Pascal's Wager and Decision-making with Imprecise Probabilities

André Neiva

(Federal University of Alagoas)

Forthcoming in *Philosophia*

Abstract

Unlike other classical arguments for the existence of God, Pascal's Wager provides a pragmatic rationale for theistic belief. Its most popular version says that it is rationally mandatory to choose a way of life that seeks to cultivate belief in God because this is the option of maximum expected utility. Despite its initial attractiveness, this long-standing argument has been subject to various criticisms by many philosophers. What is less discussed, however, is the rationality of this choice in situations where the decision-makers are confronted with greater uncertainty. In this paper, I examine the imprecise version of Pascal's Wager: those scenarios where an agent's credence that God exists is imprecise or vague rather than precise. After introducing some technical background on imprecise probabilities, I apply five different principles for decision-making to two cases of state uncertainty. In the final part of the paper, I argue that it is not rationally permitted to include zero as the lower probability of God's existence. Although the conditions for what makes an act uniquely optimal vary significantly across those principles, I also show how the option of wagering for God can defeat any mixed strategy under two distinct interpretations of salvation.

Decision Theory; Imprecise Probabilities; Mixed Strategies; Pascal's Wager.

1 Introduction

Pascal's Wager is a pragmatic argument for the existence of God. In its most standard form, which consists in a decision-theoretic approach, Pascal's argument leads us to conclude that rationality requires one to wager for the existence of God because this is the option that uniquely maximizes the expected utility.[1] If you bet on the existence of God and it turns out that God exists, then you get an infinite reward. Otherwise, if you bet against the existence of God when he exists, then you get either a negative infinite or finite utility value. You are only finitely rewarded for either betting on or against God when he does not exist. Given that your subjective probability for God's existence is some positive real number, your expected utility of betting on God will be infinite, making it much greater than your expected utility of betting against God, which will be at most equal to some positive finite number. Thus, assuming that such a decision matrix is right, the principle of expected utility maximization requires you to bet on God's existence.

As appealing as this might be, philosophers have challenged Pascal's argument in a variety of ways over the years. The argument also seems to depend on various crucial and subtle assumptions. In particular, it is widely held that you should assign a point-valued credence to the proposition that God exists. But maybe your confidence in God's existence is just far from being precise, even though it indicates some degree of uncertainty about whether this proposition is true. So what happens to Pascal's argument if we drop credal precision, permitting agents to have imprecise or unsharp credences instead? In what follows, I will be exploring Pascal's argument in the context of imprecise decision theory. More specifically, I will assess how different principles of practical rationality behave when it comes to a decision matrix in which it is rationally permissible to have an imprecise credence in God's existence. Before I delve into the application of those rules, I provide some technical background on the nuts and bolts of imprecise probabilities. Later, I distinguish between two kinds of scenarios for decision-making under uncertainty that may be relevant to Pascal's Wager. After applying those candidate rules for decision-making with imprecise probabilities to Pascal's Wager, I discuss whether rational agents

[1] This argument has been called *the argument from generalized expectations* or, more simply, *Pascal's Canonical Wager*. For more on different versions of Pascal's Wager, see Jordan (2006) and Hájek (2018*b*).

should include 0 as their lower probability of God's existence. Finally, I discuss how the mixed strategies objection makes trouble for Pascal's Wager with imprecise probabilities. I show that although Schlesinger's criterion cannot rescue the imprecise version of the Wager, there are two standard approaches to salvation—viz. orders of infinity and finite utilities—that can help to overcome this challenge (provided that certain conditions hold).

2 Pascal's Canonical Wager

As I have indicated, it has been widely accepted that Pascal's Canonical Wager is best represented in decision-theoretic terms. Given a partition $\Omega = \{G, \neg G\}$ of two possible states of the world, where G stands for the proposition that God exists and $\neg G$ stands for the proposition that he does not exist, one can wager either for or against G. For present purposes, W_G and $W_{\neg G}$ symbolize these two options at one's disposal at a given time. Following a common suggestion in the literature, wagering or betting on God amounts to adopting a continuing action that attempts to foster the belief in God, whereas wagering or betting against God involves not adopting a way of life where the agent tries to cultivate this belief.[2] While an agent who sticks to the former will follow a kind of life that is in accordance with a set of religious practices that might raise her confidence in G, a person who prefers the latter will not go along with this course of action, which means that she is not willing to make any great practical effort to increase her confidence in G.[3]

Moreover, here are some important assumptions of Pascal's Wager:

The possible states of the world are independent of the acts and each act-state pair $\langle A_i, s_i \rangle$ determines a consequence or outcome;

$\forall r \in \mathbb{R} \quad -\infty < r < +\infty$;

$\forall r \in \mathbb{R} \ \pm \infty + r = \pm\infty$;

[2] As other philosophers have already emphasized, this move avoids problems with doxastic voluntarism. After all, atheists and agnostics will be hardly capable of voluntarily believing in the existence of God, even if this is the prospect of maximum expected utility. How could rationality require a person to believe in God's existence if choosing this prospect appears to be impossible for him or her? For reasons like that, authors such as Hájek (2003), Jordan (2006), and Bartha (2018) all construe the Wager as providing acts that are under one's control. Because the set of acts available to the agent facing a decision problem typically forms a partition, she cannot but decide one way or the other.

[3] As the God of Pascal's Wager is the Christian God, I concentrate on a decision matrix with only this traditional deity in this paper. So I will not deal with a more fine-grained partition of possibilities, which extends Pascal's original matrix so as to include other theistic worldviews.

$\forall r \in \mathbb{R}_{>0} \ \pm\infty \times r = \pm\infty;$

$\infty - \infty$ has an indeterminate form.

