



An Epistemic Version of Pascal's Wager

ABSTRACT: *Epistemic consequentialism is the view that epistemic goodness is more fundamental than epistemic rightness. In this article, I examine the relationship between epistemic consequentialism and theistic belief. I argue that in an epistemic consequentialist framework, there is an epistemic reason to believe in God. Imagine having an unlimited amount of time to ask an omniscient being anything you wanted. The potential epistemic benefits would be enormous. Considerations like these point to an epistemic version of Pascal's wager. I compare and contrast the epistemic wager with the traditional wager, and argue that the epistemic wager has several notable advantages.*

KEYWORDS: Pascal's wager, theistic belief, epistemic consequentialism, epistemic value, heaven and hell

Introduction

Most theistic arguments are for the conclusion that God exists. Pascal's Wager, in contrast, concludes that one ought to believe in God. The advocate of Pascal's wager reasons that if God exists and one believes in God, one has infinity to gain. If God exists and one does not believe in God, one has an infinite loss. If God does not exist, what one gains or loses is finite either way. Thus, one should believe in God (table 1).

In table 1, ω represents the infinite utility or goodness associated with heaven, and $-\omega$ represents the infinite pain, disutility, or badness associated with hell. F_1 and F_2 are finite values. To account for the possibility of annihilationism—and illustrate that one need not believe in hell for the argument to go through—included is the possibility of 0 utility if one does not believe in God and God exists. Note that the probability that God exists, n , needs to be non-zero and non-infinitesimal.

There are many objections to the traditional version of Pascal's wager. A classic problem is the many-gods objection; advocates point out that there are a variety of mutually exclusive religions, and the wager does not tell us how to pick between them (Mougin and Sober 1994). Other objections include the worry that we cannot control our beliefs—and even if we could, beliefs formed for practical reasons are irrational (Mackie 1982: 64, 201; Jones 1998: 173). Still others have

Thanks to John Hawthorne for a conversation that inspired this article. Thanks to Dustin Crummett, Andrew Moon, Wes Siscoe, and two anonymous referees at this journal for comments on earlier drafts. Thanks to Josh Barthuly, Josh Brecka, Justin D'Ambrosio, Perry Hendricks, Brett Lunn, Justin Mooney, Daniel Muñoz, Andy Rogers, Eric Sampson, Rohan Sud, Kevin Vallier, Johnny Waldrop, and the audiences at the Toronto Metropolitan University Philosophy of Religion work-in-progress group and the New Theists Workshop for helpful comments and questions.



Table 1 The Traditional Wager

	God exists (n)	God does not exist (1-n)	Expected value
Believe in God	ω	F1	ω
Don't believe in God	$-\omega$ or 0	F2	$-\omega$ or $F2(1-n)$

worried about the use of infinities in decision theory in general (Duff 1986; Hájek 2003). While there are responses to these objections (see Jackson and Rogers 2019; Chen and Rubio 2020; Jackson 2023a), I set them aside for now. My goal is not to defend the traditional version of the wager, and I will address these objections insofar as they apply to my reformulated wager.

Important for my purpose is that, insofar as the traditional wager is successful, it provides a *practical* reason for belief in God—one can maximize expected value or utility in the long run by being a theist. In what follows, I defend an alternative version of the wager that instead provides an *epistemic* reason to believe in God. This alternative wager centrally relies on a normative theory in epistemology called *epistemic consequentialism*. More specifically, if epistemic consequentialism is true, then there is *epistemic* reason to believe in God.

I proceed as follows: In section 1, I explain epistemic consequentialism, give some basic motivations for it, and distinguish various versions of it. In section 2, I apply epistemic consequentialism to theistic belief and describe how this results in an epistemic version of Pascal's wager. I compare and contrast my wager with the traditional wager and explore various upshots of the argument. In section 3, I respond to objections to the reformulated wager, and in section 4, I conclude.

1. Epistemic Consequentialism

Ethical consequentialism is traditionally defined as the view that the good is more fundamental than, or explains, the right. Epistemic consequentialism applies this theory to the epistemic realm: namely, that epistemic goodness explains epistemic rightness. In their volume on epistemic consequentialism, Ahlstrom-Vij and Dunn define epistemic consequentialism as 'epistemic rightness—denoted by terms such as "justification" or "rationality"—is to be understood in terms of conduciveness to epistemic goods' (2018: 2). On this view, epistemic goods are *value-promoting*: conforming to epistemic norms is more likely to cause or lead to epistemically valuable states of affairs than is nonconformity. My goal in this section is threefold. One, to give the reader an overview of epistemic consequentialism as a normative theory, and two, to indicate key choice points that have later implications. My third goal is to show that epistemic consequentialism is a defensible and reasonable view; however, I will not give a full-fledged defense of it.

Several authors, including Daniel Singer and Alvin Goldman, have argued that it is natural to understand epistemology in a consequentialist framework. Singer says, 'Epistemic consequentialism is an intuitive natural starting point for theorizing a view of epistemic normativity' (2019: 269). Goldman similarly writes, 'People have interest, both intrinsic and extrinsic, in acquiring knowledge (true belief) and

avoiding error. It therefore makes sense to have a discipline that evaluates intellectual practices by their causal contributions to knowledge or error. This is how I conceive of epistemology: as a discipline that evaluates practices along truth-linked (veritistic) dimensions. Social epistemology evaluates specifically social practices along these dimensions' (1999: 69).

Both Singer and Goldman suggest that epistemic consequentialism is an intuitive starting place for epistemology. Why think this? An attractive idea is that our epistemic goal is to maximize the number of truths we believe and minimize the number of falsehoods we believe. 'Believe truth; shun error' is a mantra that many epistemologists have considered to be at the core of epistemic normativity since William James ([1896] 1979: 13). Goldman (1999), for instance, conceives of epistemology as a consequentialist enterprise on which the fundamental goods are seeking truth and avoiding error. This picture lends itself to a consequentialist view of epistemic normativity.

However, we need not make epistemic consequentialism merely about believing the truth and avoiding error. One thing we learn from the ethical case is that consequentialism is flexible because it does not require us to commit to a certain view about what is good or valuable—it simply says that, whatever the good is, we should promote it (Portmore 2009; Brown 2011; Muñoz 2021). This leads naturally to the question: What is epistemically good?

