

# Belief Gambles in Epistemic Decision Theory

Mattias Skipper

Penultimate draft, forthcoming in *Philosophical Studies*

**Abstract:** Don't form beliefs on the basis of coin flips or random guesses. More generally, don't take *belief gambles*: if a proposition is no more likely to be true than false given your total body of evidence, don't go ahead and believe that proposition. Few would deny this seemingly innocuous piece of epistemic advice. But what, exactly, is wrong with taking belief gambles? Philosophers have debated versions of this question at least since the classic dispute between William Clifford and William James near the end of the 19th century. Here I reassess the normative standing of belief gambles from the perspective of *epistemic decision theory*. The main lesson of the paper is a negative one: it turns out that we need to make some surprisingly strong and hard-to-motivate assumptions to establish a general norm against belief gambles within a decision-theoretic framework. I take this to pose a dilemma for epistemic decision theory: it forces us to either (i) make seemingly unmotivated assumptions to secure a norm against belief gambles, or (ii) concede that belief gambles can be rational after all.

*Keywords:* Belief gambles, epistemic risk, epistemic consequentialism, epistemic value, epistemic decision theory

## 1. Introduction

Mrs. Cautious is about to flip a fair coin. What will it be? Heads or tails? She's reluctant to commit to an answer. After all, she takes it, the coin is just as likely to come up heads as it is tails. So, rather than forming an arbitrary opinion on the matter, she much prefers to suspend judgment.

Mr. Hazard is pondering the universe. Is the number of stars odd or even? He can't be sure. In fact, he sees no reason to be inclined one way or the other. "But," he thinks to himself, "nothing ventured, nothing gained." So, rather than forming no opinion at all, he much prefers to form an arbitrary opinion on the matter.

When *you* find yourself flipping coins or pondering the universe, should you follow the lead of Mrs. Cautious or Mr. Hazard? The vast majority of epistemologists—at least among those who maintain that it is your evidence that determines what you should believe—would

side with Mrs. Cautious. According to them, you shouldn't take 50-50 belief gambles (henceforth just *belief gambles*): when a proposition is no more likely to be true than false given your total evidence, don't go ahead and believe that proposition. Few would deny this seemingly innocuous piece of epistemic advice.<sup>1</sup>

But what, exactly, is wrong with taking belief gambles? This question echoes the subject matter of the classic dispute between William Clifford and William James near the end of the 19th century.<sup>2</sup> Clifford famously summed up his position by deeming it “wrong always, everywhere, and for anyone to believe anything upon insufficient evidence.” It's notoriously vague what Clifford took to count as “sufficient evidence.” But he presumably had in mind an evidential threshold well above 50-50 chance.

James didn't buy it. According to him, such an exacting evidential standard has implausible consequences for how vigorously we should pursue what he conceived of as the “two great commandments” of inquiry: gaining true beliefs and avoiding false ones. James' central observation was that the pursuit of truth need not go hand in hand with the avoidance of error. In fact, the best means to the former aim—excessive credulity, the willingness to believe anything, however improbable—is also the worst means to the latter aim. And, conversely, the best means to the latter aim—excessive caution, the unwillingness to believe anything, however probable—is also the worst means to the former aim. Thus, we inevitably face a challenge of how to weigh truth-seeking against error-avoidance.

How to navigate this balancing act? Clifford's evidence norm is an effective guard against false belief, since it only permits us to form beliefs that are highly likely to be true. But it often comes at the cost of losing out on the truth, since many true propositions will fail to meet the high evidential standard. This led James to conclude that Clifford's evidence norm effectively asks us to place more weight on error-avoidance than on truth-seeking.

According to James, however, reason doesn't by itself compel us to accept any particular weighing of the fundamental goals of inquiry. As he saw it, this must be decided by our own “passional nature.” While Mrs. Cautious (and her fellow conservatives) will tend to place

---

<sup>1</sup> Here is a potential exception: some epistemologists have recently argued that there can be more than one rational response to a given body of evidence (e.g., Kelly 2013 and Schoenfield 2014). When taken to the extreme, this “permissivist” view opens up the possibility that an agent is sometimes be permitted to believe  $q$  and permitted to believe  $\sim q$  at the same time, in which case the agent is permitted to believe at least one proposition that is no more likely to be true than false (since  $q$  and  $\sim q$  can't both be more likely to be true than false). I'll return to the connection between permissivism and belief gambles in §4.3.

<sup>2</sup> Clifford (1877), James (1896).

more weight on error-avoidance than on truth-seeking, Mr. Hazard (and his fellow liberals) will tend to place more weight on truth-seeking than on error-avoidance. Neither weighing is any more rational than the other. People simply have different temperaments, and this is reflected in the way they trade off the fundamental goals of inquiry.

I suspect that few contemporary epistemologists would be prepared to follow James in placing this amount of epistemic significance on our passional lives. But the Jamesian challenge doesn't become obsolete for that reason. Far from it. Anyone who wants to maintain that belief is governed by some kind of evidence norm ultimately owes a story about why evidence gets to have such normative significance. And, as James pointed out, this may prove to be a difficult task. In particular, it's far from obvious why we should consider it more important to avoid false beliefs than to gain true ones.

So much for the historical debate between Clifford and James. How does all of this relate to the normative standing of belief gambles? We can understand the claim that belief gambles are irrational as a very minimal version of Clifford's evidence norm. Whatever else might be said about what counts as "sufficient evidence," I take it that the evidential threshold should, at the very least, lie above 50-50 chance. Of course, Clifford presumably had a much higher evidential standard in mind. But even a very minimal evidence norm—one that merely renders belief gambles irrational—will be afflicted by the Jamesian challenge. After all, someone who considers it just as important to gain true beliefs as to avoid false ones *should*, it seems, be prepared to join Mr. Hazard in taking belief gambles. Thus, anyone who wants to maintain that belief gambles are irrational faces essentially the same challenge that James leveled against Clifford's evidence norm. And given that most of us think that belief gambles are irrational, this is a challenge that most of us must face.

Hence our question: what, exactly, is wrong with taking belief gambles? Or, to put the question in a more exploratory spirit, how might we go about vindicating a general norm against belief gambles? Here I'll approach this question from the perspective of *epistemic decision theory*. The motivation to do so stems from the recent surge of interest in providing decision-theoretic arguments for various putative epistemic norms such as Probabilism (Joyce 1998), Conditionalization (Oddie 1997; Greaves & Wallace 2006), Rational Reflection (Easwaran 2013), the Principle of Indifference (Pettigrew 2016a), the Lockean Thesis (Dorst 2019), and more. To my mind, many of these norms seem much less innocuous than the requirement to not form beliefs on the the basis of coin flips or random guesses. Thus, it's natural to hope that epistemic decision theory might have the resources to vindicate a norm against belief gambles.

However, and this is the main lesson of the paper, it turns out that we need to make some surprisingly strong and hard-to-motivate assumptions to establish a general norm against belief gambles. The upshot will be that epistemic decision theory faces a dilemma that hasn't been satisfactorily addressed: it forces us to either (i) make seemingly unmotivated assumptions to secure a norm against belief gambles, or (ii) concede that belief gambles can be rational after all.