In order to make sense of infinite utility, we might extend the real numbers by including $+\infty$ and $-\infty$ as members. By doing this, we get the affinely extended real number system. As suggested above, the usual arithmetic operations also apply to elements of the set $\mathbb{R} \cup \{-\infty, +\infty\}$. The operations above will be almost everything we need throughout this paper.[4] This approach is perhaps the more straightforward way of representing the outcome of salvation that, according to Pascal, results from the combination of the act of wagering for God and the state in which God does exist. However, it is worth mentioning that this strategy is not entirely uncontroversial.[5] The main problem is that infinite utility conflicts with a central feature imposed on the preference relation, known as *continuity*.[6] For now, let us assume some sort of naïve infinite decision theory.[7] Given those assumptions, we can state a standard version of Pascal's decision matrix:

	G	$\neg G$
W_G	∞	u_1
$W_{\neg G}$	u_2	u_3

Table 1: Pascal's 2×2 decision matrix.

In contrast with the outcome resulting from betting on an existent God $\langle W_G, G\rangle$, it is often assumed that u_1, u_2, and u_3 are all finite utilities. To be sure, there is some discussion about whether u_2 should be either a finite number or negative infinite, since there are passages in his work Pensées (1960 [1670]) where Pascal uses the word *misery* to refer

[4] I consider the particular case of multiplying 0 by ∞ in section 4.

[5] It is not totally clear what Pascal meant by "there is an infinity of infinitely happy life to win" when he refers to religious salvation (1960 [1670]: 343). As a matter of fact, there is no passage of his Pensées in which he uses the lemniscate symbol to refer to the utility associated with this outcome.

[6] Continuity says that if an agent $\mathcal{S}$ prefers A to B to C, then there is a gamble between A and C (where the probability of getting A is p and that of getting C is $1-p$) such that $\mathcal{S}$ is indifferent between accepting that gamble and B. Most expected utility theories require the preference relation to obey continuity. For instance, the classical representation theorem proposed by von Neumann and Morgenstern (1944) requires continuity, which (along with other requirements) guarantees that an agent's preferences are representable as maximizing expected utility. However, we cannot represent one's set of preferences using continuity if there is a consequence X whose value is infinitely better than another Y according to that set. It is worth noting that there are alternative, rigorous models of expected utility that accommodate the use of infinite utilities, most notably the multidimensional approach developed by Hausner (1954) and the relative utility theory advanced by Bartha (2007). See also McClennen (1994) for more information about the incompatibility between infinite utilities and classical decision theory.

[7] I borrow this term from Bartha & Pasternack (2018).

to the outcome associated with it, and others where he speaks of *hell*. To simplify our discussion, I assume that God will not inflict such an awful punishment (viz. $-\infty$) on those who prefer to bet against his existence. So even though agents who stick to wagering against God will not get an infinite reward, they will not be doomed to hell for eternity either. As a result, u_2 has some finite utility value.

It may be worth emphasizing that Pascal's Wager has been typically thought of as a decision problem under uncertainty. Unlike decisions under risk, in the context of decisions under uncertainty, it is up to the agent to assign probabilities to the different possible states of the world. Probabilities are not fixed beforehand.[8] Instead of objective probabilities, a distribution *pr* is a subjective probability function. We define it as a mapping from the states of the world to real numbers between 0 and 1 (inclusive), where the outputs denote one's various degrees of confidence in those states under consideration.[9]

Let $\mathcal{O}$ be a partition of acts (or options) available to an agent $\mathcal{S}$. It is rationally permissible for $\mathcal{S}$ to choose $A \in \mathcal{O}$ just in case A is the act that maximizes the expected utility, namely $EU(A) \geq EU(B)$ for any other $B \in \mathcal{O}$, where EU is an expected utility function.[10] When the expected utility of choosing A is uniquely maximal, the set of permissible acts consists of the singleton $\mathcal{O}^* = \{A\}$. Accordingly, $\mathcal{S}$ is rationally required to choose the option A, since $EU(A) > EU(B)$ for any other $B \in \mathcal{O}$.

If G's probability is positive, then the expected utility of W_G has an infinite value:

$$EU(W_G) = \infty \ . \ p \ + \ u_1 \ . \ (1-p) = \infty.$$

By contrast, the expected utility of $W_{\neg G}$ will receive a finite value only:

$$EU(W_{\neg G}) = u_2 \ . \ p \ + \ u_3 \ . \ (1-p) = \alpha, \text{ where } \alpha \text{ is some positive finite number.}$$

[8] It is common to say that lotteries are the objects of one's preferences in decisions under risk—for example, von Neumann and Morgenstern's framework (1944)—whereas acts are the objects of one's preferences in decisions under uncertainty—for instance, Savage's decision theory (1954). In the former case, the assumption is that the decision-maker knows the probabilities of all possible outcomes, so they are fixed or given in the decision problem. This differs from situations involving uncertainty, where the probabilities are unknown or unavailable. See Buchak (2013) for more information about this distinction.

[9] The type of uncertainty I am interested in is sometimes called *state uncertainty*, which is the most common case discussed in the literature. It has to do primarily with the decision maker's uncertainty about what possible state of the world turns out to be true. See Bradley & Drechsler (2014) for more details on both different kinds and levels of uncertainty.

[10] Recall that the expected utility of an act is the weighted sum of the utilities of the outcomes associated with that act, where the weights are given by a probability distribution. Formally speaking, $EU(A) = \sum_{i=1}^{n}[u(A, S_i) \times pr(S_i)]$, provided that S_i belongs to a partition consisting of possible states of the world and the probabilities of each S_i do not vary with the acts.

Therefore, W_G is the option that uniquely maximizes the expected utility:

$$EU(W_G) = \infty > EU(W_{\neg G}) = \alpha.$$

Consequently, wagering for G is strictly preferred to wagering against G. If this reasoning is correct, one is rationally required to choose W_G. Now we can formulate Pascal's traditional argument in defense of wagering for God in more rigorous terms:[11]

(P1) It is rationally required for $\mathcal{S}$ to choose the act that uniquely maximizes the expected utility (if there is one available to $\mathcal{S}$).

(P2) It is rationally required for $\mathcal{S}$ to assign a positive probability to G.

(P3) Pascal's decision matrix is correct: the act W_G uniquely maximizes the expected utility, that is, the set of permissible acts consists of $\mathcal{O}^* = \{W_G\}$.