Before surveying various answers to this question, note that there are two senses of epistemic value. The first is *instrumental* epistemic value: something is valuable in this sense just in case it causes or leads to things that are themselves epistemically valuable. For example, some argue that we ought to follow our evidence because it is a good guide to true belief. Nonetheless, evidentially supported belief is not good in and of itself; it is merely instrumentally good (BonJour 1985).

Second, there is *final* or *intrinsic* epistemic value. A good is finally or intrinsically epistemically valuable if and only if it is epistemically valuable in and of itself. The most popular view of final epistemic value is *veritism*, discussed above; on this view, accuracy (that is, truth) is the final epistemic value, and accuracy is understood in terms of gaining accurate beliefs (or credences) and avoiding inaccurate beliefs (or credences). Goldman (1999: 5; 2001: 31) and Singer (2018, 2019) express sympathy for this view, as do many advocates of epistemic utility theory.

But there are numerous other views of final epistemic value. On another view, knowledge is the final epistemic value (Littlejohn 2018), and those who are convinced, as was Plato, that knowledge is more valuable than true belief will likely be sympathetic to this view (see DePaul 2001). Others argue that evidentially supported belief (or credence) is of final epistemic value (Feldman 2000). On yet another view, understanding is finally epistemically valuable (Kvanvig 2003; Pritchard 2010). Others argue that possessing epistemic virtue is finally epistemically valuable (Turri and Alfano 2021). There are also views on which (some of) the states above are valuable, but only when they contain *important* or *significant* content (Friedman 2018; Talbot 2019). Finally, there are pluralist views that maintain there are several final epistemic values. All of these views additionally include commitments about what is epistemically

disvaluable—such as false belief, lack of understanding, beliefs that do not amount to knowledge. It is not necessary to take a stand on any of these views, for, as I discuss below, my argument is consistent with all of them.

Two final distinctions play a key role in determining epistemic value. The first is the question of whether epistemic norms are *categorical* and apply to everyone, or whether they are *instrumental* and vary person to person, depending on each individual's preferences or desires. Consider moral norms: it is commonly thought that they are categorical—universally applicable, and not goal- or preference-relative. On some views, such as Kelly (2003), epistemic norms are similarly universally applicable.

On the instrumentalist view, epistemic norms depend on human goals or desires, and thus are a means-ends norm. Means-ends norms are derivable from one's mental states (as long as those states satisfy certain constraints, like transitivity). Decision theory, widely taken to be our best theory of means-ends rationality, does not tell one what to desire or believe, but how to act *given* one's beliefs or credences and desires or utility function. On a similar instrumentalist view in epistemology, epistemic goods are valuable insofar as one desires them.

On either view, there is some agreement for what counts as epistemically valuable, as people often desire the candidate universal epistemic goods, such as certain pieces of knowledge or true or rational belief. My argument relies on the overlap between the two views, so it is not necessary to take a stand on this controversy regarding epistemic norms.

Finally, there are diachronic issues concerning epistemic value. Consider an axiology on which true beliefs are categorically epistemically valuable. There are two ways to measure this value. (For simplicity, I bracket disvalue.) On a *synchronic* measure of epistemic value, each instance of learning a truth is valuable, no matter how long a token belief is held. To calculate epistemic value, simply add up the total number of true beliefs one had, for any amount of time. On this view, forgetting a previously known truth does not detract from epistemic value, which is (arguably) counterintuitive. On a *diachronic* measure of epistemic value, the value of true beliefs is based on the length of time they are held, so, such as having a true belief for a hundred days is better than having a true belief for one day. This view might initially seem more plausible, but it has the counterintuitive consequence that living longer makes one epistemically better off. Again, while this distinction is relevant below, it is not necessary to take a stand on the best way to measure epistemic value over time.¹

It is also worth surveying some epistemological views that are consequentialist in nature. The most basic form of epistemic consequentialism is analogous to act-consequentialism in ethics. This view is defended by Singer (2018), among others, and summarized by Ahlstrom-Vij and Dunn: 'a belief is epistemically justified (or rational) just in case the total causal consequences of holding or forming that belief contain more epistemic good than the total causal consequences of holding or forming any alternative doxastic attitude' (2018: 2).

¹ For more on diachronic issues concerning epistemic value, see Kawall 2010; Treanor 2013; Whitcomb 2018; Aronowitz 2021; Singer and Aronowitz 2022. Thanks to Dustin Crummett for helpful discussion.

A related epistemic consequentialist view is reliabilism (Goldman 1979). On reliabilism, a belief is justified just in case it is reliably formed, where reliability is understood in terms of yielding a certain ratio of true to false beliefs (Berker 2013a: 368). This can be understood either over time or across possible worlds. Either way, on reliabilism, epistemic justification is understood in terms of achieving a favorable ratio of true to false beliefs in the long run—which may be analogous to epistemic rule-consequentialism (Ahlstrom-Vij and Dunn 2023).

Epistemic consequentialism goes beyond reliabilism, however. Selim Berker (2013b) argues that a large number of views in epistemology are consequentialist, such as Plantinga's proper functionalism (1993), Sosa's virtue epistemology (2011), Foley's egocentric epistemology (1993), and Alston's epistemic desiderata view (1993). Berker also claims 'there is a way of viewing much of the debate between externalists and internalists about epistemic justification over the past few decades as a debate within epistemic consequentialism, roughly corresponding to the debate between objective and subjective consequentialists in the ethical realm' (Berker 2013b: 368). Some versions of evidentialism and coherentism can be understood as consequentialist (that is, insofar as they ask us to believe the thing that leads to the maximum number of evidentially supported beliefs or coherent beliefs, respectively).² Epistemic consequentialism also has defenders in the philosophy of science (Solomon 1992; Laudan 1984).