To set up this dilemma, I'll begin by introducing the decision-theoretic framework that I'll be working with, and use it to formulate a precise norm against belief gambles. This is the business of §2. In §§3-4, I'll go on to explore various proposals for how to validate this norm, and argue that each proposal rests on implausible or unmotivated assumptions. In §5, I close with some thoughts on how the main lessons of the paper bear more generally on the prospects for epistemic decision theory.

## 2. Background on Epistemic Decision Theory

It has long been commonplace to use formal, decision-theoretic methods to reason about norms of *action*. Less so when it comes to norms of *belief*. However, over the past two decades, a growing number of epistemologists have started to explore how decision-theoretic methods might be used to justify various putative epistemic norms. Continuing this line of work, I'd like to embed the requirement against taking belief gambles within a simple decision-theoretic framework.

Our first ingredient is a set,  $W$ , of *epistemically possible worlds* (henceforth just "worlds"), which represent all the different ways the world might be for all the agent under consideration can be certain of. For example, if the agent is uncertain about whether the number of stars in the universe is odd or even, there will be worlds in  $W$  representing the number of stars as odd, and there will be worlds in  $W$  representing the number of stars as even. By contrast, if the agent is certain that the number of stars is greater than the number of galaxies, every world in  $W$  will represent the number of stars as being greater than the number of galaxies.<sup>3</sup>

Propositions are modeled as subsets of  $W$ : that is, we think of a proposition,  $q$ , as the set of worlds in  $W$  that represent  $q$  as being true. Truth is represented in terms of set-membership: that is,  $q$  is true at a world  $w$  iff  $w \in q$ . Logical operations are defined in terms

---

<sup>3</sup> Throughout, I'll assume that the number of worlds in  $W$  is finite. This is not to say that agents can only ever represent finitely many possibilities. However, the infinite case introduces some technical complications that I'd like to sidestep for present purposes.

of set-theoretic ones:  $\sim q = W - q$ ,  $p \wedge q = p \cap q$ , and so on. The agent's evidential situation is represented by a probability function,  $P$ , which assigns real numbers between 0 and 1 to propositions. For each proposition,  $q$ ,  $P(q)$  is the *evidential probability* of  $q$ : that is, the probability of  $q$  given the agent's total body of evidence.

While decision theorists typically define a "menu" of options for agents to choose from, we will define a set,  $A$ , of *doxastic attitudes* (henceforth just "attitudes") that agents can adopt towards any given proposition. If  $a$  is an attitude and  $q$  is a proposition, I'll write " $a(q)$ " to mean that the agent takes attitude  $a$  towards  $q$ . When context allows, I'll omit reference to the proposition and simply write " $a$ " to mean that the agent takes attitude  $a$  towards the contextually given proposition.

I'll make three assumptions about the attitudes in  $A$ . First, I'll be working with a relatively coarse-grained distinction between *belief* ("b"), *suspension of judgment* ("s"), and *disbelief* ("d"). This is not to say that the Jamesian challenge presupposes any particular individuation of doxastic attitudes. Indeed, I suspect that a very similar challenge will arise in a credence-based setting.<sup>4</sup> But to keep things relatively simple, I'll focus on the belief case here. Second, I'll assume that there is no difference between disbelieving  $q$  and believing  $\sim q$ . Finally, I'll assume that suspension of judgment is incompatible with belief (and disbelief): that is, if you believe (or disbelieve)  $q$ , you thereby don't suspend judgment about  $q$ . Otherwise, I won't make any substantive assumptions about what combinations of attitudes it is possible for agents to have. For example, I won't foreclose the possibility of believing a proposition and its negation at the same time.

---

<sup>4</sup> For those interested, here is a brief excursion into accuracy-first epistemology: it might be tempting to think that we automatically get a norm against "credence gambles" as long as we use a *proper scoring rule*, since proper scoring rules have the property of being "immodest" or "self-recommending:" roughly, they expect themselves to be most accurate. For example, an agent with a credence of 50% in  $q$  won't expect any other credence to be more accurate, and so won't be prepared to "gamble" by moving to a credence of 80% in  $q$  (assuming that the agent seeks to maximize expected accuracy calculated with respect to his or her own credence function).

However, the self-recommendation property of proper scoring rules doesn't by itself get us a norm against credence gambles, since it says nothing about how the rational credence relates to the available evidence. To get a norm against credence gambles, one would need to add a further assumption to the effect that the rational credence in  $q$  is identical to the evidential probability of  $q$ . But in doing so, one effectively assumes that it is more important to avoid high credences in falsehoods than it is to gain high credences in truths. Hence, we are left with essentially the same Jamesian challenge that I've set out to address in this paper.

Further discussion of how the accuracy-based approach to epistemology relates to the debate between Clifford and James can be found in Joyce (2009, §9) and Pettigrew (2016b). See also Williamson (2000, ch. 10) for discussion of the relation between evidential probability and rational credence.

Importantly, epistemic decision theory—or, in any case, the version of epistemic decision theory that I’ll be focusing on here—is based on a *consequentialist* view of rational belief. To a first approximation, epistemic consequentialism is the view that the rational status of any given doxastic attitude is determined solely by its conduciveness (or lack thereof) to that which has fundamental *epistemic value*, whatever it is. In principle, this view of epistemic rationality is compatible with many different theories about what has epistemic value. However, the predominant value theory among epistemic decision theorists is one according to which *true belief* is the sole bearer of fundamental epistemic value, whereas *false belief* is the sole bearer of fundamental epistemic disvalue. Following Goldman (1999), I’ll refer to this version of epistemic consequentialism as “Veritism.”

To implement these consequentialist commitments, let’s define a *value function*,  $V$ , which assigns real numbers to attitudes at worlds. If  $a$  is an attitude (towards a contextually given proposition), and  $w$  is a world, then  $V(a, w)$  is the epistemic value of having  $a$  in  $w$ . At this point, I won’t make any other assumptions about  $V$ , but later on I’ll consider various additional constraints that one might want to impose on  $V$ .

Finally, let  $R$  be a *decision rule* that assigns sets of attitudes to propositions. If  $q$  is a proposition,  $R(q)$  is the set of attitudes that the agent in question is *rationally permitted* (or just “permitted”) to adopt towards  $q$ . So, for example, if  $b \in R(q)$ , the agent is permitted to believe  $q$ ; and, if  $s \notin R(q)$ , the agent isn’t permitted to suspend judgment about  $q$ .

With these ingredients in hand, we can give a precise formulation of the claim that belief gambles are irrational:

### **Belief Gamble Aversion**

You’re permitted to believe a proposition only if the evidential probability of that proposition is greater than 50%:

$$[b \in R(q)] \rightarrow [P(q) > .5]$$

Note that Belief Gamble Aversion provides only a *necessary* condition for permissible belief: it says that you aren’t permitted to believe a proposition if the evidential probability of that proposition is no greater than 50%, but it says nothing about whether you are permitted to believe a proposition if the evidential probability of that proposition is greater than 50%.

