, and if it is legitimate to assume that any agent can see the entailment from (i)-(iv) to (v), and that is something they should never give up, then we can simply turn our Strict Logical Dominance argument for (BP8) into a Strict Doxastic Dominance argument for (BP11). Now, notice that the multi-premise entailment in question is a four-premise entailment. So in order to run our dominance argument for it, we need to assume a version of Extreme Epistemic Conservatism, namely, that $4R < W$. But if we have that, then we can conclude (BP11).

Thus, we have two putative solutions to our problems. On the Two-Tier solution, we retain the narrow scope norm (BP10) by saying that some propositions—including those that pertain to the correct logic and the consequence relation of that logic—fix the doxastic possibilities, while others are determined after those possibilities have been fixed by considerations of epistemic utility. On the Wide Scope solution, we do not assign the logical propositions any special role, and instead treat them just like other propositions, giving us the wide scope version (BP11).

So much for our solutions to the problems raised above. In the remainder of this section, we make a handful of further observations on the move from logic-rationality bridge principles with purely logical antecedents to the versions with doxastic antecedents. First, we note that there are two ways in which you might not know or believe all the logical truths. You might know what the correct logic is—for instance, you might know that classical logic governs A and B —but you might not know that A entails B within that logic. But you might not even know what the correct logic is—you might not know whether strong Kleene logic, classical logic, or Logic of Paradox governs A and B . Above, we focussed on the first sort of case, assuming that there was some particular logic that our agents

believed to be the correct one. But there are some things to say about the second case as well.

Suppose, for instance, that our agent believes that the correct logic is either strong Kleene logic or classical logic; suppose she knows that A entails B ; and suppose $N + Z > R$ and $W \geq R$; then it would be irrational for her to believe A and disbelieve B —that is, irrational for her to have belief function b_1 . After all, if we pool all of the worlds that are possible relative to classical logic and all of the worlds that are possible relative to strong Kleene logic, then b^* dominates b_1 . The reason is that every classically possible world is also logically possible from the point of view of strong Kleene logic. Thus, since b^* dominates b_1 relative to the strong Kleene worlds, it dominates b_1 relative to all the strong Kleene *and* classical worlds.

That might tempt us to think that if a logic-rationality bridge principle can be justified by logical dominance reasoning relative to one logic and justified by logical dominance reasoning also relative to another logic, then it can be justified by doxastic dominance reasoning for someone who believes that one or other of these logics is correct, but isn't certain which. But that is not the case. The reason: it might be that a belief function b is dominated by b' and not by b'' relative to the first logic, while it is dominated by b'' and not by b' relative to the second. In that case, neither b' nor b'' dominates b relative to the disjunction of the logics. In this case, we have a situation akin to the Miners Paradox.³⁴ If the first logic is actual, the agent ought not to choose b (since it is dominated by b'); if the second logic is actual, the agent ought not to choose b (since it is dominated by b''); the first logic is actual or the second is; but it does not follow that the agent ought not to choose b .³⁵

My next observation on logic-rationality bridge principles with doxastic antecedents concerns the argument due to MacFarlane that we should not be satisfied with them. The problem with these principles, MacFarlane argues, is this: if they are the strongest norms in the vicinity, it seems that the less you know, logically speaking, the less restricted are your beliefs; by remaining ignorant of logical facts, you are less likely to be irrational, since less stringent restrictions are placed upon you. And this seems counterintuitive. It seems to give an incentive to remain logically uninformed. But this should be nothing new. For many philosophers—subjectivist Bayesian epistemologists, for instance—the less evidence you have, whether logical or not, the fewer restrictions are placed upon you. But this only gives an incentive to remain uninformed if avoiding irrationality

³⁴ Derek Parfit, "What We Together Do" (unpublished manuscript).