Finally, consider epistemic utility theory (Joyce 1998; Konek and Levinstein 2019). Epistemic utility theory employs a model similar to decision theory, and its defenders argue that belief states that meet certain constraints are to be preferred, as they lead to better epistemic outcomes, where 'better' is normally understood in terms of accuracy. Epistemic utility theory has many applications, including a new, non-pragmatic argument that coherent credences are rationally required. The accuracy-dominance argument for probabilism shows that, for all incoherent credences, there are always alternative credences that are guaranteed to be more accurate (see Greaves and Wallace 2006; Leitgeb and Pettigrew 2010a, 2010b; Greaves 2013; Easwaran 2013; Easwaran and Fitelson 2015; Pettigrew 2016; for an overview of epistemic utility theory, see Pettigrew 2013).

Thus, epistemic consequentialism is a popular and fruitful view that many epistemologists take seriously, and it is consistent with numerous mainstream epistemological views. While, of course, objections have been raised to epistemic consequentialism—notably, by Berker (2013a, 2013b) and Kurt Sylvan (2020a, 2020b)—it is nonetheless both plausible and widely held.

2. The Epistemic Wager

In the next two sections, I defend my main claim: if epistemic consequentialism is true, then we epistemically ought to believe in God. In this section, I present an alternative decision matrix that represents the epistemic wager. I explain the

² One might reply that if all these views count as consequentialist, then we need a narrower definition of epistemic consequentialism (see Ho 2022). My argument does not hang on which views in particular should be understood as consequentialist. Thanks to Dustin Crummett and Andrew Moon for helpful discussion.

relevant aspects of the matrix and defend the required premises. I also discuss the difference between finite and infinite versions of the epistemic wager and compare the epistemic wager to the traditional one. In the next section, I respond to objections.

2.1 Explaining the Epistemic Wager

Consider the following decision matrix, which represents the epistemic wager (table 2):

In table 2, U represents the value of the epistemic goods associated with going to heaven. Notice that those who believe in God receive U , unless God does not exist, but those who do not believe in God receive U only if universalism is true and everyone goes to heaven (but, on this outcome, they are dinged for having a false belief). V is the value of the epistemic goods associated with going to hell. There may be *some* epistemic goods associated with hell, assuming people in hell learn things. Nonetheless, I argue that U is greater than V . Also, as I discuss below, U and V may be infinite or finite.

F_1 , F_2 , F_3 , and F_4 are all finite. F_2 and F_3 , while similar, have different values because theistic beliefs have different truth values on each outcome. F_4 represents the badness of having a false theistic belief during one's pre-afterlife existence; its value may be negligibly small.³ P_1 – P_4 are probabilities. P_1 is the probability of [theism and the traditional view of hell], P_2 is the probability of [theism and the annihilationist view of hell, that is, that nonbelievers cease to exist after death], and P_3 is the probability of [theism and universalism, that is, that all go to heaven]. P_4 is the probability of atheism.

This table assumes a couple things. First, given theism, there are three possible afterlives for nonbelievers that are mutually exclusive and exhaustive. If we understand these afterlives broadly—that is, heaven represents the net good afterlives, hell represents the net bad afterlives, and annihilationism represents the net neutral afterlives—it is plausible that they exhaust the probability space. Second, for the argument to work, the disjunction of P_2 or P_3 —that is, either annihilationism or the traditional view of hell—must be (at least) non-zero and non-infinitesimal. Note that this is consistent with thinking universalism is quite likely, and even with *believing* universalism, as long as one does not assign universalism credence 1. On certain finite versions of the epistemic wager, this assumption may need to be slightly stronger, but it is in most cases consistent with believing universalism.

The argument requires that either (1) the value of P_2 , the probability of annihilationism, is sufficiently high; or (2) $U > V$, that is, it is epistemically better to be in heaven than to be in hell. On most credence functions, the value of P_2 will be high enough for the argument to go through, but there are also good reasons to think that (2) holds—that U is in fact greater than V . Here are a few. Consider the beatific vision—the doctrine that those in heaven are face to face with God and see God's essence. This is commonly understood to have an

³ Thanks to Josh Brecka.

Table 2 The Epistemic Wager

	God exists and nonbelievers go to hell (P ₁)	God exists and nonbelievers are annihilated (P ₂)	God exists and universalism is true (P ₃)	God does not exist (P ₄)
Believe in God	U	U	U	F ₁
Don't believe in God	V (< U)	F ₃	U - F ₄	F ₂

intellectual component, where one gains new beliefs, knowledge, or understanding about God, themselves, and the world in an indescribably rich and glorious way (Aquinas 1265–1274; Boersma 2018). Further, on many theological traditions, those in heaven have direct access to God. They can ask God any question, about any topic that they want. Not only is this objectively epistemically valuable on almost any theory of epistemic value—one has access to unlimited true beliefs, knowledge, accurate credences, understanding, and the like—but it is also an epistemic good that is likely to be desirable to many. Thus, it is also valuable if epistemic norms are instrumental, derived from the one's epistemic goals. One can learn any piece of information by simply asking God about it.

Further, while it is possible that people in hell learn things, and potentially even attain significant knowledge, if hell is in other ways undesirable, unpleasant, or involves serious suffering, people there are less likely to pursue epistemic goods. There is also the possibility that hell is epistemically bad because it is a place of confusion and self-deception; those in hell may actually gain false, unjustified, or incoherent beliefs. Thus, hell may have an overall negative epistemic value. Finally, recall that V is the epistemic value of not believing in God if God exists. Those who do not believe in God thus miss out on at least one true belief (but potentially more, as they are also likely to miss out on true beliefs that follow from or cohere with theism).

This raises an important, related question: whether U (and V) are infinite or finite. Is afterlife *infinitely* epistemically good? Suppose those in heaven gain a finite number of epistemically valuable states per day (or per finite period) for all eternity. On a synchronic measure of epistemic value, then, heaven is infinitely epistemically valuable.⁴ If one also has epistemically valuable states for an infinite time period, then heaven is also infinitely epistemically good on a diachronic measure of epistemic value. I do not completely rule out the possibility that hell is infinitely epistemically valuable for similar reasons, even though I gave some reasons to doubt this above. If heaven and hell are both infinitely good, but heaven is better, then to compare them, we need a decision theory that enables the comparison of different sizes or kinds of infinities.

⁴ There is one notable exception to this: if people simply learn, forget, relearn, and re-forget the same truths over and over. However, this afterlife seems unlikely, and the argument can go through as long as we do not assign it credence 1. Thanks to Dustin Crummett.