$\therefore$ Rationality requires $\mathcal{S}$ to strictly prefer W_G over $W_{\neg G}$.

As compelling as this line of reasoning might look, Pascal's canonical argument depends upon various crucial assumptions. Closer to the matter at hand, it is widely assumed that rational agents should assign a fixed, point-valued subjective probability to the proposition that God exists. I want to relax this core assumption. The ultimate aim of this work is to scrutinize whether Pascal's Wager continues to justify the option of wagering for God as the rationally mandatory act when we turn to decision-making with imprecise probabilities. Note that the argument's second premise says it is not permissible to assign 0 probability to God's existence. At this point, one may reasonably argue that agnostics and even atheists cannot be rationally certain of G's falsehood. Perhaps it is rationally permissible to have a very low subjective probability in G or even assign any real number in the unit interval except the extreme values of 0 and 1 to G.[12] This issue will guide much of our discussion later on. Before engaging in a more extensive discussion about permissible probability distributions, I will spend a bit of time motivating why we should take Pascal's Wager as a decision problem under uncertainty in which the probabilities in question are actually unsharp rather than perfectly precise. Then, I will present the

[11] Here I am following Hájek's exposition (2018*b*) of the canonical argument.

[12] A further condition is that the distribution does not give an infinitesimal probability to G either.

standard approach for modeling one's confidence through imprecise probabilities. This will lead us to consider a group of candidate rules for decision-making in those specific sorts of cases, particularly when the degree of specificity of one's evidence seems to be lower than those situations involving precise probabilities. Later in this paper, I examine more closely how each of those rules behaves in the face of a mixed strategy that randomizes the decision between wagering for God and wagering against God.

3 Imprecise Probabilities

One of the most well-known controversies in the Bayesian camp is whether it is sometimes rationally permissible for an agent to adopt an imprecise or vague subjective probability. According to precise Bayesianism, one's entire doxastic state is best represented by a single probability distribution. This makes an agent's doxastic attitudes sharp. If rationality requires credal precision, then an agent should ascribe numerically definite credences to all propositions of her opinion set. On the other hand, imprecise Bayesians do not impose such a demanding constraint on a rational agent's attitudes, allowing her to have interval-valued credences in at least some propositions. For imprecise Bayesians, rational agents do not need to be opinionated about every proposition of their opinion set. Putting it differently, perhaps there are some sorts of occasions on which agents ought—or at least they are rationally permitted—to be non-committal to credal precision.[13]

Arguably, there are independent motivations for the claim that rationality allows one to have imprecise probabilities. I mention two. First, there may be cases of incomparability in one's doxastic attitudes. You may be more confident in P than $\neg P$ and more confident in Q than $\neg Q$, but you might fail to compare P and Q because either you regard them as incomparable or it has never occurred to you to compare your levels of confidence in P and Q. We cannot do justice to those kinds of judgments unless we abandon credal precision. Second, and more importantly, it seems that your evidence for P can be sufficiently unspecific to deserve a precise credence. Given that rationality demands responsiveness

[13]Some authors—like Levi (1974)—may prefer the expression "indeterminate probabilities" rather than "imprecise probabilities" for referring to the idea of having doxastic attitudes that fall short of credal precision. Be that as it may, I am taking the expression "imprecise probabilities" in a somewhat loose way to cover the whole family of approaches that represent an agent's doxastic attitudes using ranged credences or sets of distributions. I define imprecise probabilities more accurately in section 3.1.

to evidence, having imprecise probabilities may be considered the appropriate, rational response in a large number of cases. After all, what precisely should be your credence in the proposition that it will snow in West Lafayette on New Year's Day 2080? How about your credence that the next toss of a particular coin of unknown bias will come up tails? Presumably, the evidence you currently possess is indeed quite unspecific about such matters. If so, having precise credences does not appear to fit well in those and other similar cases. More generally, we may think that when an agent has sparse or unspecific evidence about P, it seems right for her to respect such an impoverished evidential condition in which she finds herself by not assigning a precise credence to P. In fact, as Joyce (2005: 167) points out, the level of specificity of one's evidence can vary greatly depending on the case in hand: one's evidence can be less than totally specific about P's truth because it is either incomplete or ambiguous.[14] Compare a coin of unknown bias with a coin you know is fair. While you have plenty of evidence about the past outcomes that make you have a flat distribution in the latter scenario, you are completely in the dark regarding the fairness of the first coin. Once we appreciate plausible cases like these, it becomes more clear that sometimes it is rationally permissible to adopt imprecise credences towards some propositions.

Just as there may be ordinary circumstances that it is rationally permissible for $\mathcal{S}$ to have an imprecise probability in a certain proposition P, so too one may think of Pascal's Wager as a scenario in which $\mathcal{S}$ is faced with greater uncertainty. Some might go even further and claim that there are situations in which it is not only rationally permissible but also rationally mandatory to adopt an imprecise attitude towards P. For the purposes of this paper, all we need is to concede the permissibility of imprecise probabilities. Also, I am not suggesting that Pascal's Wager should be viewed as a typical example of a decision under massive uncertainty. However, it is not implausible to say that some people lack enough evidence to ascribe a sharp subjective probability to God's existence. Even if the extent of uncertainty may vary significantly across people, one might claim that credal precision should be the exception rather than the rule in this particular case. That is to say, some of us—including, most notably, agnostics and atheists—are usually more

[14] To make this distinction more vivid, Joyce says that a body of evidence E is incomplete whenever E does not help us to discriminate P from other incompatible propositions, and it is ambiguous whenever E's evidential impact on P changes across different reasonable construals of it.

confident in $\neg G$ than G, but it may be that our evidence is not so specific to require credal precision. Let us suppose this is correct, so rational agents can have unsharp probabilities that God exists. In that case, the question should be: would the pure act of wagering for God continue to be rationally mandatory? In order to answer this question, we should first see what kind of decision theory we can get from cases involving credal imprecision. Then, we should investigate if there is any rational constraint on imprecise probabilities so as to rule out particular interval-valued assignments as irrational.