³⁵For a related discussion, see J. R. G. Williams, "Rational Illogicality."

is the only thing you care about. And of course it is not. You also care about making good decisions and having accurate beliefs. In the case of non-logical facts, the *value of learning theorem* due to I. J. Good and Frank Ramsey shows that you can expect to make better decisions after you have learned those facts than before;³⁶ and it is straightforward to adapt the epistemic utility-based argument for Conditionalization to show that you can expect to have more accurate beliefs after you learn non-logical facts than before.³⁷ Now, there is no reason why these arguments shouldn't apply to learning logical facts as well. Thus, we can take the doxastic versions of the logic-rationality bridge principles to be the strongest principles in the vicinity, whilst also thinking that agents should try to know as many logical facts as possible, and should then manage their beliefs in line with the logical facts that they believe. However, their reasons for doing so are just their usual reasons for learning, and then managing their beliefs in line with what they've learned.

This leads us to our final point in this section. Given a particular logic-rationality bridge principle, we can ask what role is played by the logicity of the fact about logical consequence that appears in the principle's antecedent.³⁸ Are there principles of rationality of the same form that feature a non-logical fact in the antecedent? Does the argument for the bridge principle pay any special attention to the logicity of the fact about logical consequence? The structure of the answer depends, I think, on whether you embrace the Two-Tier solution or the Wide Scope solution above; but the conclusion doesn't. Whichever of those solutions you choose, the logicity of the logical facts plays no special role. Let's see why. First, suppose you opt for the Two-Tier solution. That is, you say that there are two different roles that beliefs can play: they can circumscribe the set of doxastic possible worlds, and they can be evaluated for their epistemic utility. What's more, you say that all logical beliefs play the first role, while some non-logical beliefs play the second. But there is no reason to suppose that it is *only* the logical beliefs that can delimit the doxastically possible worlds. And, of course, if we spell out this solution by saying that it is beliefs near to the centre of the web of belief that play the delimiting role, then presumably beliefs concerning analytic or conceptual truths, mathematical truths, or metaphysical necessities will fit the bill

³⁶ I. J. Good, "On the Principle of Total Evidence," *The British Journal for the Philosophy of Science* 17 (1967): 319–322, Frank P. Ramsey, "Weight or the Value of Knowledge," *The British Journal for the Philosophy of Science* 41 (1990): 1–4.

³⁷ Hilary Greaves and David Wallace, "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility," *Mind* 115, 459 (2006): 607–632.

³⁸This question also runs through Harman, *Change in View*.

just as well as logical beliefs do. Thus, the logicity of the logical beliefs plays no special role—what is relevant is their location in the web of belief, and plenty of non-logical beliefs occupy nearby locales. Next, suppose you opt for the second solution to the problems outlined at the beginning of this section. Then the irrelevance of the logicity of those beliefs is even more stark. After all, on that solution, we justify the wide scope norm (BP11). But our argument for that appeals only to the obvious entailment from (i)-(iv) to (v) above. It does not appeal at any point to the logicity of (i)-(iii). The argument would run just as well if (i)-(iii) were any premises for which the entailment from those, together with (iv), to (v) is sufficiently obvious. Thus, on both solutions there is nothing about the logicity of the logic-rationality bridge principles that plays a role in our arguments.

Bridge Principles for Partial Beliefs

In the previous section, we asked what we can learn about logic-rationality bridge principles for categorical doxastic states, such as full belief, full disbelief, and suspension of judgment, by looking at the epistemic utility of those doxastic attitudes. In this section, we turn our attention to partial beliefs, or credences as we will call them.

As before, we begin with the now-standard story about the classical case.³⁹ Suppose our agent has a credence function c defined on a set of propositions \mathcal{F} . For each proposition A in \mathcal{F} , $c(A)$ gives the agent's credence in A . By convention, we take maximal credence to be 1 and minimal credence to be 0. Thus, $c: \mathcal{F} \rightarrow [0,1]$. Now, as above, we must define an epistemic utility function EU that takes a credence function c and a possible world w and returns $EU(c, w)$, a measure of the epistemic utility of having credence function c at world w . As above, we take it to be additive. That is, we assume that there is a local epistemic utility function $eu: \{\mathbf{t}, \mathbf{f}\} \times [0,1] \rightarrow [-\infty, \infty]$ such that

$$EU(c, w) = \sum_{X \in \mathcal{F}} eu(w(X), c(X))$$

How do we define eu ? We do so in two steps. First, for each proposition A and each possible world w , we take there to be an *ideal* or *perfect* or *vindicated* credence in A at w —we call this $v_w(A)$. Now, as in the case of categorical doxastic attitudes, the standard story in the credal case takes a veritist approach—that is, it assumes that the sole fundamental source of epistemic value for doxastic states is