Not all infinite epistemic goods are decision-theoretically equal. For example, getting one hundred pieces of new knowledge a day is better than getting one per day, which is better than getting one every thousand years, even though these are all, strictly speaking, infinite. Further, if the content of what is learned matters, then we should prefer gaining a significant piece of knowledge every day to gaining an insignificant piece of knowledge every day. Finally, we should prefer a 99 percent chance at getting an infinite epistemic good than a 1 percent chance at the same good, so both utility and probability matter in the infinite case. (Plausibly, we can use ordinally ranked infinities to compare these cases.)

The lesson—that we should not treat all infinities the same in decision theory—is one that many have drawn in response to the mixed strategies objection to the traditional wager, and there are numerous formal models to compare infinities that can easily be translated over to the epistemic wager. These models utilize mathematical tools such as limits and ratios or surreal numbers (see, for instance, Bartha 2007; Jackson and Rogers 2019; Chen and Rubio 2020). The epistemic wager is consistent with a variety of formalisms; we need not commit to a specific way to compare infinities here, as they agree on the main results we need. On all of these models, given an infinite version of the wager, the value of F_1 and F_2 and the probability of [hell or annihilationism] will not matter, as long as the probability of [hell or annihilationism] is non-zero and non-infinitesimal.

Nonetheless, some are skeptical of infinities in general—or at least skeptical of the use of infinities in decision theory. Others might simply question whether the epistemic gains in the afterlife are infinitely valuable. To account for this, we can utilize an alternative model on which U and V are finite but nonetheless quite large. On the finite version of the wager, the probability of the disjunction of P_2 or P_3 —annihilationism or the traditional view of hell—would need to be larger than simply non-zero/non-infinitesimal. However, we can utilize arbitrarily large finite values for U and V , so that the probability of the disjunction of P_2 and P_3 can be extremely small. Thus, the epistemic wager is consistent with afterlife epistemic goods being either finite or infinite.

Finally, for the epistemic wager to be successful, one needs to care about promoting long-term expected epistemic value (in a synchronic or diachronic sense). As with the traditional wager, if one is an epistemic future discounter, then, depending on the exact discount function, the epistemic wager may not have purchase. Some of the above epistemic consequentialist views described above (such as some versions of epistemic utility theory) focus only on the short-term epistemic consequences of, for example, adopting a certain credal state. Thus, not *all* versions of epistemic consequentialism pay attention to long-term consequences (for example, Joyce 1998).⁵ But these restricted views face other challenges. When should we stop caring about the epistemic consequences? Is there a non-ad hoc way to determine the cut-off point? Second, most philosophers think discounting in the practical case is irrational; why would it be rational in the epistemic case? (See Jackson and Rogers 2019: 79–82.) Finally, as I argue below, if one does not

⁵ Thanks to Rohan Sud.

consider total consequences of one's decisions when acting, it is not clear that one is following a consequentialist norm.

In sum, the epistemic wager requires three things: (1) a sufficient probability of [annihilationism or the traditional view of hell]—but this probability can nonetheless be extraordinarily low (how low depends on whether the wager is finite or infinite); (2) either (a) a sufficiently high probability of annihilationism (but again, this can still be relatively low), or (b) that the epistemic value of heaven is greater than that of hell; and (3) long-term expected epistemic consequences to be normatively relevant.

2.2 The Traditional Wager Verses the Epistemic Wager

There are at least four salient differences between the traditional wager and the epistemic wager, several of which are clear benefits of the epistemic wager.

The first difference involves the irrationality objection to the traditional wager. According to this objection, beliefs formed for practical reasons are epistemically irrational. If I believe *p* because you have paid me money to believe *p*, or simply because I want *p* to be true, my belief is not epistemically rational. Beliefs formed on the basis of the traditional wager are similarly based on practical reasons, and thus—so the objection goes—also epistemically irrational. This general line of reasoning is advocated by several in the literature, including Clifford (1877: 289); Flew (1976: 64); Mackie (1982: 202); Oppy (1991: 167); Schroeder (2012: 266). (For responses to the irrationality objection, see Jackson 2016; Benton 2018, Jackson and Rogers 2019; Jackson 2023c; Jackson 2023a.)

But this objection does not apply to the epistemic wager. If, as assumed, epistemic consequentialism is the correct theory of epistemic normativity, then the epistemic wager provides an *epistemic* reason to believe in God, not just a practical reason. Belief in God maximizes expected epistemic value, as it promotes epistemically valuable states and helps us avoid epistemically invaluable states. One of the most important features of the epistemic wager is that it provides epistemic, rather than practical, reasons for belief. Against the backdrop of epistemic consequentialism, the irrationality objection does not apply.

To drive this point home, consider an analogy. For a long time, probabilism, the view that one's credences should be probabilistically coherent, was taken to be justified by Dutch book arguments: one could lose money in a series of bets if they had incoherent credences. However, this view faces a key objection: this a pragmatic reason to have probabilistic credences, not an epistemic one. In response, the new accuracy arguments for probabilism use epistemic decision theory to show that for all non-probabilistic credences, some probabilistic credences are guaranteed to be more accurate (Joyce 1998). Thus, one of the major applications of epistemic decision theory is that it provides a non-pragmatic justification of probabilism. This gives us an *epistemic* reason to have coherent credences. Along similar lines, the epistemic wager provides an *epistemic* reason for belief in God, rather than merely a pragmatic one.

Consider a second way the epistemic wager differs from the traditional wager. Another worry for Pascal's wager involves the motives exemplified by taking the

traditional wager. The thought is that if one simply believes in God to avoid hell or get the reward of heaven, this seems unvirtuous—something that God would not be pleased with. While I am not convinced this objection is correct (see Jackson 2023b), those sympathetic to this worry may prefer the epistemic wager. Consider the person who wagers on God because they take it to be their best shot of getting at the truth. Being motivated by epistemic goods, such as gaining knowledge or finding answers to significant questions, is admirable. Thus, the motives involved in the epistemic wager may be more virtuous than those involved in the traditional wager.