Philosophers such as Hájek (2000), Duncan (2003), and Rinard (2018) have already discussed how Pascal's Wager fares when agents have imprecise rather than precise probabilities. In this paper, I want to bring new developments to the table. More specifically, I apply a set of fairly reasonable principles for decision-making under uncertainty to two distinct scenarios for the imprecise version of the Wager. Before we get to all that, let me introduce the formal machinery of imprecise probabilities into our discussion.

3.1 Modeling Imprecise Probabilities

The standard approach to imprecise probabilities represents an agent's confidence by appealing to a set of functions rather than just by a single probability distribution:[15]

$$\mathcal{P}r = \{pr : 0 \leq pr(P) \leq 1 \ \wedge \ P \in \mathcal{L}\}.$$

As van Fraassen (1980) suggests, the set $\mathcal{P}r$ is an agent's representor inasmuch as all functions in $\mathcal{P}r$ work together in order to describe her credal state. But which distributions are included in an agent's representor? Above all, it depends on what credal facts about one's doxastic state we want to represent. For instance, consider the case of an agnostic who is strictly more confident in $\neg G$ than G. We represent such a credal fact about his or her state by establishing that $pr(\neg G) > pr(G)$ for every $pr \in \mathcal{P}r$. Therefore, each distribution in the representor set will share this common property about his or her state.

Once we characterize $\mathcal{P}r$ as representing an agent's credal state, we define the upper

[15] Remember that a precise subjective probability distribution pr on a language $\mathcal{L}$ is a real-valued mapping from $\mathcal{L}$ to the unit interval, where $\mathcal{L}$ is a classical propositional language constructed out of a countable set of atomic propositions and the logical operators $\vee$ and $\neg$. As usual, a probability distribution satisfies non-negativity, normality, and finite additivity. In our case, each pr is defined on the same $\mathcal{L}$.

probability of P as the least upper bound of one's credence in P and the lower probability of P as the greatest lower bound of one's credence in P. Both the upper and lower probabilities are defined relative to the representor $\mathcal{P}r$.[16] The greater the distance between P's lower and upper probabilities, the more impreciseness there is in $\mathcal{S}$'s attitude towards P. In place of assigning one single value to $\mathcal{S}$'s credence in P, her attitude relative to P will form a numerical interval. For example, if the lower and upper probabilities of P are equal to a and b respectively, then the probabilities will range over $[a, b]$. As usually put, we say that an agent's credence in P corresponds to the interval $[a, b]$.

It is important to notice that the representor by itself does not guarantee the existence of credences whose outputs are intervals rather than points. As far as we have been able to determine, it is possible that $\mathcal{P}r$ includes only functions assigning a and b to P. And how about all other real numbers ranging from a to b? In order to ensure that $\mathcal{S}$'s credal state amounts to a genuine numerical interval, her representor must be convex. Besides ruling out discontinuities in an agent's imprecise attitudes, convexity implies that any mixture of functions in $\mathcal{P}r$ is included in this set as well.[17] Here I will focus on the simplest case—namely, Pascal's 2×2 matrix—where the relevant domain consists of $\{G, \neg G\}$.

3.2 Decision Theory

Decision theory with precise probabilities is relatively straightforward. However, things get a little more tricky and unclear when it comes to imprecise probabilities. For starters, every distribution belonging to the representor yields some expected utility for each option under consideration. Applying it to our case in hand, if those distributions assign different probability values to $\neg G$, then the expected utilities of $W_{\neg G}$ may differ from each other as well. In general, the expected utility of an act will no longer be a particular number or value. Since we now have a representor—and, thereby, interval-valued credences generated by it—we will most of the time not get merely one expectation for each option, but an

[16] Slightly abusing the notation, $\mathcal{P}r^*(P) = \sup\{\zeta : \zeta \in \mathcal{P}r(P)\}$ and $\mathcal{P}r_*(P) = \inf\{\zeta : \zeta \in \mathcal{P}r(P)\}$ are the upper and lower probabilities of P, respectively. If P and R are mutually exclusive, then $\mathcal{P}r^*(P \vee R) \leq \mathcal{P}r^*(P) + \mathcal{P}r^*(R)$ (subadditivity) and $\mathcal{P}r_*(P \vee R) \geq \mathcal{P}r_*(P) + \mathcal{P}r_*(R)$ (superadditivity). The upper and lower probabilities are also connected by $\mathcal{P}r^*(P) = 1 - \mathcal{P}r_*(\neg P)$.

[17] For any two distributions $pr, pr' \in \mathcal{P}r$, the representor $\mathcal{P}r$ is said to be convex just in case $\lambda pr + (1-\lambda)pr'$ also belongs to $\mathcal{P}r$, where $\lambda \in (0, 1)$.

interval of expected values instead.[18] Therefore, we cannot simply apply the principle of expected utility maximization to find the optimal decision as before.

It seems at first sight that imprecise probabilities bring complications for those who are interested in decision-making. They reflect significantly on the expected utilities. Considering that we have a set or a numerical range of expected values, how can we assess the merits of the acts available to the agent? How do we compare those different expectations? More importantly, what rule does govern a rational agent's choice in those situations? Philosophers and decision theorists have been given distinct answers to these questions. In fact, there has been a lot of disagreement in the literature about how best to choose among many incompatible options when one's probabilities are imprecise. Although I will not develop a comprehensive discussion around the foundations and challenges to imprecise decision theory, which would lead us far afield, I think there are a few good candidates that could fit the bill. I will not argue for any of those principles as a requirement of practical rationality, though I do think that each constitutes a natural extension of the expected utility theory.[19] In this respect, Troffaes (2007) explores a number of different alternatives for decision-making with imprecise probabilities. I implement those criteria in Pascal's Wager. My aim is to examine whether those putative rules vindicate the option of wagering for God under several circumstances.