³⁹ Joyce, “A Nonpragmatic Vindication,” Joyce, “Accuracy and Coherence,” Pettigrew, *Accuracy and the Laws of Credence*.

the accuracy with which they represent the world. In the classical case, this suggests:

$$v_w(A) = \begin{cases} 1 & \text{if } v_w(A) = \mathbf{t} \\ 0 & \text{if } v_w(A) = \mathbf{f} \end{cases}$$

We might call this assumption about the ideal credences *Vindicated is Omniscient*.

Second, having defined the ideal credence in a given proposition at a given possible world, we can then define the epistemic utility of credence $c(A)$ at world w to be its proximity to $v_w(A)$. That is, the epistemic *disutility* of $c(A)$ at w is the distance from $v_w(A)$ to $c(A)$. How are we to measure distance between credence functions? There are various arguments for measuring such distances using the so-called *Bregman divergences*.⁴⁰ A divergence is a function \mathfrak{d} that takes a pair of real numbers x and y and returns $\mathfrak{d}(x, y)$, a non-negative real number or ∞ . We say that \mathfrak{d} is a divergence iff $\mathfrak{d}(x, y) \geq 0$ with equality iff $x = y$. And we say that \mathfrak{d} is a *Bregman divergence* if there is a strictly convex, continuously differentiable function $\varphi: [0,1] \rightarrow [0, \infty)$ such that:

$$\mathfrak{d}(x, y) = \varphi(x) - \varphi(y) - \varphi'(y)(x - y)$$

where φ' is the derivative of φ . That is, the divergence from x to y is the difference between the value at x of φ and the value at x of the tangent to φ taken at y . If eu is a local epistemic utility function, we demand that it is generated by a Bregman divergence \mathfrak{d} as follows:

$$eu(w(A), c(A)) = -\mathfrak{d}(v_w(A), c(A))$$

That is, the epistemic utility of a credence $c(A)$ in proposition A at world w is the negative of the divergence from $v_w(A)$, the ideal credence in A at w , to $c(A)$. Putting all of this together, we have that

$$EU(c, w) = \sum_{X \in \mathcal{F}} eu(w(X), c(X)) = - \sum_{X \in \mathcal{F}} \mathfrak{d}(v_w(X), c(X))$$

A well known result shows that the epistemic utility functions defined in this way are precisely the so-called *additive and continuous strictly proper inaccuracy measures*. We might call this assumption about epistemic utility *Bregman Divergence*.

⁴⁰ Joyce, "Accuracy and Coherence," Pettigrew, *Accuracy and the Laws of Credence*, Benjamin Levinstein, "A Pragmatist's Guide to Epistemic Utility," *Philosophy of Science* (forthcoming), Sophie Horowitz, "Accuracy and Educated Guesses," *Oxford Studies in Epistemology* (forthcoming).

Now, it is natural to ask what logic-rationality bridge principles follow from this account of the epistemic utility of credences when we apply the decision-theoretic principles that we considered above. The following is a well-known result:⁴¹

Theorem 4

Suppose:

- i. EU is an additive and continuous strictly proper inaccuracy measure.

That is, $EU = \sum_{X \in \mathcal{F}} eu(w(X), c(X)) = -\sum_{X \in \mathcal{F}} \mathfrak{d}(v_w(X), c(X))$, where \mathfrak{d} is a Bregman divergence.

Then:

- a. If c is not a probability function on \mathcal{F} , then there is a probability function c^* on \mathcal{F} such that c^* strictly logically dominates c relative to EU—that is, $EU(c, w) < EU(c^*, w)$, for all logically possible worlds w .