As we will see, Belief Gamble Aversion isn’t automatically rendered valid or invalid by the decision-theoretic framework outlined above. This is important because it allows us to explore various ways of validating Belief Gamble Aversion by imposing additional constraints on our framework. The goal of the next two sections is to do just that.

### 3. Dominance and Minimal Veritism

Plausibly, you shouldn't choose an option if there is an alternative option that is better *no matter what*. In decision-theoretic parlance: if option A is *dominated* by option B, you shouldn't choose option A. Doing so would effectively amount to ignoring an option that is guaranteed to make you better off—and this seems patently irrational.<sup>5</sup>

When adapted to the present framework, the claim that dominated options are always impermissible becomes:

#### **Dominance**<sup>6</sup>

You're permitted to have an attitude,  $a$ —that is,  $a \in R(q)$ —only if no other attitude,  $a'$ , *dominates*  $a$ , where  $a'$  dominates  $a$  iff:

- (i)  $V(a',w) \geq V(a,w)$ , for all worlds  $w$ ; and
- (ii)  $V(a',w) > V(a,w)$ , for some world  $w$ .

As stated, Dominance states merely a necessary condition on permissible attitudes: it says that dominated attitudes aren't permitted, but it says nothing about whether undominated attitudes are permitted. In §4.1, I'll consider alternative decision rules that include a sufficient condition as well. But I'd like to begin by considering whether a minimal decision rule like Dominance might be enough to vindicate Belief Gamble Aversion.

To go along with Dominance, I'll impose a few different constraints on the value function, which I take to capture a very minimal commitment to Veritism:<sup>7</sup>

#### **Minimal Veritism**

*Ranking:* If  $w \in q$ , then  $V(b,w) > V(s,w) > V(d,w)$ .

*Extensionality:* If  $\{w,w'\} \subseteq q$  or  $\{w,w'\} \subseteq \sim q$ , then  $V(a,w) = V(a,w')$ .

*Equivalence:* If  $w \in q$  and  $w' \notin q$ :

- (i)  $V(b,w) = V(d,w')$ ; and
- (ii)  $V(d,w) = V(b,w')$ ; and
- (iii)  $V(s,w) = V(s,w')$ .

---

<sup>5</sup> Complications arise in cases where every available option is dominated by some other option. But for present purposes, we can safely ignore this subtlety. See Pettigrew (2016c) and Briggs (2018) for discussion of this point.

<sup>6</sup> Decision theorists often distinguish between *weak* dominance (as defined here) and *strong* dominance (defined as strict inequality of value at all worlds). For dialectical reasons, I'll focus on weak dominance here, since this makes for a stronger decision rule (that is, one that renders fewer attitudes permitted).

<sup>7</sup> Similar constraints have been employed by, e.g., Dorst (2019) and Pettigrew (2016b).

Here is a quick run-through of these assumptions.

The first assumption, Ranking, says two things: (i) true belief is more valuable than suspension of judgment; and (ii) false belief is less valuable than suspension of judgment. These are meant as purely ordinal claims. Nothing is said about *how much* better (or worse) true (or false) belief is than suspension of judgment. It's not even assumed that the *difference* in value between true belief and suspension of judgment stands in some particular ordinal relation to the *difference* in value between false belief and suspension of judgment.

The second assumption, Extensionality, says that the value or disvalue of a given attitude depends solely on the *truth-value* of the relevant proposition. For example, the value or disvalue of believing  $q$  depends solely on whether  $q$  is true or false. This is what makes Veritism a form of monism about epistemic value: only true belief has fundamental epistemic value, and only false belief has fundamental epistemic disvalue.

The third assumption, Equivalence, consists of three indifference claims: (i) there is no difference in value between believing a true proposition and disbelieving a false one; (ii) there is no difference in value between disbelieving a true proposition and believing a false one; and (iii) there is no difference in value between suspending judgment about a true proposition and suspending judgment about a false one.

These three assumptions—Ranking, Extensionality and Equivalence—are ones that presumably all advocates of Veritism should be willing to accept. Later, in §4.2, I'll consider some stronger and more controversial constraints that one might want to impose on the value function. But first I want to consider whether Minimal Veritism, when combined with Dominance, is sufficient to vindicate Belief Gamble Aversion.

The answer turns out to be negative: even if the evidential probability of  $q$  is 50% or lower, it is consistent with Minimal Veritism that believing  $q$  isn't dominated by suspending judgment about  $q$  (modulo the trivial case where the evidential probability of  $q$  is 0%). This is our first result (all proofs can be found in the Appendix):

### **First Negative Result**

Dominance + Minimal Veritism  $\not\Rightarrow$  Belief Gamble Aversion.

Given how weak Dominance and Minimal Veritism both are, this result may not seem all that surprising. But it serves to illustrate a challenge that will recur below: it turns out to be much harder to give a decision-theoretic vindication of Belief Gamble Aversion than one might have initially expected. In particular, Dominance and Minimal Veritism aren't strong



enough. Hence, we are left to explore how these minimal assumptions might be strengthened so as to validate Belief Gamble Aversion.

#### 4. Vindicating Belief Gamble Aversion: Three Strategies

In this section, I'll explore three natural strategies for giving a decision-theoretic vindication of Belief Gamble Aversion. I'll say more to clarify and motivate the strategies as we go, but here they are in rough outline:

- (1) Replace Dominance by a stronger decision rule.
- (2) Replace Minimal Veritism by a stronger set of constraints on the value function.
- (3) Appeal to "exogenous" principles that might be thought to have a bearing on the normative standing of belief gambles.

Below I consider each strategy in turn, establish various results about them, and use these results to draw out some general lessons about the prospects for a decision-theoretic vindication of Belief Gamble Aversion.

##### 4.1. First Strategy: Risk-Averse Decision Rules

As illustrated by our initial comparison of Mrs. Cautious and Mr. Hazard, we can understand Belief Gamble Aversion as a *risk-averse* norm—one that tells us to avoid error more vigorously than we pursue the truth. No wonder, then, that Dominance is of little help when it comes to validating Belief Gamble Aversion. After all, Dominance isn't a risk-averse decision rule. It isn't risk-seeking either. Dominance is simply silent on how risk-averse or risk-seeking you should be. For example, if you're offered a gamble that gives you a 1% chance of winning \$1 and a 99% chance of losing \$1M, it would be extremely risk-seeking to take the gamble. Yet, for all Dominance says, doing so is permitted. Conversely, if you're offered a gamble that gives you a 1% chance of losing \$1 and a 99% chance of winning \$1M, it would be extremely risk-averse to refuse to take the gamble. Yet, for all Dominance says, doing so is permitted. This makes it natural to think that we might be able to vindicate Belief Gamble Aversion by replacing Dominance with a risk-averse decision rule. But which one?

A familiar candidate is the *Maximin* rule, which plays a central role in John Rawls' (1971) theory of justice, and which has more recently been used by Richard Pettigrew (2016a) to give a decision-theoretic vindication of the Principle of Indifference. According to Maximin, you're permitted to choose an option only if its worst outcome is at least as good as the worst outcome of any alternative option. For example, if you're offered a gamble that

gives you a 99% chance of winning \$1M and a 1% risk of losing \$1, you aren't permitted by Maximin to take the gamble, since the worst outcome of taking the gamble (losing \$1) is worse than the worst outcome of not taking the gamble (losing nothing and winning nothing).