Before proceeding, some clarifications are in order. For ease of exposition, I formulate the rules in terms of rational permissibility, using the symbol $\geq$ to identify when a particular act is permissible. By and large, I will be using $\geq$ to compare different expected utilities. In addition to this variety of choice-worthiness, there are cases where a certain option A is strictly preferred to another option B. As these conditions vary across the principles under the current examination, I will point out their differences when I address the imprecise

[18] I say "most of the time" because this is what happens to cases involving interval probability assignments and finite utilities. As we will see, introducing infinite utilities into a decision problem deeply affects the calculation of the expectations in a significant way.

[19] One of the main challenges to the imprecise approach is to provide a response to what is known as the problem of sequential choices, which consists largely in a set of two bets that leaves the agent with a guaranteed profit if she decides to accept each of those bets and nothing if she chooses to reject both of them. The question is whether imprecise models can give the verdict that it is rationally impermissible for the agent to reject both bets. There is a vast literature dedicated to this problem, originally posed by Elga (2010, 2012) in the philosophical community, yet a similar challenge traces back at least to the work of Hammond (1988). Chandler (2014) discusses extensively the nature of Elga's claims and how the rules for decision-making with unsharp probabilities I apply to Pascal's Wager behave in such scenarios of sequential decisions.

version of Pascal's Wager. It should also be noted that the multiple distributions in a representor, taken together, generate the range of expected utilities. Since there may be an interval of expected values for each act, the divergence among the principles boils down to how we compare those values using the distributions in $\mathcal{P}r$.

The first suggestion is a version of the MaxiMin rule.[20] The classical MaxiMin criterion was originally proposed to cope with decision problems under ignorance, namely those situations in which it does not make sense to have any probabilities whatsoever. Since we are operating with imprecise probabilities, this principle takes a second-order form: the MaxiMin rule for expected utilities. It says that one should look at the minimum expected utility for each act and prefer the option with the maximum expected utility among those worst expected values. Simply put, the agent should maximize the minimum expectation. I characterize this rule as follows, where the subscript *pr* in the notation means "by the lights of *pr*" or "according to *pr*":

($\mathcal{P}r-$MaxiMin)
An act A is rationally permissible just in case $\min\{EU_{pr}(A) : pr \in \mathcal{P}r\} \geq \min\{EU_{pr}(B) : pr \in \mathcal{P}r\}$ for any other act $B \in \mathcal{O}$.[21]

If this is a conservative rule, the next principle goes in exactly the opposite direction. Just as MaxiMin, the MaxiMax rule has been considered a criterion for making decisions under ignorance. In our case, nonetheless, the MaxiMax rule for expected utilities states that the optimal decision is one whose upper expected value is the greatest. On this view, the agent should strive for the best-case optimization:

($\mathcal{P}r-$MaxiMax)
An act A is rationally permissible just in case $\max\{EU_{pr}(A) : pr \in \mathcal{P}r\} \geq \max\{EU_{pr}(B) : pr \in \mathcal{P}r\}$ for any other act $B \in \mathcal{O}$.

The MaxiMin and the MaxiMax rules locate the optimal decision at two different extremes of the spectrum, the most pessimistic and most optimistic scenarios. For those who think

[20] A helpful overview of the principles for decision-making explored here can be found in Troffaes (2007). See also Weatherson (1998). Chandler (2014) uses a conditional notation for the expectations, but we can suppress it because the acts and states are taken to be independent of each other in Pascal's Wager.

[21] Versions of the MaxiMin principle for expected utilities have been widely discussed in the literature: see, for instance, Gilboa & Schmeidler (1993*a*, 1993*b*) and Gärdenfors & Sahlin (1982).

that they are way too restrictive, there are two permissive rules that do not evaluate the choice-worthiness of the acts by inspecting either the maximum or minimum expected utilities. If you are looking for a more lenient principle, here's one plausible suggestion:

($E-$Admissibility)

An act A is rationally permissible just in case there is a $pr \in \mathcal{P}r$ such that $EU_{pr}(A) \geq EU_{pr}(B)$ for any other act $B \in \mathcal{O}$.

A slightly weaker principle arises if we switch the order of the quantifiers:

(Maximality)

An act A is rationally permissible just in case, for any other act $B \in \mathcal{O}$, there is a $pr \in \mathcal{P}r$ such that $EU_{pr}(A) \geq EU_{pr}(B)$.

As Troffaes (2007: 24) correctly identifies, both $E-$Admissibility and Maximality give precisely the same results whenever there are only two possible acts within the partition under consideration. In other words, those rules agree on the set of rationally permissible choices if $|\mathcal{O}| = 2$. Despite this caveat, the latter usually ranks more permissible acts than the former in cases involving more than two possible options.

The last principle that I will consider invokes the minimum and maximum expectations as the first two rules did, though in a different way. More carefully, it tells the agent to take the greatest expected utility of the target option and to compare this particular value with the lowest expected utilities of the other options available to her. If the former expected value is greater or equal to the latter ones, then the agent is permitted to choose that specific act. It is not hard to see that such a principle is the weakest of the lot. As it turns out, Interval Dominance includes more elements in the set of permissible acts than each of the previous four principles for decision-making:

(Interval Dominance)

An act A is rationally permissible just in case $\max\{EU_{pr}(A) : pr \in \mathcal{P}r\} \geq \min\{EU_{pr}(B) : pr \in \mathcal{P}r\}$ for any other act $B \in \mathcal{O}$.[22]

[22] Levi (1974, 1980) was the most famous advocate of $E-$Admissibility, whereas Seidenfeld (2004) provides a detailed comparison between $E-$Admissibility and $\mathcal{P}r-$MaxiMin. As far as I can tell, there is no explicit defense of $\mathcal{P}r-$MaxiMax, Maximality, and Interval Dominance among philosophers and decision theorists. It should be noted that Troffaes (2007) explores the properties and differences between those five principles. According to him, which criterion is more appropriate in situations of severe uncertainty depends ultimately on the decision maker's goals.