Now, a probability function on \mathcal{F} is a credence function c that satisfies the following conditions:

(BP11a) If \models_{cl} governs A , and $\models_{cl} A$, then $c(A) = 1$

(BP11b) If \models_{cl} governs A , and $A \not\models_{cl}$, then $c(A) = 0$

(BP12) If \models_{cl} governs A, B , and $A \models_{cl} B$, then $c(A) \leq c(B)$

(BP13) If \models_{cl} governs A, B , then $c(A \& B) + c(A \vee B) = c(A) + c(B)$

The first three are logic-rationality bridge principles: they concern how credences should behave given facts about the consequence relation. The fourth is not: it concerns the interaction between credences in propositions of different logical forms. Williams calls (BP12) the *No Drop principle*. It is the credal analogue to principles like (Wo-) from MacFarlane⁴² and (BP4-7) from above. Thus, we have the following epistemic utility argument:

(EU9) Bregman Divergence + Vindicated is Omniscient + Strict Logical Dominance \Rightarrow (BP11-13).

Before we leave the classical case, it is worth sketching the proof of Theorem 4, since that will show us how that proof might be adapted to the non-classical case. The proof is based on two lemmas. First:

⁴¹ Joel Predd, Robert Seiringer, Elliott Lieb, Daniel Osherson, Vincent Poor, and Sanjeev Kulkarni, “Probabilistic Coherence and Proper Scoring Rules,” *IEEE Transactions of Information Theory* 55, 10 (2009): 4786–4792.

⁴² MacFarlane, “In What Sense (If Any).”

Lemma 5 The set of probability functions is precisely the *closed convex hull* of the set of vindicated credence functions, v_w , for possible worlds w .⁴³

Second:

Lemma 6 Suppose \mathfrak{d} is a Bregman divergence, and $\mathcal{X} \subseteq [0,1]^n$ is a set of n -dimensional vectors. Then, if z is a point in $[0,1]^n$ that lies outside the closed convex hull of \mathcal{X} , then there is a point z^* inside the convex hull of \mathcal{X} such that $\sum_{i=1}^n \mathfrak{d}(x_i, z_i^*) < \mathfrak{d}(x_i, z_i)$ for all x in \mathcal{X} .

Thus, suppose c is a credence function that is not a probability function. Then, by Lemma 5, c lies outside the closed convex hull of the vindicated credence functions. Then, by Lemma 6, there is c^* in the convex hull of the vindicated credence functions such that $\mathfrak{d}(v_w(X), c^*(X)) < \mathfrak{d}(v_w(X), c(X))$. And thus, $EU(c, w) < EU(c^*, w)$, for all w .

Breaking down the result into these two component parts allows us to see how it might be generalised. Suppose we move to a different logic. And, as a result, we take different credences to be vindicated—that is, we define $v_w(A)$ differently from how we defined it above. Suppose further that we continue to measure the local epistemic utility of a credence $c(A)$ as its proximity to the vindicated credence: that is, $eu(w(A), c(A)) = -\mathfrak{d}(v_w(A), c(A))$. Then we can find the bridge principle for which strict dominance provides an epistemic utility argument as follows:

First, we characterise the closed convex hull of those new vindicated credence functions—that is, we provide an analogue of Lemma 5.

Second, we note that, if a credence function c lies outside this closed convex hull, then there is an alternative c^* that is closer to each of the vindicated credence functions than c , and thus epistemically better than c at all logically possible worlds—that is, we deploy Lemma 6.

We will see exactly this strategy in action below.

Before we see it in action in the non-classical case, we first observe it in the classical case. We can use Lemma 6 to justify the following bridge principle, (BP14). (BP14) is the general bridge principle for credences that Field⁴⁴ defends, drawing on Adams and Edgington⁴⁵.

⁴³ If \mathcal{X} is a set of credence functions, then its convex hull is the smallest convex set that contains \mathcal{X} ; that is, the smallest set that contains \mathcal{X} and contains every mixture of two credence functions whenever it contains those credence functions. We denote the convex hull \mathcal{X}^+ . The closure of a set is the union of that set with the set of its limit points.