When adapted to our framework, the Maximin rule becomes:

### **Maximin**

You're permitted to have an attitude,  $a$ , only if no other attitude,  $a'$ , *worst-case dominates*  $a$ , where  $a'$  worst-case dominates  $a$  iff  $\min_{w:P(\{w\})>0} [V(a',w)] > \min_{w:P(\{w\})>0} [V(a,w)]$ .

Like Dominance, Maximin states merely a necessary condition on permissible attitudes: it says that worst-case dominated attitudes aren't permitted, but it says nothing about whether attitudes that aren't worst-case dominated are permitted. Maximin is sometimes formulated with a sufficient condition as well, but we only need the necessary condition here, since it turns out to be enough to validate Belief Gamble Aversion, when combined with Minimal Veritism. More precisely: whenever the evidential probability of  $q$  is 50% or lower, it follows from Minimal Veritism that believing  $q$  will be worst-case dominated by suspension of judgment about  $q$ . That's our first positive result:

### **First Positive Result**

Maximin + Minimal Veritism  $\Rightarrow$  Belief Gamble Aversion.

This result shows—for what it's worth—that it is indeed possible to validate Belief Gamble Aversion within our decision-theoretic framework. But, obviously, not any old *derivation* of Belief Gamble Aversion can serve as a satisfying *vindication* of the norm. At the very least, the assumptions that go into the derivation shouldn't be too implausible or unmotivated. Unfortunately, the First Positive Result falls short in this regard, I think.

The problem is that Maximin is an *extremely* risk-averse decision rule. It effectively tells us to only care about error-avoidance and forget all about truth-seeking. Consequently, someone who conforms to Maximin will only be prepared to believe propositions the truth of which she can be absolutely certain.<sup>8</sup> And since there is almost always, if not always, at least a miniscule risk of being wrong, this means that someone who conforms to Maximin will hardly ever, if ever, form any beliefs at all. Presumably, rationality doesn't demand such an

---

<sup>8</sup> More precisely: it follows from Maximin and Minimal Veritism that believing  $q$  is permitted only if  $P(q) = 1$ . The proof of this result is analogous to that of the First Positive Result, and I omit the details.

extreme level of epistemic risk-aversion. It might well be advisable to conform to Maximin in *some* contexts—say, when choosing one’s urpriors (cf. Pettigrew 2016a)—but it seems highly implausible to suppose that we should *always* conform to Maximin when making epistemic decisions. The same goes in the practical case: it might well be advisable to conform to Maximin in *some* practical contexts—say, when negotiating behind the veil of ignorance (cf. Rawls 1971)—but it seems highly implausible to suppose that we should *always* conform to Maximin when making practical decisions.

Instead, one might try out a decision rule, which is stronger than Dominance, but which doesn’t share its problematic risk-averse features. The most salient candidate here is the standard rule of *Expected Utility Maximization*, which lies at the heart of orthodox decision theory. According to this decision rule, you should always choose an option whose expected value is at least as great as the expected value of any other option. For example, if you’re offered a gamble that gives you a 50% chance of winning \$2 and a 50% risk of losing \$1, you should take the gamble, since the expected value of taking the gamble is positive (more specifically, \$1). By contrast, if you’re offered a gamble that gives you a 50% chance of winning \$1 and a 50% risk of losing \$2, you shouldn’t take the gamble, since the expected value of doing so is negative (more specifically, minus \$1).

When adapted to our framework, the rule of Expected Utility Maximization becomes:

#### **Risk-Neutral EUM**

You’re permitted to have an attitude,  $a$ , iff no other attitude,  $a'$ , has greater expected epistemic value than  $a$ , where the expected value,  $E_P(a)$ , of  $a$  relative to a probability function,  $P$ , is defined as:  $E_P(a) = \sum_w V(a,w)P(w)$ .

Unlike Dominance and Maximin, Risk-Neutral EUM includes both a necessary and a sufficient condition on permissible attitudes. Even so, however, Risk-Neutral EUM isn’t strong enough to validate Belief Gamble Aversion, when combined with Minimal Veritism. More precisely: even if the evidential probability of  $q$  is 50% or lower, it’s consistent with Minimal Veritism that the expected epistemic value of believing  $q$  is no smaller than the expected epistemic value of suspending judgment about  $q$  (modulo the trivial case where the evidential probability of  $q$  is 0%). This is our next result:

## Second Negative Result<sup>9</sup>

Risk-Neutral EUM + Minimal Veritism  $\not\Rightarrow$  Belief Gamble Aversion.

I take this result to suggest that if we want to validate Belief Gamble Aversion without imposing any additional constraints on the value function, we need to adopt a decision rule with at least *some* risk-averse features. An entirely risk-neutral decision rule like Risk-Neutral EUM won't do it.

The previous two results—that is, the First Positive Result and the Second Negative Result—might be taken to suggest that we face a problem of *moderation*. Extremely risk-averse decision rules like Maximin are overly prohibitive. Completely risk-neutral decision rules like Risk-Neutral EUM aren't prohibitive enough. Is there a way to strike a middle ground? A natural idea would be to draw on Lara Buchak's (2013) recent work on "risk-weighted" utility theory. Very roughly, she characterizes a spectrum of decision rules, ranging from extremely risk-averse rules (akin to Maximin) to extremely risk-seeking ones (akin to Maximax). Perhaps we should be looking for a decision rule within this spectrum, which is neither too close to the risk-neutral middle, nor too close to the risk-averse extreme.

However, I don't find this strategy very promising—for at least two reasons. First, no moderately risk-averse decision rule is going to be strong enough to validate Belief Gamble Aversion, unless we impose additional constraints on the value function (that is, ones that go beyond Minimal Veritism).<sup>10</sup> And, as we will see in §4.2, doing so quickly gets us into contentious territory. Secondly, and perhaps more importantly, any moderately risk-averse decision rule faces essentially the same problem that led us to reject Maximin in the first place. Even if it is *sometimes* advisable to conform to a risk-averse decision rule, it's hard to see why one should *always* do so. Why, for example, think that someone who conforms to Risk-Neutral EUM on a given occasion is necessarily irrational for that reason? I don't see any independent reasons to think so. Of course, there might be reasons to doubt that one should *always* conform to Risk-Neutral EUM. But to deny that it can *ever* be rational to conform to Risk-Neutral EUM seems like a desperate move.

---

<sup>9</sup> Since Risk-Neutral EUM is stronger than Dominance, this result can be seen as a corollary of the Second Negative Result.

<sup>10</sup> The formal details behind this claim would take us too far astray, but the informal story goes as follows: Minimal Veritism says nothing about how the value of true belief compares to the disvalue of false belief. Hence, we can always make the risk-weighted expected utility of believing  $q$  greater than the risk-weighted expected utility of suspending judgment about  $q$ , simply by making the value of true belief high enough compared to the disvalue of false belief. To eliminate this possibility, we would have to go beyond Minimal Veritism.