Third, the traditional wager is often viewed as a simple trumping argument. Infinite value is always better than finite value, so one should wager on God, even if their credence in theism is extraordinarily low. The epistemic wager is more complex, partially because it is not obvious that the epistemic value is heaven is *infinitely* better than that of hell. Thus, the epistemic wager does not involve the simple trumping of the traditional wager. However, several authors have argued that even the traditional wager cannot rely on simple trumping principles, due to the mixed strategies objection (see Duff 1986; Hájek 2003; Jackson and Rogers 2019; Chen and Rubio 2020). This may be a feature of all plausible versions of Pascal's wager.

The fourth and final difference is that the epistemic wager does better the higher one's credence that non-believers are annihilated, whereas the traditional wager does better the higher one's credence that nonbelievers go to hell (unless hell has negative epistemic value). This is because hell may be epistemically better than being annihilated, but is pragmatically worse. Epistemically and pragmatically, annihilation has utility 0. By contrast, going to hell has a negative pragmatic utility, but may have a positive epistemic utility.

3. Objections

In this section, I consider four objections, involving, respectively: belief and control, trade-offs and long-term consequences, many gods, and hell versus annihilationism.

3.1 Belief and Control

One might object that, given the epistemic wager, I cannot simply make myself believe in God, even if I know that doing so will lead to a lot of epistemic goods down the line. In this, the epistemic wager faces a control problem: believing for long-term epistemic goods seems just as difficult as believing for practical reasons. In reply, I merely argue above that we epistemically ought to believe in God, and I argue for this only *conditionally*, on the supposition of epistemic consequentialism. I have not even argued that the antecedent holds, much less argued for anything about control over belief.

The objector might reply that the epistemic *ought* implies *can*. In other words, if S epistemically ought to *phi*, then S can *phi*. Given this principle, if we cannot control our beliefs, one can reject my claim that we should believe in God via modus tollens reasoning. I have two preliminary replies. First, again, this is not an objection to the conditional claim, which is my primary thesis. And second, many (Feldman 2000,

Hieronimi 2009, McHugh 2012) deny that epistemic ought implies can, so epistemic obligations can govern our beliefs even if we cannot control them.

That said, this argument would be less interesting if the antecedent and the consequent of the conditional thesis were false, and I do not want the argument to hang on whether epistemic ought implies can. This brings me to a third reply, which focuses on the acts in the epistemic decision matrix. Suppose first that the acts are as they are in the matrix above, simply *believing* and *not believing* in theism. One interesting thing about the epistemic wager is that there is less pressure to move away from this straightforward interpretation, since it provides epistemic reasons, and those apply to belief. On the straightforward interpretation, we either need to deny that ought implies can or argue that we can control beliefs. I argue elsewhere (Jackson 2023a) that when it comes to belief in God, many of us are in an *epistemically permissive* case: our evidence underdetermines what we epistemically ought to believe. I also argue that, in permissive cases, we have a much more direct kind of control than in other cases. I argue that we have either direct or what I call “semi-direct” control over belief in permissive cases, and thus we can, in some meaningful sense, choose to believe based on Pascalian considerations. Further, almost everyone who works on issues about belief and control agrees that we have *indirect* doxastic control: we can *try* to believe certain things, and, in many cases, successfully influence our beliefs over time. This is even more plausible when we consider propositions that are candidates for being permissive (like *God exists*) rather than propositions that are clearly true or false, like $1+1=2$ or *the Earth is flat*.

Nonetheless, I am also open to interpreting the acts in the epistemic wager more broadly. We might frame the wager in terms of things besides belief, such as desire or action. Alternatively, the wager might be understood not in terms of belief but in terms of *faith* (assuming faith does not entail belief, as contended by Pojman 1986; Audi 1991; Alston 1996, among others). Or the act in question might be a *commitment* to God, a commitment to pursue a relationship with God, or a commitment to accept (that is, act as if) God exists (see Golding 1994; Jones 1998; Rota 2016a, 2016b). While this requires that there are epistemic reasons for action—which is controversial; see Simion (2018)—this is another viable version of the epistemic wager.

There are thus at least three ways out of this objection to the epistemic wager. One, we frame the epistemic wager about believing or trying to believe, and argue that we have the relevant kind of control, even if indirect. Two, we deny that epistemic ought implies can, or three, we focus on epistemic reasons for action or commitment.

3.2 Trade-Offs and Long-Term Consequences

According to the epistemic wager, it is epistemically rational to believe in God because it leads to a large expected epistemic benefit. However, one might argue that leading to a bunch of other true or rational beliefs about *q* and *r* cannot make my belief that *p* epistemically rational; this is what Berker (2013a) calls ‘the separateness of propositions’. In other words, simply because believing in God

will lead to a bunch of expected epistemic value down the line does not make my belief in God epistemically rational, especially if that value concerns other propositions.

In reply, first note that the trade-off worry is an objection to epistemic consequentialism, which I am not arguing for. This objection parallels a classic objection to ethical consequentialism, namely, that ethical consequentialism does not respect the separateness of persons (but in the epistemic case, it is the separateness of *propositions*). Of course, what one's view says about this case depends on their epistemic axiology. However, many epistemic consequentialists do not merely value the accuracy of a single attitude, but the accuracy of a belief or credal state as a whole.

And it is hard to see how one could have a genuine version of epistemic consequentialism without committing to the permissibility of *some* epistemic trade-offs. Littlejohn (2018) and Wedgwood (2018) both emphasize that, in a genuine consequentialist theory, the total consequences must be used to derive the normative property. As Ahlstrom-Vij and Dunn (2018: 6) point out, if all that matters for the thing being evaluated is just the thing itself, it is not clear how this is a version of consequentialism. On most versions of consequentialism, we do not restrict the relevant consequences, as long as they concern things in our value theory. And being concerned only with one's attitude toward a single proposition does not seem like a sensible epistemic value theory.

Further, if I am right that many of us are in an epistemically permissive case with respect to theism, so that we could be rational as theists, atheists, or agnostics, then it is plausible that downstream epistemic goods could be a tiebreaker. If multiple attitudes to *p* are evidentially on a par, then one could consider other epistemic consequences that go beyond one's attitude toward *p*. To consider an analogy from the ethical case, consider the choice between two actions: action 1 saves only *S*'s life, and action 2 saves *S*'s life and a bunch of other lives. Even if *S* has special axiological consideration, action 2 seems better, as it dominates action 1.