⁴⁴ Hartry Field, "What Is the Normative Role of Logic?" *Proceedings of the Aristotelian Society* (Supplementary Volumes) 83 (2009): 251–268.

⁴⁵ Ernest W. Adams, *The Logic of Conditionals* (Dordrecht: Reidel, 1975), Dorothy Edgington,

Richard Pettigrew

(BP14) If \models_{cl} governs A_1, \dots, A_n, B , and $A_1, \dots, A_n \models_{cl} B$, then $c(A_1) + \dots + c(A_n) - (n - 1) \leq c(B)$

Or, equivalently and more intuitively:

(BP14) If \models_{cl} governs A_1, \dots, A_n, B , and $A_1, \dots, A_n \models_{cl} B$, then $\bar{c}(B) \leq \bar{c}(A_1) + \dots + \bar{c}(A_n)$

where $\bar{c}(X) = 1 - c(X)$ measures an agent's degree of disbelief in X when $c(X)$ measures her degree of belief in X .

Here's the argument: it is easy to see that, if \models_{cl} governs A_1, \dots, A_n, B and $A_1, \dots, A_n \models_{cl} B$, then, for each logically possible world w , the ideal credence function v_w satisfies (BP14). That is,⁴⁶

$$v_w(A_1) + \dots + v_w(A_n) - (n - 1) \leq v_w(B)$$

What's more, whenever two credence functions satisfy (BP14), so does every convex combination of them. And whenever each credence function in an infinite sequence satisfies (BP14), so does the limit of that sequence. Thus, every credence function in the closed convex hull of the ideal credence functions satisfies (BP14). And thus, by Lemma 6, any credence function that violates (BP14) is strictly logically dominated. This establishes (BP14).

Non-Classical Logics

What happens when we move from classical logic to a non-classical alternative? The key issue here is to determine, for each possible world w , what the vindicated credence function v_w is at that world. In the classical case,

$$v_w(A) = \begin{cases} 1 & \text{if } w(A) = \mathbf{t} \\ 0 & \text{if } w(A) = \mathbf{f} \end{cases}$$

But what about the non-classical case? Here is one suggestion:⁴⁷

$$v_w(A) = \begin{cases} 1 & \text{if } w(A) \text{ is designated} \\ 0 & \text{if } w(A) \text{ is not designated} \end{cases}$$

We might call this the *Vindicated is Designated* condition on epistemic utility. Notice that this is analogous to the suggestion in the full belief case that, if a proposition has designated truth value, then belief is the ideal categorical doxastic

⁴⁶ "On conditionals," *Mind* 104 (1995): 235–329.

⁴⁶ *Proof.* There are two cases. First, if there is A_i such that $v_w(A_i) = 0$, then, since $v_w(A_j) \leq 1$ for all $1 \leq j \leq n$, $v_w(A_1) + \dots + v_w(A_n) - (n - 1) \leq 0 \leq v_w(B)$. Second, if $v_w(A_i) = 1$, for all A_i , then $v_w(B) = 1$ and $v_w(A_1) + \dots + v_w(A_n) - (n - 1) = 1 = v_w(B)$. QED.

⁴⁷ J. R. G. Williams, "Gradational Accuracy and Non-Classical Semantics," *Review of Symbolic Logic* 5, 4 (2012): 513–537.

attitude, with value R , while disbelief takes value $-W$, and if that proposition has a non-designated truth value, then disbelief is the ideal attitude, with value R , while belief takes value $-W$. In the credal case, we can then appeal to a result due to Jeff Paris to provide epistemic utility arguments for various bridge principles for a wide variety of non-classical logics.⁴⁸

Theorem 7

Suppose:

- i. The logical consequence relation for a many-valued logic k is defined in terms of the preservation of designated truth values.
- ii. $w(X \& Y)$ is designated iff $w(X)$ and $w(Y)$ are both designated.
- iii. $w(X \vee Y)$ is designated iff $w(X)$ or $w(Y)$ is designated.
- iv. $v_w(A) = \begin{cases} 1 & \text{if } w(A) \text{ is designated} \\ 0 & \text{if } w(A) \text{ is not designated} \end{cases}$
- v. $EU(c, w) = \sum_{X \in \mathcal{F}} eu(w(X), c(X)) = -\sum_{X \in \mathcal{F}} d(v_w(X), c(X))$

Then:

- a. c is strictly logically dominated if it is not a *generalized probability function for logic k*.