In sum, then, the first strategy doesn't seem promising to me. We'd better look elsewhere for a decision-theoretic vindication of Belief Gamble Aversion.

#### 4.2. Second Strategy: Risk-Averse Value Functions

Rather than adopting a risk-averse *decision rule*, perhaps we should adopt a risk-averse *value function* instead. Here is the weakest strengthening of Minimal Veritism that meets this objective:<sup>11</sup>

##### **Risk-Averse Veritism**

In addition to Minimal Veritism, the value function,  $V$ , must satisfy the following constraint, for all  $w \in q$ :

$$|V(b,w) - V(s,w)| < |V(d,w) - V(s,w)|$$

According to Risk-Averse Veritism, the absolute difference in value between true belief and suspension of judgment is smaller than the absolute difference in value between false belief and suspension of judgment. On the assumption that suspension of judgment has no value or disvalue (i.e.,  $V(s,w) = 0$ ), this amounts to saying, more simply, that the value of true belief is smaller than the disvalue of false belief (i.e.,  $|V(b,w)| < |V(d,w)|$ ). Like Ranking, this is meant as a purely ordinal claim. Nothing is said about *how much* smaller the value of true belief is than the disvalue of false belief. All Risk-Averse Veritism says is that true belief isn't as valuable as false belief is disvaluable.

If we combine Risk-Averse Veritism with Risk-Neutral EUM, Belief Gamble Aversion follows. That is, whenever the evidential probability of  $q$  is 50% or lower, it follows from Risk-Averse Veritism that the expected epistemic value of believing  $q$  is smaller than the expected epistemic value of suspending judgment about  $q$ . That's our next result:

##### **Second Positive Result**

Risk-Neutral EUM + Risk-Averse Veritism  $\Rightarrow$  Belief Gamble Aversion.

---

<sup>11</sup> Similar constraints have been employed by Dorst (2017), Easwaran (2016), and Pettigrew (2016b) to capture the idea that one should place more weight on error-avoidance than on truth-seeking.

The upshot of this result is that *any* risk-averse value function is strong enough to establish a prohibition against belief gambles, when combined with Risk-Neutral EUM.<sup>12</sup>

However, it seems to me that Risk-Averse Veritism raises more questions than it answers. What reason is there to think that the intrinsic disvalue of false belief is greater than the intrinsic value of true belief? Why not instead say that the intrinsic value of true belief is equal to or greater than the intrinsic disvalue of false belief? This question cries out for an answer. After all, it's far from obvious that there *should* be an asymmetry between the value of true belief and the disvalue of false belief. On the contrary, positing such a value asymmetry strikes me with an air of puzzlement or mystery. What might the Veritist say to alleviate this reaction?

I anticipate a temptation here to fall back on practical considerations to justify Risk-Averse Veritism. Perhaps false beliefs tend to be more harmful than true beliefs are beneficial, from a practical point of view. The idea may sound plausible. But it isn't available to the veritist. Whatever we say on behalf of the veritist to motivate Risk-Averse Veritism must be consistent with the view that truth and falsity are the sole bearers of fundamental epistemic value and disvalue, and that these values alone are what grounds epistemic norms. It might well be that the practical disvalue of false belief is (typically) greater than the practical value of true beliefs. But from the veritist's perspective, we should resist the temptation to motivate Risk-Averse Veritism on practical grounds—to do so would be to give up the game.

Relatedly, it might be tempting to invoke the Jamesian idea that the asymmetry between truth-seeking and error-avoidance is simply a feature of our human psychology, which isn't subject to normative evaluation. Just as subjective Bayesians treat an agent's "urpriors" as an arational matter, the idea goes, we should treat an agent's propensity for epistemic risk aversion as an arational matter. This obviously raises the question of whether people in fact exhibit a sufficiently general propensity for epistemic risk aversion. But whatever the answer to this question is, this proposal isn't available to the Veritist either, since it once again grounds the value or disvalue of a given belief in something other than the truth or falsity of that belief—this time in certain psychological features of the believer.

An initially more promising suggestion, due to Kevin Dorst (2019, pp. 185), goes as follows. Consider two agents, Ann and Bob, whose doxastic states differ only in their attitude

---

<sup>12</sup> What if we combine Risk-Averse Veritism with Dominance instead of Risk-Neutral EUM? The answer is that Belief Gamble Aversion no longer follows, since Ranking implies that believing  $q$  won't be dominated by suspending judgment about  $q$  (modulo the trivial case where  $P(q) = 0$ ). The formal details behind this result are left to the reader.

towards  $q$ : Ann suspends judgment about  $q$ , whereas Bob believes both  $q$  and  $\sim q$ . Suppose also that Ann and Bob possess the exact same evidence and that  $q$  and  $\sim q$  are equally likely on that evidence. We can then ask: which belief state would a rational agent prefer to be in? Surely, Ann's! But if Risk-Averse Veritism were false, it's hard to see why this would be so. After all, if Risk-Averse Veritism were false, Bob's belief state would instantiate just as much (if not more) epistemic value as Ann's. So, Risk-Averse Veritism must be true—or so the suggestion goes.

I don't find this line of reasoning quite convincing, for reasons also put forth by Florian Steinberger (2019, §3). The problem, as I see it, is that it gets the explanatory direction wrong: it uses a premise about what is intuitively *rational* to establish a conclusion about what is *epistemically valuable*, whereas we're looking for a premise about what is intuitively *epistemically valuable* to establish a conclusion about what is *rational*. This is obviously not to deny that people, when confronted with the case above, have a strong intuition that Ann is somehow "better off" than Bob, epistemically speaking. But what drives this intuition (at least insofar as I can tell) is that Bob is being irrational, whereas Ann isn't (or, at any rate, need not be). And this is very different from saying that Ann's belief state instantiates more epistemic value than Bob's.

At this point, it might be wondered whether I'm being too hard on the Veritist. Am I asking for an explanation where none can be found? This would be the case if Risk-Averse Veritism encoded a "brute fact" about epistemic value, one that cannot be grounded in anything more fundamental. At the end of the day, explanation must stop *somewhere*. Why not here?

My reluctance to treat Risk-Averse Veritism as a fundamental "bedrock" assumption stems from what I take to be a fairly uncontroversial methodological commitment, whereby the fundamental assumptions on which we base our epistemological theories should, to the extent possible, be relatively plausible on their face. At the very least, they shouldn't strike us as outright puzzling or mysterious. But Risk-Averse Veritism does just that. I don't consider this to be a *decisive* reason not to treat Risk-Averse Veritism as a fundamental assumption. Perhaps other considerations will tip the scale. But it seems to me that Risk-Averse Veritism isn't the kind of principle that we should easily accept as part of our theoretical bedrock.