Though those rules are certainly not beyond dispute—and you might have your own favorite one—they constitute a set of alternative extensions of the traditional expected utility theory, normatively construed. I want to put any controversy about the foundations of imprecise decision theory aside and apply those candidate principles to the Wager. With them in mind, let us see whether Pascal's argument continues to succeed.

4 Pascal's Wager and Imprecise Probabilities

As noted already, it seems that people facing a decision like that of Pascal's Wager find themselves in a situation of uncertainty. Moreover, perhaps the level of uncertainty they are exposed to may be enough to prevent them from having a perfectly precise credence in the proposition that there is a God: their evidence might be sufficiently unspecific to allow for multiple permissible assignments to G. Assuming that the latter claim is correct, what exactly is the extent of an agent's uncertainty in this context? Should we put rational constraints on her lower and upper probabilities? Does a rational agent's representor rule out certain distributions as unreasonable? Before I consider the epistemic side of having imprecise credences, I show in some detail what results follow from applying those putative rules for decision-making to Pascal's decision matrix.

Let us now turn to Pascal's matrix with imprecise probabilities. To distinguish it from the original matrix, I have placed the interval-valued probabilities together with the corresponding outcomes in each square. I will work mostly with closed intervals so that the endpoints a and b of one's imprecise credence in G are also in $\mathcal{P}r$. Because of convexity, $\mathcal{P}r$ assigns a numerical range to G.[23] Here's the new version of Pascal's matrix:

	G	$\neg G$
W_G	∞ $[a,b]$	u_1 $[1-b,1-a]$
$W_{\neg G}$	u_2 $[a,b]$	u_3 $[1-b,1-a]$

Table 2: Pascal's 2×2 decision matrix with IP.

Having established that, I discriminate between two main types of scenarios in which the decision-maker is confronted with greater uncertainty about the possible states of the

[23]The same holds for $\neg G$. Recall from section 3.1 (see footnote 16) that $\mathcal{P}r^*(G) = 1 - \mathcal{P}r_*(\neg G)$ and $\mathcal{P}r_*(G) = 1 - \mathcal{P}r^*(\neg G)$. This explains the range of values assigned to $\neg G$.

world. In the first kind, the agent has probabilities ranging from a to b, but they are not extreme values. Hence, the representor assigns neither 0 nor 1 to G:

(**Uncertainty 1**)

$\mathcal{S}$'s credence in G ranges over $[a, b]$, provided that $a \neq 0$ and $b \neq 1$.

Putting it more formally, there is no $pr \in \mathcal{P}r$ such that $pr(G) = 0$ or $pr(G) = 1$. The greater the range's size, the more imprecise is an agent's attitude towards the proposition that God exists. Note that the size of $[a, b]$ depends on the representor. More specifically, the upper and lower probabilities of G in $\mathcal{P}r$ determine how large the interval is.

In the second scenario, the representor is even more permissive, containing at least one distribution that gives G an extreme probability. For the sake of argument, let us suppose this time that 0 is placed at the bottom of our numerical range:

(**Uncertainty 2**)

$\mathcal{S}$'s credence in G ranges over $[0, b]$, provided that $b \leq 1$.

Contrary to the first uncertainty-type case, the second scenario implies that:

$$\exists pr \in \mathcal{P}r \text{ such that } pr(G) = 0.$$

If $b = 1$, then we get the widest possible range of probabilities. Were it rationally permissible for the agent to have these distributions in her representor, her credence in G would be vague over the whole unit interval. Nevertheless, maybe rational agents are not permitted to have a credence ranging from 0 to 1 (inclusive). There seem to be good reasons for rejecting this sort of radical imprecision. I will be concerned with this issue soon. For the time being, I focus predominantly on the case where $b < 1$. Even though this stipulation simplifies things a little bit, it has no bearing on the results I present now.

We can narrow different scenarios of *state uncertainty* down to only the two above. Those are particular situations where the agent has some information about the possible states of the world to make at least an imprecise judgment on how probable they are. With this in place, we are able to make use of our set of candidate rules for decision-making with imprecise credences. Let us take the first uncertainty-type case. If $\mathcal{S}$'s credence in God's existence is vague over the range $[a, b]$—granted that $a \neq 0$ and $b \neq 1$—then the set of expected utilities of wagering for God will form a singleton, which is $\{\infty\}$. Therefore,

$$EU_{pr}(W_G) = \infty \text{ for every } pr \in \mathcal{P}r.$$

On the other hand, the expectation for $W_{\neg G}$ will range over a numerical interval whose minimum and maximum expected values are both finite numbers. If $u_2 > u_3$, then:

$$[u_2 \,.\, a \;+\; u_3 \,.\, (1-a),\; u_2 \,.\, b \;+\; u_3 \,.\, (1-b)].$$

Otherwise, if $u_3 > u_2$, then the expectation for $W_{\neg G}$ will range over:

$$[u_2 \,.\, b \;+\; u_3 \,.\, (1-b),\; u_2 \,.\, a \;+\; u_3 \,.\, (1-a)].$$ [24]

Whatever imprecise rule one chooses here, it is clear that Pascal's argument in defense of wagering for God works equally well in our first scenario. Both MaxiMin and MaxiMax vindicate the act of wagering for God as rationally mandatory. The strict inequality ensures such a conclusion because these two principles compare particular expectations for W_G and $W_{\neg G}$ by the lights of a particular pr in the representor:

($\mathcal{P}r-$MaxiMin)
$\min\{EU_{pr}(W_G) : pr \in \mathcal{P}r\} > \min\{EU_{pr}(W_{\neg G}) : pr \in \mathcal{P}r\}$.

($\mathcal{P}r-$MaxiMax)
$\max\{EU_{pr}(W_G) : pr \in \mathcal{P}r\} > \max\{EU_{pr}(W_{\neg G}) : pr \in \mathcal{P}r\}$.

$E-$Admissibility and Maximality yield the following results:[25]

($E-$Admissibility)
$\exists pr \in \mathcal{P}r$ such that $EU_{pr}(W_G) \geq EU_{pr}(X)$ for any other act $X \in \mathcal{O}$.