Finally, note that the epistemic wager may be successful apart from a full-blown commitment to epistemic consequentialism. There are at least two possibilities. First, one might endorse epistemic consequentialism with side constraints. Maybe one should believe the thing that has the best epistemic consequences unless their evidence suggests it is definitely or very likely false. For example, believe that *God exists* because of the large expected epistemic value, unless your credence that *God exists* falls below some threshold. This would avoid the consequence that *S*'s believing *p* can be epistemically rational even if, say, *S* thinks *p* is obviously false; *S*'s epistemic position toward *p* must be sufficiently favorable. (Relatedly, Michael Rota [2016b] defends a version of the wager on which one should wager if their credence in theism is at least 0.5. The threshold need not be that high, however.)⁶

Second, there are deontological views in ethics on which one ought not to commit certain acts unless the consequences are extreme. For example, one might think torture is wrong unless, for example, a million lives are at stake. One could adopt a similar view in epistemology: epistemic consequences do not matter for

⁶ Thanks to Justin Mooney.

rationality unless the relevant consequences are sufficiently weighty. When the stakes are quite high, however, epistemic rationality depends on consequences. Because there is so much to gain in the epistemic wager, we ought to believe in God, even if the consequences are not relevant when the epistemic stakes are lower (compare: we should not torture someone to save three lives).⁷

3.3 Many Gods

One might wonder what the epistemic wager implies about mutually exclusive religions. Suppose I take the epistemic wager and believe in the God of religion 1; I may not get the epistemic goods of religion 2.

My response is twofold. First, it is sensible to simply ask the question of whether one ought to be a theist, and then leave the choice between religions for later. The main purpose of this essay is not to adjudicate between religions, but to draw a connection between general theism and epistemic consequentialism.

Second, if all religions offer roughly the same epistemic goods, then according to the epistemic wager, one should practice the religion that one thinks is most likely to be true. It is not obvious, however, that we should quantify the epistemic value of all religions equally; the epistemic goods to be gained in the afterlife vary widely by religious tradition. For example, as discussed above, on many versions of Christianity, heaven involves the beatific vision—seeing God's essence and gaining a complete understanding of the world. Alternatively, by practicing certain versions of Buddhism, one can eventually reach enlightenment, the state of comprehending reality as it is (see Wright 2016), which also seems quite epistemically valuable. Mormonism teaches that humans can eventually become divine, as they 'progress toward perfection and ultimately realize their divine destiny as heirs of eternal life' (Church of Jesus Christ of the Latter-Day Saints 1995). The epistemic benefits of being divine seem quite significant, to say the least. Thus, various religions at least offer *different* epistemic goods; there is a further question of how to weigh these goods against each other. If one thinks that a certain religion offers better epistemic goods than another, one can also factor that into the decision table; if one takes the goods of various religions to be infinite, one can utilize one of the decision-theoretic methods for comparing infinities discussed above. A detailed theological comparison of the epistemic goods of various religions is a promising area of further research.⁸ (For more on using the wager to compare religions, see Jackson and Rogers 2019.)

3.4 Hell versus Annihilationism

Above, I suggest that hell may, in some sense, be better than annihilationism. One may wonder: Is this right? This seems to have counterintuitive consequences. For instance, if we take heaven out of the picture, then, from an epistemic perspective, I should maximize my chance of going to hell.

⁷ Thanks to Eric Sampson and Kevin Vallier.

⁸ Thanks to Wes Siscoe for helpful discussion.

In reply, there are two views of hell relevant for this objection. On the first, hell is genuinely epistemically better than annihilationism. Nonetheless, this does not mean that one should all-things-considered prefer hell—hell is merely better from an epistemic point of view. To decide what one should all-things-considered prefer, one may need to weigh epistemic and practical goods against each other. (This is also potentially relevant to the many-gods issue, because there is also the possibility that one religion's afterlife is epistemically better, and another's is practically better. This would also require weighing different types of goods against each other.) Second, as mentioned above, there is also the possibility that hell is a place of self-deception and confusion. On this picture, the epistemic wager looks more like the traditional wager—hell has a negative epistemic value. Then one should prefer annihilation to hell.

4. Conclusion

In conclusion, I've defended a cousin of Pascal's wager, conditional on the truth of epistemic consequentialism. More specifically, I argued that if epistemic consequentialism is true, then there is epistemic reason to believe in God. I explained epistemic consequentialism, and then showed how it gives us the resources to defend an epistemic wager. I contrasted the epistemic wager with the traditional wager, noted that the epistemic wager has some significant benefits, and responded to objections.

Of course, some may simply take this as a *reductio* of epistemic consequentialism. Rejecting epistemic consequentialism is a straightforward way to avoid the epistemic wager. But, as shown above, epistemic consequentialism has many defenders, and some maintain that it is a natural starting point for epistemology. If that is so, then my argument has the surprising conclusion that epistemic consequentialists should be theists—even if it motivates others to reject the consequentialism they defend. My primary goal is simply to defend a conditional claim; one can choose to *ponens* or *tollens* it.

A final possibility involves combining the traditional pragmatic wager with the epistemic wager. So, for instance, if there are both strong epistemic and practical reasons to believe in God—because of the enormous epistemic *and* practical benefits of the afterlife—then we might have an all-things-considered reason to be a theist. This possibility merits further exploration.⁹

Traditionally, many have thought that Pascal's wager is a distinctively pragmatic argument for theism. If this version of the wager is successful, however, then what sounds like an oxymoron may actually be true: there are *Pascalian epistemic* reasons to believe in God.

ELIZABETH JACKSON 

TORONTO METROPOLITAN UNIVERSITY

lizjackson111@gmail.com

⁹ Thanks to Wes Siscoe.