Before I proceed to the third and final strategy, let me consider, for the sake of completeness, the following *risk-neutral* strengthening Minimal Veritism:

### Risk-Neutral Veritism

In addition to Minimal Veritism, the value function,  $V$ , must satisfy the following constraint, for all  $w \in q$ :

$$|V(b, w) - V(s, w)| = |V(d, w) - V(s, w)|$$

According to Risk-Neutral Veritism, the value of true belief equals the disvalue of false belief (assuming that suspending judgment has no value or disvalue). On the face of it, this assumption might seem less *ad hoc* than its risk-averse cousin. However, as will be unsurprising by now, Risk-Neutral Veritism isn't strong enough to validate Belief Gamble Aversion, when combined with a risk-neutral decision rule like Risk-Neutral EUM. That's our third negative result:

### Third Negative Result<sup>13</sup>

Risk-Neutral EUM + Risk-Neutral Veritism  $\not\Rightarrow$  Belief Gamble Aversion.

In sum, then, the second strategy doesn't seem promising to me either.

### 4.3. Third Strategy: "Exogenous" Constraints

So far, I've examined various *endogenous* constraints on the decision-theoretic framework: that is, constraints that apply to various parts of the framework itself (specifically: the decision rule and the value function). But one might also try to constrain the framework in an *exogenous* way, by appealing to various independently plausible principles that aren't integral to the decision-theoretic framework itself. Below I examine two such principles, which might be thought to have a bearing on the normative standing of belief gambles.

The first principle is a familiar logical norm:

#### Non-Contradiction

You're never permitted to believe a proposition and its negation both at once:

$$b \notin R(q \wedge \sim q)$$

---

<sup>13</sup> This result has the Second Negative Result as a corollary, since Risk-Neutral Veritism is stronger than Minimal Veritism. Moreover, since Dominance is weaker than Risk-Neutral EUM, it is a corollary of the Third Negative Result that Dominance and Risk-Neutral Veritism do not jointly entail Belief Gamble Aversion.



Why think that Non-Contradiction bears on the normative standing of belief gambles? A tempting thought might be that Non-Contradiction straightforwardly *entails* Belief Gamble Aversion, because denying Belief Gamble Aversion amounts to saying that an agent can be at once permitted to believe  $q$  and permitted to believe  $\sim q$ . However, this would be too quick: even if an agent is both permitted to believe  $q$  and permitted to believe  $\sim q$ , the agent need not be permitted to believe both  $q$  and  $\sim q$  at the same time. This is sometimes put by saying that permissions don't (necessarily) *agglomerate*. That's what drives the following result:

**Fourth Negative Result**<sup>14,15</sup>

Risk-Neutral EUM + Risk-Neutral Veritism + Non-Contradiction  $\not\Rightarrow$  Belief Gamble Aversion.

I take this result to suggest that Non-Contradiction is not going to provide us with a straightforward vindication of Belief Gamble Aversion. Still, one might wonder whether Non-Contradiction can provide us with some kind of *indirect* vindication of Belief Gamble Aversion. This would arguably be the case if Belief Gamble Aversion turned out to be necessary in order to vindicate Non-Contradiction. However, it doesn't turn out that way. We can derive Non-Contradiction from Dominance and Minimal Veritism alone:

**Third Positive Result**<sup>16</sup>

Dominance + Minimal Veritism  $\Rightarrow$  Non-Contradiction.

---

<sup>14</sup> This result has the Third Negative Result as a corollary, since Non-Contradiction is added as an additional premise. Moreover, since Dominance and Minimal Veritism are weaker than Risk-Neutral EUM and Risk-Neutral Veritism, respectively, there are two immediate corollaries of the Fourth Negative Result: first, Dominance and Risk-Neutral Veritism do not jointly entail Belief Gamble Aversion; and, second, Risk-Neutral EUM and Minimal Veritism do not jointly entail Belief Gamble Aversion.

<sup>15</sup> Easwaran and Fitelson (2015, Theorem 6) present a similar result, which essentially shows that a norm against belief gambles can be used to derive a dominance principle and a non-contradiction principle, whereas the opposite entailment doesn't hold. A detailed comparison to their discussion is beyond the scope of this paper, but it's worth noting that their framework (as they define it) doesn't allow for suspension of judgment: that is, agents must either believe or disbelieve any given proposition.

<sup>16</sup> Here is a subtlety: although Dominance and Minimal Veritism entail a norm against believing the *conjunction* of  $q$  and  $\sim q$ , they don't straightforwardly entail a norm against believing  $q$  and  $\sim q$  *separately*. This is because believing  $q \wedge \sim q$  amounts to having *one* false belief, whereas believing  $q$  and believing  $\sim q$  amounts to having *one* false belief and *one* true belief. Thus, given only Minimal Veritism, the former option is dominated by suspension of judgment about  $q$ , whereas the latter isn't. However, we can get a norm against believing  $q$  and  $\sim q$  separately by invoking a (widely held, albeit not uncontroversial) closure principle, whereby you're

It is also worth noting that, even if Belief Gamble Aversion *did* turn out to feature in our best vindication of Non-Contradiction, this would not (it seems to me) do much to vindicate Belief Gamble Aversion itself. It might give us additional reason to think that Belief Gamble Aversion is *true*. But we have already granted this much. What we are after is a way to ground Belief Gamble Aversion in other, more basic principles that we find independently plausible.

Moving on to the second exogenous principle, consider another familiar norm:

**Impermissivism**<sup>17</sup>

You're never permitted to believe a proposition and permitted to disbelieve that proposition at the same time:

$$\{b, d\} \not\subseteq R(q)$$

Obviously, Impermissivism is contested—much more so than Non-Contradiction—and I won't take a stance on it here.<sup>18</sup> But I find it worth observing that there is a close relationship between Impermissivism and Belief Gamble Aversion:

**Fourth Positive Result**

Risk-Neutral EUM + Risk-Neutral Veritism + Impermissivism  $\Rightarrow$  Belief Gamble Aversion.

What this result illustrates is that we can derive Belief Gamble Aversion from Impermissivism without making any risk-averse assumptions about our decision rule or value function.<sup>19</sup>

---

permitted to believe  $q$  and  $\sim q$  separately only if you're permitted to believe  $q \wedge \sim q$ . So, in any case, Belief Gamble Aversion is hardly needed to vindicate the idea that contradictory beliefs are irrational.

<sup>17</sup> Some epistemologists have discussed a stronger impermissivist thesis, which says that there is always a *unique* permissible attitude to have towards any given proposition (White 2005; Kelly 2013). The present formulation doesn't imply this much, since it leaves open the possibility that one can be at once permitted to suspend judgment about  $q$  and permitted to believe (or disbelieve)  $q$ . I focus on this weaker version of the thesis, since it turns out to be strong enough to establish the Fourth Positive Result.

<sup>18</sup> For an overview of the recent epistemological debate on permissivism vs. uniqueness, see Kopec and Titelbaum (2016).

<sup>19</sup> What happens if we replace Risk-Neutral EUM by Dominance? Answer: Belief Gamble Aversion no longer follows, since Ranking ensures that believing  $q$  isn't dominated by disbelieving  $q$  (modulo the trivial case where  $P(q) = 0$ ). What happens if we replace Risk-Neutral Veritism by Minimal Veritism? Answer: Belief Gamble Aversion no longer follows, since Minimal Veritism says nothing about how the value of true belief compares to the disvalue of false belief. The details of both results are left out.