(Maximality)
For any other act $X \in \mathcal{O}$, $\exists pr \in \mathcal{P}r$ such that $EU_{pr}(W_G) \geq EU_{pr}(X)$.

[24] To see why the upper and lower bounds reverse the order whenever $u_3 > u_2$, assume that $\mathcal{S}$'s credence in G is vague over $[0.7, 0.9]$, that is, $a = 0.7$ and $b = 0.9$. If $u_3 = 2$ and $u_2 = 1$, we get the following: $(1 \times 0.7 + 2 \times 0.3) = 1.3$ and $(1 \times 0.9 + 2 \times 0.1) = 1.1$. The expectation for $W_{\neg G}$ will range over $[1.1, 1.3]$.

[25] There are only two acts available to $\mathcal{S}$, that is, $X = W_{\neg G}$.

What guarantees that W_G is strictly preferred to $W_{\neg G}$ in the first scenario of uncertainty? For these two rules, it is the fact that there is no pr in the representor according to which the option of wagering against God's existence is rationally permissible:

$$\neg\exists pr \in \mathcal{P}r \text{ such that } EU_{pr}(W_{\neg G}) \geq EU_{pr}(W_G).$$[26]

Lastly, one can reach the same conclusion if one makes use of the most lenient rule:

(Interval Dominance)
$\max\{EU_{pr}(W_G) : pr \in \mathcal{P}r\} > \min\{EU_{pr}(W_{\neg G}) : pr \in \mathcal{P}r\}$,
$\max\{EU_{pr}(W_{\neg G}) : pr \in \mathcal{P}r\} < \min\{EU_{pr}(W_G) : pr \in \mathcal{P}r\}$.

According to those principles, the option of wagering for God is not only rationally permissible but also rationally mandatory. As before, all rules we have been considering here require the agent to choose W_G largely because $W_{\neg G}$ is rationally impermissible for those rules. The former act is strictly preferred to the latter one. The very same feature that has played a significant role in Pascal's original argument for W_G comes into play again. It is the prospective prize of salvation as having infinite utility. As Hájek identifies (2018*a*: 131), the swamping effect of ∞ is the major factor responsible for that remarkable accomplishment.[27] Our first scenario of state uncertainty is not free from this influence.

Unlike the previous situation, the second sort of scenario involves attaching 0 as the lower probability of G. Recall that we have been working under the assumption of the extended real number line, which basically enlarges $\mathbb{R}$ by adding $-\infty$ and $+\infty$ as elements of this set. As I said in section 2, this is probably the more direct way of making sense of Pascal's idea of salvation. Furthermore, it is commonly held that $0 \times \infty = 0$ and $0 \times -\infty = 0$ in areas such as probability and measure theory. Otherwise, the minimum expectation for W_G would be undefined, and we would get undecidable results. Such standard definition avoids complications, yet I recognize that this might be a contentious issue. Taking this assumption for granted, if $\mathcal{S}$'s credence in G is vague over $[0, b]$, where $b < 1$, then the expected utilities of wagering for God will consist of the following set:

[26] Alternatively, $EU_{pr}(W_G) = \infty > EU_{pr}(W_{\neg G})$ for every $pr \in \mathcal{P}r$.

[27] As we shall see, this benign consequence of introducing ∞ turns out to be especially problematic if one takes into account the possibility of choosing W_G based on a mixed strategy.

$\{u_1, \infty\}$. Despite the fact that the expectations for wagering against God do not remain unchanged through our two uncertainty-type scenarios, they still form a numerical range.

If $u_2 > u_3$, then the following holds:

$$[u_3,\ u_2\ .\ b\ +\ u_3\ .\ (1-b)].$$

It is worth reemphasizing, nonetheless, that the range of expectations for $W_{\neg G}$ will change if $u_3 > u_2$, making u_3 the greatest value of that interval:

$$[u_2\ .\ b\ +\ u_3\ .\ (1-b),\ u_3].$$

The value of u_3 corresponds to the putative gain associated with the outcome of wagering against God in the state where he does not exist. Since the lower probability of G is fixed at 0, the minimum (or maximum) expected utility of $W_{\neg G}$ amounts to u_3.

Pascal's argument goes through according to the MaxiMax rule, where $b \in (0, 1)$:[28]

($\mathcal{P}r-$MaxiMax)

$\max\{EU_{pr}(W_G) : pr \in \mathcal{P}r\} > \max\{EU_{pr}(W_{\neg G}) : pr \in \mathcal{P}r\}$.

$\infty >\ u_2\ .\ b\ +\ u_3\ .\ (1-b)$ or $\infty > u_3$.

W_G ✓

$W_{\neg G}$ X

So the act of wagering for God is rationally permissible, whereas wagering against God is impermissible. For that reason, this principle vindicates Pascal's argument for W_G in both scenarios, Uncertainty 1 and Uncertainty 2. Notice that taking the second scenario for granted has a decisive consequence. The second premise of the Wager would no longer require that one should assign only positive probabilities to G. Thus, if we assumed that 0 is the lower probability of G, then MaxiMax would tell us to strictly prefer W_G to $W_{\neg G}$.

Yet this does not necessarily hold for the other principles we have been exploring. To be more precise, it is true that $E-$Admissibility, Maximality, and Interval Dominance all agree that the act of betting on God is (at least) rationally permissible. Nevertheless,

[28] I am adopting the symbols ✓ and X for rational permissibility and impermissibility, respectively. When it is unclear whether an act is permissible or impermissible, I use the symbol ? to indicate it.

it is far from being clear what precisely these rules say about the act of betting against God. I begin by highlighting the results we get from $E-$Admissibility and Maximality:

($E-$Admissibility)

$\exists pr \in \mathcal{P}r \; EU_{pr}(W_G) \geq EU_{pr}(W_{\neg G})$.

W_G ✓

$\exists pr \in \mathcal{P}r \; EU_{pr}(W_{\neg G}) \geq EU_{pr}(W_G)$.