That is, c is strictly logically dominated if it fails to satisfy any of the following bridge principles:

- (BP14a) If \models_k governs A , then $\models_k A$, then $c(A) = 1$
- (BP14b) If \models_k governs A , then $A \models_k$, then $c(A) = 0$
- (BP15) If \models_k governs A, B , and $A \models_k B$, then $c(A) \leq c(B)$
- (BP16) If \models_k governs A, B , then $c(A \& B) + c(A \vee B) = c(A) + c(B)$

In fact, we only require (ii) and (iii) in order to infer (BP16), which is not a logic-rationality bridge principle. Thus, if we are interested only in the bridge principles, we can prove a more general theorem. This is the credal analogue to Theorem 2:⁴⁹

⁴⁸*Proof.* Paris proves that, if logic k satisfies (i), (ii), and (iii), and if v_w is defined as in (iv), then the closed convex hull of the set of vindicated credence functions is precisely the set of credence functions that satisfy (BP14-16). We then simply apply Lemma 6 to obtain the theorem. QED.

⁴⁹*Proof.* Note that, if (i) and (ii) hold, then (BP14a), (BP14b), and (BP15) are all satisfied by each of the vindicated credence functions. What's more, when those conditions are satisfied by two credence functions, they are also satisfied by any convex combination of them; and when they are satisfied by each credence function in a sequence, they are also satisfied by the limit, if such exists. Thus, they are satisfied by everything in the closed convex hull of the vindicated credence

Theorem 8

Suppose:

- i. The logical consequence relation for a many-valued logic k is defined in terms of the preservation of designated truth values.
- ii. $v_w(A) = \begin{cases} 1 & \text{if } w(A) \text{ is designated} \\ 0 & \text{if } w(A) \text{ is not designated} \end{cases}$
- iii. $EU(c, w) = \sum_{X \in \mathcal{F}} eu(w(X), c(X)) = - \sum_{X \in \mathcal{F}} d(v_w(X), c(X))$

Then:

- a. c is weakly dominated if it fails to satisfy any of the following bridge principles:

(BP14a) If \models_k governs A , then $\models_k A$, then $c(A) = 1$

(BP14b) If \models_k governs A , then $A \models_k$, then $c(A) = 0$

(BP15) If \models_k governs A, B , and $A \models_k B$, then $c(A) \leq c(B)$

Thus, we have the following epistemic utility argument for the logics in question:

(EU10) Bregman Divergence + Vindicated is Designated + Strict Logical Dominance \Rightarrow (BP14-15).

And, as in the classical case, we can also establish

(BP17) If \models_k governs A_1, \dots, A_n, B , and $A_1 \dots, A_n \models_k B$, then $c(A_1) + \dots + c(A_n) - (n - 1) \leq c(B)$

That is,

(BP17) If \models_k governs A_1, \dots, A_n, B , and $A_1 \dots, A_n \models_k B$, then $\bar{c}(B) \leq \bar{c}(A_1) + \dots + \bar{c}(A_n)$

So we obtain Field's logic-rationality bridge principle for all such logics.

However, as in the full belief case, while this may be the correct account of epistemic value for Field's use of strong Kleene logic or Priest's use of Logic of Paradox, it is not obviously the correct account for the application of strong Kleene logic to vague propositions nor the application of Logic of Paradox to propositions concerning future contingents. But, as the following theorem shows, as soon as we abandon this account of epistemic value, we lose the No Drop principle, (BP15), and with it the logic-rationality bridge principle that Field endorses, namely, (BP17). This is the analogue of Theorem 3.⁵⁰

functions. QED.