References

- Ahlstrom-Vij, Kristoffer, and Jeffrey Dunn. (2023) 'Why No True Reliabilist Should Endorse Reliabilism'. *Episteme*, 20, 39–56.
- Alston, William. (1993) 'Epistemic Desiderata'. *Philosophy and Phenomenological Research*, 53, 527–51.
- Alston, William. (1996) 'Belief, Acceptance, and Religious Faith'. In Jeff Jordan and Daniel Howard-Snyder (eds.), *Faith, Freedom, and Rationality: Philosophy of Religion Today* (Lanham: Rowman & Littlefield), 3–27.
- Aquinas, Thomas. (1265–1274) 'Question 92. Of the Vision of the Divine Essence in Reference to the Blessed'. In *Summa Theologica*, Supplement (XP) to the Third Part [. . .] Christian Classics Ethereal Library. 6645–64. <https://www.ccel.org/ccel/a/aquinas/summa/cache/summa.pdf>. Accessed September 4, 2023.
- Aronowitz, Sara. (2021) 'Exploring by Believing'. *Philosophical Review*, 130, 339–83.
- Audi, Robert. (1991) 'Faith, Belief, and Rationality'. *Philosophical Perspectives*, 5: 213–39.
- Bartha, Paul. (2007) 'Taking Stock of Infinite Value: Pascal's Wager and Relative Utilities'. *Synthese*, 154, 5–52.
- Benton, Matthew. (2018) 'Pragmatic Encroachment and Theistic Knowledge'. In Matthew A. Benton, John Hawthorne, and Dani Rabinowitz (eds.), *Knowledge, Belief, and God: New Insights in Religious Epistemology* (Oxford: Oxford University Press), 267–87.
- Berker, Selim. (2013a) 'Epistemic Teleology and the Separateness of Propositions'. *Philosophical Review*, 122, 337–93.
- Berker, Selim. (2013b) 'The Rejection of Epistemic Consequentialism'. *Philosophical Issues*, 23, 363–87.
- Boersma, Hans. (2018) *Seeing God: The Beatific Vision in Christian Tradition*. Grand Rapids: Eerdmans.
- BonJour, Laurence. (1985) *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Brown, Campbell. (2011) 'Consequentialize This'. *Ethics*, 121, 749–71.
- Chen, Eddy Keming, and Daniel Rubio. (2020) 'Surreal Decisions'. *Philosophy and Phenomenological Research*, 100, 54–74.
- The Church of Jesus Christ of Latter-Day Saints. (1995) 'The Family: A Proclamation to the World'. General Relief Society Meeting, September 23, 1995, Salt Lake City, Utah. <https://www.churchofjesuschrist.org/study/general-conference/1995/10/the-family-a-proclamation-to-the-world?lang=eng>. Accessed August 31, 2023.
- Clifford, William K. 'The Ethics of Belief'. (1877) *Contemporary Review*, 29: 289–309.
- DePaul, Michael R. (2001) 'Value Monism in Epistemology'. In Matthias Steup (ed.) *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (Oxford: Oxford University Press), 170–84.
- Duff, Anthony. (1986) 'Pascal's Wager and Infinite Utilities'. *Analysis*, 46, 107–9.
- Dunn, Jeffrey, and Kristoffer Ahlstrom-Vij. (2018) 'Introduction: Epistemic Consequentialism'. In Kristoffer Ahlstrom-Vij and Jeffrey Dunn (eds.), *Epistemic Consequentialism* (Oxford: Oxford University Press), 1–22.
- Easwaran, Kenny. (2013) 'Expected Accuracy Supports Conditionalization—And Conglomerability and Reflection'. *Philosophy of Science*, 80, 119–42.
- Easwaran, Kenny, and Brandon Fitelson. (2015) 'Accuracy, Coherence, and Evidence'. In Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*. Vol. 5 (Oxford: Oxford University Press), 61–96.
- Feldman, Richard. (2000) 'The Ethics of Belief'. *Philosophy and Phenomenological Research*, 60, 667–95.
- Flew, Antony. (1976) *The Presumption of Atheism and Other Philosophical Essays on God, Freedom, and Immortality*. New York: Harper and Row.
- Foley, Richard. (1993) *Working without a Net: A Study of Egocentric Epistemology*. New York: Oxford University Press.