Despite this, however, I'm doubtful that Impermissivism can provide us with a satisfying vindication of Belief Gamble Aversion. For one thing, there is something methodologically suspect about invoking a highly controversial norm like Impermissivism to ground a relatively innocuous norm like Belief Gamble Aversion. For another thing, it turns out that we can straightforwardly derive Impermissivism from Belief Gamble Aversion, with no additional assumptions:

#### **Fifth Positive Result**

Belief Gamble Aversion  $\Rightarrow$  Impermissivism.

I take this result to provide yet another reason not to ground Belief Gamble Aversion in Impermissivism. If anything, it seems more natural to treat Belief Gamble Aversion as the more basic norm of the two.

There are no doubt other exogenous principles aside from Non-Contradiction and Impermissivism that one might try to bring to bear on Belief Gamble Aversion. But in light of the above, I see little reason for optimism on behalf of the third strategy.

## **5. Conclusion**

Belief gambles are irrational. Few would deny it. But what makes it so? I've tried to address this question from the perspective of epistemic decision theory. Needless to say, I haven't exhausted all imaginable strategies for providing a decision-theoretic vindication of Belief Gamble Aversion. But I have examined three natural candidates, with unpromising results: each strategy turned out to rest on implausible or unmotivated assumptions. That's why I said in the beginning that epistemic decision theory faces a dilemma: it forces us to either (i) make unmotivated assumptions to secure a norm against belief gambles, or (ii) concede that belief gambles can be rational after all.

Where to go from here? I suspect that few epistemologists—especially among those who think that it is your evidence that determines what you should believe—would be prepared to concede that belief gambles can be rational. Those who share this sentiment may take the upshot of this paper as an invitation to search for alternative ways of vindicating the idea that belief gambles are irrational. Perhaps deontological theories of rational belief are better suited for this purpose (Sylvan 2020). Or perhaps a knowledge-based account is the way to go (Williamson 2000). I won't try to assess such alternatives here—they all deserve a much fuller discussion than I can hope to provide in these closing remarks. But whichever way we go, it

seems to me that we need to face up to the dilemma presented in this paper. Until we do, a notable gap will remain in our understanding of rational belief.

## Appendix

This appendix contains proofs of the results in the main text. The formal framework is introduced in §2. Further definitions are introduced in §§3-4. Results are repeated here for convenience. To avoid clutter, I'll write " $P(w)$ " to denote the evidential probability of  $\{w\}$ .

**First Negative Result:** Dominance + Minimal Veritism  $\not\Rightarrow$  Belief Gamble Aversion.

*Proof.* We must define a model that satisfies Dominance and Minimal Veritism, but not Belief Gamble Aversion. Let  $W = \{w_1, w_2\}$ ,  $q = \{w_1\}$ , and  $P(q) = .5$ . The following specification of  $V$  is consistent with Minimal Veritism:

$$\begin{aligned} V(b, w_1) &= V(d, w_2) = 1 \\ V(s, w_1) &= V(s, w_2) = 0 \\ V(d, w_1) &= V(b, w_2) = -1 \end{aligned}$$

Since no attitude dominates any other attitude relative to this value assignment, the following specification of  $R$  is consistent with Dominance:

$$R(q) = \{b, s, d\}$$

Thus,  $b \in R(q)$  and  $P(q) \not\geq .5$ , contrary to Belief Gamble Aversion. □

**First Positive Result:** Maximin + Minimal Veritism  $\Rightarrow$  Belief Gamble Aversion.

*Proof.* We must show that any model that satisfies Maximin and Minimal Veritism must satisfy Belief Gamble Aversion as well. Begin by assuming that  $P(q) \leq .5$ . Given this, there must be at least one world  $w$  such that  $q$  is false at  $w$ . By Ranking, it follows that  $V(b, w) < V(s, w')$ , for all  $w'$  (where  $w'$  may or may not be identical to  $w$ ). But this is just to say that  $s$  worst-case dominates  $b$ . Thus, by Maximin,  $b \notin R(q)$ , as Belief Gamble Aversion dictates. □

**Second Negative Result:** Risk-Neutral EUM + Minimal Veritism  $\not\Rightarrow$  Belief Gamble Aversion.

*Proof.* We must define a model that satisfies Minimal Veritism and Risk-Neutral Expected Value Maximization, but not Belief Gamble Aversion. Let  $W = \{w_1, w_2\}$ ,  $q = \{w_1\}$ , and  $P(q) = .5$ . The following specification of  $V$  is consistent with Minimal Veritism:

$$\begin{aligned}
V(b, w_1) &= V(d, w_2) = 1 \\
V(s, w_1) &= V(s, w_2) = 0 \\
V(d, w_1) &= V(b, w_2) = -1
\end{aligned}$$

Given this specification of  $P$  and  $V$ , it follows trivially that  $E_P(b) = E_P(s) = E_P(d) = 0$ . Hence, the following specification of  $R$  is consistent with Risk-Neutral EUM:

$$R(q) = \{b, s, d\}$$

Thus,  $b \in R(q)$  and  $P(q) \not\geq .5$ , contrary to Belief Gamble Aversion.  $\square$

**Second Positive Result:** Risk-Neutral EUM + Risk-Averse Veritism  $\Rightarrow$  Belief Gamble Aversion.

*Proof.* We must show that any model that satisfies Risk-Neutral EUM and Risk-Averse Veritism must satisfy Belief Gamble Aversion as well. Begin by letting  $P(q) \leq .5$ , and assume (without loss of generality) that suspension of judgment has no value or disvalue, that is,  $V(s, w) = 0$ , for all  $w$ . Given this specification of  $P$  and  $V$ , it follows trivially that  $E_P(s) = 0$ . Due to Extensionality, we can define two constants,  $\nu_1$  and  $\nu_2$ , such that  $\nu_1$  denotes the value of  $V(b, w)$ , for each  $w \in q$ , whereas  $\nu_2$  denotes the value of  $V(b, w)$ , for each  $w \in \sim q$ . The expected value of  $b$  is then given by:

$$\begin{aligned}
E_P(b) &= \sum_{w \in W} V(b, w) \cdot P(w) \\
&= \sum_{w \in q} V(b, w) \cdot P(w) + \sum_{w \in \sim q} V(b, w) \cdot P(w) \\
&= \sum_{w \in q} \nu_1 \cdot P(w) + \sum_{w \in \sim q} \nu_2 \cdot P(w) \\
&= \nu_1 \cdot P(q) + \nu_2 \cdot P(\sim q)
\end{aligned}$$

Since  $P(q) \leq .5$  and  $P(q) + P(\sim q) = 1$ , it follows that  $P(q) \leq P(\sim q)$ . By Risk-Averse Veritism,  $\nu_2 + \nu_1 < 0$ . It is then easily shown that  $\nu_1 \cdot P(q) + \nu_2 \cdot P(\sim q) < 0$ , which means that  $E_P(b) < E_P(s)$ . Thus, by Risk-Neutral EUM,  $b \notin R(q)$ , as dictated by Belief Gamble Aversion.  $\square$

**Third Negative Result:** Risk-Neutral EUM + Risk-Neutral Veritism  $\not\Rightarrow$  Belief Gamble Aversion.

*Proof.* The proof is immediate from the proof of the Second Negative Result.  $\square$

**Fourth Negative Result:** Risk-Neutral EUM + Risk-Neutral Veritism + Non-Contradiction  $\not\Rightarrow$  Belief Gamble Aversion.