⁵⁰ *Proof.* The truth value i from (iv) is either designated or undesignated. Suppose first that i is

Theorem 9

Suppose:

- i. The logical consequence relation for a many-valued logic k is defined in terms of the preservation of designated truth values.
- ii. $A \models_k B$
- iii. If $w(X) = \mathbf{t}$, then $v_w(X) = 1$; and if $w(X) = \mathbf{f}$, then $v_w(X) = 0$.
- iv. There is truth value \mathbf{i} such that, if $w(X) = \mathbf{i}$, then $0 < v_w(X) < 1$.
- v. $EU(c, w) = \sum_{X \in \mathcal{F}} eu(w(X), c(X)) = -\sum_{X \in \mathcal{F}} \delta(v_w(X), c(X))$

Then:

- a. There is an undominated credence function c such that $c(A) > c(B)$.

For instance, suppose we define v_w as follows for strong Kleene logic:

$$v_w(A) = \begin{cases} 1 & \text{if } w(A) = \mathbf{t} \\ \frac{1}{2} & \text{if } w(A) = \mathbf{u} \\ 0 & \text{if } w(A) = \mathbf{f} \end{cases}$$

This seems natural when the propositions in question include vague properties and they are governed by strong Kleene logic. If we do this, we can no longer establish the following version of (BP15):

(BP15_{skl}) If \models_{skl} governs A, B , and $A \models_{skl} B$, then $c(A) \leq c(B)$.

Suppose $A \models_{skl} B$. This does not preclude a possible world w such that $w(A) = \mathbf{u}$, but $w(B) = \mathbf{f}$. But in that world $v_w(A) = \frac{1}{2}$ and $v_w(B) = 0$. Thus, $v_w(A) > v_w(B)$.

The upshot of this section is similar to the upshot of our earlier section on bridge principles for beliefs in the presence of non-classical logics: the fate of logic-rationality bridge principles is sensitive to the logic that governs the propositions in question, the interpretation of the truth values in that logic, and the credences we thereby identify as vindicated.

designated. Then there is a logical possibility w^* , where A takes \mathbf{t} and B takes \mathbf{i} . And in this case $v_{w^*}(B) < v_{w^*}(A)$, by (iv). Since v_{w^*} is a vindicated credence function, it is in the closed convex hull of the vindicated credence functions. Thus, it is undominated. Next, suppose that \mathbf{i} is undesignated. Then there is a logical possibility w^\dagger , where A takes \mathbf{i} and B takes \mathbf{f} . And in this case $v_{w^\dagger}(B) < v_{w^\dagger}(A)$, by (iv). Since v_{w^\dagger} is a vindicated credence function, it is in the closed convex hull of the vindicated credence functions. Thus, it is undominated. QED.

Conclusion

In this paper, we have explored a novel way to adjudicate between the vast variety of putative logic-rationality bridge principles that purport to govern our full beliefs, disbeliefs, and suspensions of judgment, as well as the bridge principles that purport to govern our credences. We have deployed epistemic utility theory to discover which bridge principles are justified by considerations of the epistemic value that accrues to our doxastic attitudes in virtue of their accuracy. Our conclusions are a mixed bag. With very weak and natural assumptions about the epistemic utility of categorical doxastic attitudes the classical single-premise case for full belief, we found compelling arguments for the principles that most of the literature agree upon: if A entails B , then you ought not to believe A and disbelieve B , you ought not to believe A and suspend on B , and thus you ought to see to it that, if you believe A , and you adopt any attitude towards B , that attitude should be belief, providing A is not a contradiction. However, the picture is more complicated when we move to the classical multi-premise case and the non-classical single- and multi-premise cases. In these cases, the ways in which we assign epistemic utility to doxastic attitudes becomes very relevant. For instance, we obtain an epistemic utility argument for the multi-premise version of the principle that we justified in the single-premise case only if we assume that the badness of believing incorrectly is much greater than the goodness of believing correctly. And, in the non-classical case, whether or not the corresponding version of this principle holds depends on our account of the ideal doxastic attitude towards propositions with non-classical truth values.