- Friedman, Jane. (2018) 'Junk Beliefs and Interest-Driven Epistemology'. *Philosophy and Phenomenological Research*, 97, 568–83.
- Golding, Joshua. (1994) 'Pascal's Wager'. *Modern Schoolman*, 71, 115–43.
- Goldman, Alvin I. (1979) 'What Is Justified Belief?' In George S. Pappas (ed.), *Justification and Knowledge: New Studies in Epistemology* (Dordrecht: D. Reidel), 1–23.
- Goldman, Alvin I. (1999) *Knowledge in a Social World*. Oxford: Clarendon Press.
- Goldman, Alvin I. (2001) 'The Unity of the Epistemic Virtues'. In Abrol Fairweather and Linda Zagzebski (eds.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility* (Oxford: Oxford University Press), 30–48.
- Greaves, Hilary. (2013) 'Epistemic Decision Theory'. *Mind*, 122, 915–52.
- Greaves, Hilary, and David. Wallace. (2006) 'Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility'. *Mind*, 115, 607–32.
- Hájek, Alan. (2003) 'Waging War on Pascal's Wager'. *Philosophical Review*, 112, 27–56.
- Hieronymi, Pamela. (2009) 'Believing at Will'. *Canadian Journal of Philosophy*, supp. vol., 35, 149–87.
- Ho, Tsung-Hsing. (2022) 'Are There Any Epistemic Consequentialists?' *Episteme*, 19, 220–30.
- Jackson, Elizabeth. (2016) 'Wagering against Divine Hiddenness'. *European Journal for the Philosophy of Religion*, 8, 85–108.
- Jackson, Elizabeth. (2023a) 'A Permissivist Defense of Pascal's Wager'. *Erkenntnis*, 88, 2315–40.
- Jackson, Elizabeth. (2023b) 'Faithfully Taking Pascal's Wager'. *The Monist*.
- Jackson, Elizabeth. (2023c) 'Pragmatic Arguments for Theism'. In Jonathan Fuqua, John Greco, and Tyler Dalton McNabb (eds.), *The Cambridge Handbook of Religious Epistemology* (Cambridge: Cambridge University Press), 70–82.
- Jackson, Elizabeth, and Andy Rogers. (2019) 'Salvaging Pascal's Wager'. *Philosophia Christi*, 21, 59–84.
- James, William. [1896] (1979) 'The Will to Believe'. In Frederick Burkhardt, Fredson Bowers, and Ignas K. Skrupskelis (eds.), *The Will to Believe and Other Essays* (Cambridge, MA: Harvard University Press), 2–32.
- Jones, Ward. (1998) 'Religious Conversion, Self-Deception, and Pascal's Wager'. *Journal of the History of Philosophy*, 36, 167–88.
- Joyce, James M. (1998) 'A Nonpragmatic Vindication of Probabilism'. *Philosophy of Science*, 65, 575–603.
- Kawall, Jason. (2010) 'Autonomy, Agency, and the Value of Enduring Beliefs'. *Canadian Journal of Philosophy*, 40, 107–29.
- Kelly, Thomas. (2003) 'Epistemic Rationality as Instrumental Rationality: A Critique'. *Philosophy and Phenomenological Research*, 66, 612–40.
- Konek, Jason, and Benjamin A. Levinstein. (2019) 'The Foundations of Epistemic Decision Theory'. *Mind*, 128, 69–107.
- Kvanvig, Jonathan L. (2003) *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Laudan, Larry. (1984) *Science and Values: The Aims of Science and Their Role in Scientific Debate*. Berkeley: University of California Press.
- Leitgeb, Hannes, and Richard Pettigrew. (2010a) 'An Objective Justification of Bayesianism I: Measuring Inaccuracy'. *Philosophy of Science*, 77, 201–35.
- Leitgeb, Hannes, and Richard Pettigrew. (2010b) 'An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy'. *Philosophy of Science*, 77, 236–72.
- Littlejohn, Clayton. (2018) 'The Right in the Good: A Defense of Teleological Non-consequentialism'. In Kristoffer Ahlstrom-Vij and Jeffrey Dunn (eds.), *Epistemic Consequentialism* (Oxford: Oxford University Press), 23–47.
- Mackie, J. L. (1982) *The Miracle of Theism: Arguments for and against the Existence of God*. Oxford: Oxford University Press.
- McHugh, Conor. (2012) 'Epistemic Deontology and Voluntariness'. *Erkenntnis*, 77, 65–94.
- Mougin, Gregory and Elliott Sober. (1994) 'Betting Against Pascal's Wager'. *Noûs*, 28/3: 382–395.
- Muñoz, Daniel. (2021) 'The Rejection of Consequentializing'. *Journal of Philosophy*, 118, 79–96.

- Oppy, Graham. (1991) 'On Rescher on Pascal's Wager'. *International Journal for Philosophy of Religion*, 30: 159–68.
- Pettigrew, Richard. (2013) 'Epistemic Utility and Norms for Credences'. *Philosophy Compass*, 8, 897–908.
- Pettigrew, Richard. (2016) *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Plantinga, Alvin. (1993) *Warrant and Proper Function*. Oxford: Oxford University Press.
- Pojman, Lewis. (1986) 'Faith without Belief'. *Faith and Philosophy*, 3, 157–76.
- Portmore, Douglas W. (2009) 'Consequentializing'. *Philosophy Compass*, 4, 329–47.
- Pritchard, Duncan. (2010) 'Understanding'. In Duncan Pritchard, Alan Millar, and Adrian Haddock (eds.), *The Nature and Value of Knowledge: Three Investigations*. (Oxford: Oxford University Press), 66–88.
- Rota, Michael. (2016a) 'A Better Version of Pascal's Wager'. *American Catholic Philosophical Quarterly*, 90, 415–39.
- Rota, Michael. (2016b) *Taking Pascal's Wager: Faith, Evidence, and the Abundant Life*. Downers Grove: InterVarsity Press.
- Schroeder, Mark. (2012) 'Stakes, Withholding, and Pragmatic Encroachment on Knowledge'. *Philosophical Studies*, 160, 265–85.
- Simion, Mona. (2018) 'No Epistemic Norm for Action'. *American Philosophical Quarterly*, 55, 231–38.
- Singer, Daniel. (2018) 'How to Be an Epistemic Consequentialist'. *Philosophical Quarterly*, 68, 580–602.
- Singer, Daniel. (2019) 'Demoting Promoting Objections to Epistemic Consequentialism'. *Philosophical Issues*, 29, 268–80.
- Singer, Daniel, and Sara Aronowitz. (2022) 'What Epistemic Reasons Are For: Against the Belief-Sandwich Distinction'. In Billy Dunaway and David Plunkett (eds.) *Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard* (Ann Arbor: Maize Books), 74–94.
- Solomon, Miriam. (1992) 'Scientific Rationality and Human Reasoning'. *Philosophy of Science*, 59, 439–55.
- Sosa, Ernest. (2011) *Knowing Full Well*. Princeton: Princeton University Press.
- Sylvan, Kurt. (2020a) 'Epistemic Non-consequentialism'. *Philosophical Review*, 129, 1–51.
- Sylvan, Kurt. (2020b) 'Reliabilism without Epistemic Consequentialism'. *Philosophical and Phenomenological Research*, 100, 525–55.
- Talbot, Brian (2019) 'Repugnant Accuracy'. *Noûs*, 53, 540–63.
- Treanor, Nick. (2013) 'The Measure of Knowledge'. *Noûs*, 47, 577–601.
- Turri, John, and Mark Alfano. (2021) 'Virtue Epistemology'. In Edward N. Zalta, (ed.) *The Stanford Encyclopedia of Philosophy*. October 26, 2021. <https://plato.stanford.edu/entries/epistemology-virtue/>.
- Wedgwood, Ralph. (2018) 'Epistemic Teleology: Synchronic and Diachronic'. In Kristoffer Ahlstrom-Vij and Jeffrey Dunn (eds.), *Epistemic Consequentialism* (Oxford: Oxford University Press), 85–112.
- Whitcomb, Dennis. (2018) 'Some Epistemic Roles for Curiosity'. In İlhan Inan, Lani Watson, Dennis Whitcomb, and Safiye Yigit, (eds.), *The Moral Psychology of Curiosity* (London: Rowman & Littlefield), 217–38.
- Wright, Dale S. (2016) *What is Buddhist Enlightenment?* Oxford: Oxford University Press.