*Proof.* We must define a model that satisfies Non-Contradiction, Risk-Neutral EUM, and Risk-Neutral Veritism, but not Belief Gamble Aversion. Let  $W = \{w_1, w_2\}$ ,  $q = \{w_1\}$ , and  $P(q) = .5$ . The following specification of  $V$  is consistent with Risk-Neutral Veritism:

$$\begin{aligned} V(b, w_1) &= V(d, w_2) = 1 \\ V(s, w_1) &= V(s, w_2) = 0 \\ V(d, w_1) &= V(b, w_2) = -1 \end{aligned}$$

Relative to this value assignment, the following specification of  $R$  is easily shown to be consistent with Non-Contradiction and Risk-Neutral EUM:

$$\begin{aligned} R(q \wedge \sim q) &= \{d\} \\ R(q) &= \{b, s, d\} \end{aligned}$$

Thus,  $b \in R(q)$  and  $P(q) \neq .5$ , contrary to Belief Gamble Aversion. □

**Third Positive Result:** Dominance + Minimal Veritism  $\Rightarrow$  Non-Contradiction.

*Proof.* We must show that any model that satisfies Dominance and Minimal Veritism must satisfy Non-Contradiction as well. Since  $q \wedge \sim q$  is false at every world in every model (that is,  $q \wedge \sim q = \{\}$  in every model), it follows as an immediate consequence of Ranking that  $V(s, w) > V(b, w)$ , for every  $w$  in every model, which means that  $s$  dominates  $b$ . Thus, by Dominance,  $b \notin R(q \wedge \sim q)$ , as dictated by Non-Contradiction. □

**Fourth Positive Result:** Risk-Neutral EUM + Risk-Neutral Veritism + Impermissivism  $\Rightarrow$  Belief Gamble Aversion.

*Proof.* We must show that any model that satisfies Risk-Neutral EUM, Risk-Neutral Veritism and Impermissivism must satisfy Belief Gamble Aversion as well. We begin by assuming that  $P(q) = .5$ . Given Risk-Neutral Veritism, it is then easy to see that  $E_P(b) = E_P(d)$ . By Risk-Neutral EUM, it follows that  $\{b, d\} \subseteq R(q)$  or  $\{b, d\} \not\subseteq R(q)$ . But Impermissivism excludes the former option, which means that  $\{b, d\} \not\subseteq R(q)$  and (*a fortiori*)  $b \notin R(q)$ , as dictated by Belief Gamble Aversion. □

**Fifth Positive Result:** Belief Gamble Aversion  $\Rightarrow$  Impermissivism.

*Proof.* We prove the contrapositive. Suppose that  $\{b, d\} \subseteq R(q)$ , contra Impermissivism. Since disbelieving  $q$  is assumed to be equivalent to believing  $\sim q$ , this supposition amounts to saying

that  $b \in R(q)$  and  $b \in R(\sim q)$ . Since  $P(q) + P(\sim q) = 1$  in any model, we get that  $P(q) \not\geq .5$  or  $P(\sim q) \not\geq .5$ . Thus, either  $b \in R(q)$  and  $P(q) \not\geq .5$  or  $b \in R(\sim q)$  and  $P(\sim q) \not\geq .5$ , contra Belief Gamble Aversion.  $\square$

**Acknowledgments:** Earlier versions of this paper were presented at a workshop at Copenhagen University and a SLLERG seminar at MIT. I'd like to thank the audiences on those occasions for helpful feedback. A special thanks to Asbjørn Steglich-Petersen for extensive discussion of the ideas developed in this paper.

## References

- Briggs, R. (2018): "The Dominance Principle in Epistemic Decision Theory." In: *Philosophy and Phenomenological Research* 66, pp. 763–775.
- Buchak, L. (2013): *Risk and Rationality*. Oxford: Oxford University Press.
- Clifford, W. (1877 [1999]): "The Ethics of Belief." In: *The Ethics of Belief and Other Essays*, T. Madigan, (ed.), Amherst, MA: Prometheus, pp. 70–96.
- Dorst, K. (2019): "Lockeans Maximize Expected Accuracy." In: *Mind* 138, pp. 175–211.
- Easwaran, K. (2016): "Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities." In: *Noûs*, 50, pp. 816–853.
- (2013): "Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection." In: *Philosophy of Science* 180, pp. 119–142.
- Easwaran, K. & B. Fitelson (2015): "Accuracy, Coherence, and Evidence." In: *Oxford Studies in Epistemology* 5, pp. 61–96.
- Goldman, A. (1999): *Knowledge in a Social World*. Oxford: Clarendon Press.
- Greaves, H. & D. Wallace (2006): "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." In: *Mind* 115, pp. 607–632.
- James, W. (1896 [1956]): "The Will to Believe." In: *The Will to believe and other essays in popular philosophy*, New York: Dover Publications, pp. 1–31.
- Joyce, J. (1998): "A Nonpragmatic Vindication of Probabilism." In: *Philosophy of Science* 65, 575–603.
- (2009): "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In: *Degrees of Belief*, F. Huber & C. Schmidt-Petri (eds.), Synthese library 342, pp. 263–297.
- Kelly, T. (2013): "Evidence Can Be Permissive." In: *Contemporary Debates in Epistemology*, M. Steup & J. Turri (eds.), Blackwell.

- Kopec, M. and M. Titelbaum (2016): "The Uniqueness Thesis." In: *Philosophy Compass* 11, pp. 189–200.
- Oddie, G. (1997): "Conditionalization, Cogency, and Cognitive Value." In: *British Journal for the Philosophy of Science* 48, pp. 533–541.
- Pettigrew, R. (2016a): "Accuracy, Risk, and the Principle of Indifference." In: *Philosophy and Phenomenological Research* 92, pp. 35–59.
- (2016b): "Jamesian Epistemology Formalized: An Explication of 'The Will to Believe'." In: *Episteme* 13, pp. 253–268.
- (2016c): *Accuracy and the Laws of Credence*. Oxford University Press.
- Rawls, J. (1971): *A Theory of Justice*. Harvard University Press.
- Schoenfield, M. (2014): "Permission to Believe: Why Permissivism is True and what it Tells us about Irrelevant Influences on Belief." In: *Noûs* 48, pp. 198–218.
- Steinberger, F. (2019): "Accuracy and Epistemic Conservatism." In: *Analysis* 79, pp. 658–69.
- Sylvan, K. (2020): "An Epistemic Non-Consequentialism." In: *The Philosophical Review* 129, pp. 1–51.
- White, R. (2005): "Epistemic Permissiveness." In: *Philosophical Perspectives* 19, pp. 445–459.
- Williamson, T. (2000): *Knowledge and Its Limits*. Oxford University Press.