$W_{\neg G}$ *?*

(Maximality)

$\forall X \in \mathcal{O}$, $\exists pr \in \mathcal{P}r$ such that $EU_{pr}(W_G) \geq EU_{pr}(X)$.

W_G ✓

$\forall Y \in \mathcal{O}$, $\exists pr \in \mathcal{P}r$ such that $EU_{pr}(W_{\neg G}) \geq EU_{pr}(Y)$.[29]

$W_{\neg G}$ *?*

As I stressed above, the results one gets from these two principles above overlap whenever the partition of acts consists of two elements. They count the act of wagering for God as rationally permissible because $EU_{pr}(W_G) = \infty$ for any $pr \in \mathcal{P}r$ such that $pr(G) \neq 0$. But it is open whether they return the same verdict with respect to the act of wagering against God. For all we know, it is unclear whether the conditions under which $W_{\neg G}$ is rationally permissible according to those rules are met. In order to know whether this option is also permissible, we have to compare u_1 and u_3. These values are the expectations produced by some pr in $\mathcal{P}r$ that assigns 0 to G, namely $EU_{pr}(W_G) = u_1$ and $EU_{pr}(W_{\neg G}) = u_3$.[30] So, for both rules, only a distribution pr that assigns 0 to G could potentially allow for $W_{\neg G}$ to meet the permissibility conditions:

The act of wagering against God is rationally permissible iff $u_3 \geq u_1$.

The question at issue seems to be this: does $W_{\neg G}$ bring you more utility (or more happiness) than W_G when there is no God? To my mind, the answer to this question needs to be empirically informed by experimental findings. Although I will not pursue it here, the key point is clear: the advocates of those rules who want to make a case for the permissi-

[29]Since the set of acts $\mathcal{O}$ consists of only two options, $X = W_{\neg G}$ and $Y = W_G$.

[30]That is, $EU_{pr}(W_G) = \infty \;.\; 0 \;+\; u_1 \;.\; 1$ and $EU_{pr}(W_{\neg G}) = u_2 \;.\; 0 \;+\; u_3 \;.\; 1$ according to $pr(G) = 0$.

bility of $W_{\neg G}$ must compare these values.[31] In any case, a move like that would represent a significant departure from Pascal's original Wager: his argument does not depend to any extent on a comparison among the outcomes associated with u_1, u_2, and u_3!

Let us see what we get from applying our last rule to Uncertainty 2. Firstly, it is clear that the extreme values of expectation of each option will settle the matter. As noted earlier, W_G is rationally permissible according to Interval Dominance, since $\max\{EU_{pr}(W_G) : pr \in \mathcal{P}r\} = \infty$. How about betting against God? Is this act rationally permissible as well? In this particular case, one has to draw a comparison between the maximum expected utility of $W_{\neg G}$ and the minimum expected utility of W_G. It will be relevant whether u_3 is greater than u_2 or the other way around:

If $u_2 > u_3$, $W_{\neg G}$ is rationally permissible iff $[u_2 \; . \; b \; + \; u_3 \; . \; (1-b)] \geq u_1$.

If $u_3 \geq u_2$, $W_{\neg G}$ is rationally permissible iff $u_3 \geq u_1$.[32]

Indeed, Interval Dominance is the least stringent rule of all those we have been exploring, and we generally expect it to deliver a larger set of permissible acts. But, even here, it is not obvious that rationality allows one to choose the act of adopting a secular way of life. The crux of the matter is that, besides the values of u_1 and u_3, the outcome associated with wagering against God when he exists will play a crucial role in this comparison. Once again, Pascal's original argument does not rely on comparing the finite utilities included in the matrix in order to succeed. However, in the second scenario of uncertainty, appealing to Interval Dominance would require us to know what precisely these particular utility values are. So here is the condition to be satisfied:

(Interval Dominance)

$\max\{EU_{pr}(W_{\neg G}) : pr \in \mathcal{P}r\} \geq \min\{EU_{pr}(W_G) : pr \in \mathcal{P}r\}$.

$W_{\neg G}$ *?*

[31] In his book *Pascal's Wager: Pragmatic Arguments and Belief in God* (2006), Jeff Jordan argues for what he calls *the Jamesian Wager*. Unlike Pascal's Wager, the Jamesian Wager is concerned only with the earthly benefits associated with wagering for God (or with theistic belief). In particular, the Jamesian Wager arises in cases with symmetrically balanced evidence for and against God's existence. One of the central claims of his book is that the Jamesian Wager leads to the conclusion that wagering for God brings more benefits than wagering against God, even in the case where God does not exist. I am indebted to an anonymous referee for drawing my attention to this.

[32] If the upper probability of G were 1 ($b = 1$), then u_2 would be the greatest expected utility of $W_{\neg G}$. So we would need to compare the outcomes of $\langle W_{\neg G}, G\rangle$ and $\langle W_G, \neg G\rangle$ (that is, u_2 and u_1).

Finally, we turn to the MaxiMin principle. We are equipped with information about the minimum expectations for W_G and $W_{\neg G}$: u_1 and u_3 or $u_2.b\ +\ u_3.(1-b)$, respectively. What we do not know, however, is whether adopting a religious life brings more happiness or utility than not adopting it when there is no God. Unless one provides a sound rationale that can be used for making those comparative judgments, proponents of the MaxiMin principle cannot say that betting on God is strictly preferred to betting against God in the second scenario (and vice versa). Summing up, according to the MaxiMin rule, there is no clear rational recommendation for the scenario posed by Uncertainty 2:[33]

($\mathcal{P}r-$MaxiMin)

$\min\{EU_{pr}(W_G) : pr \in \mathcal{P}r\} = u_1.$

$\min\{EU_{pr}(W_{\neg G}) : pr \in \mathcal{P}r\} = u_3$ or $u_2\ .\ b\ +\ u_3\ .\ (1-b).$

W_G *?*

$W_{\neg G}$ *?*

Let us take stock. In the first scenario of uncertainty, the option of wagering for God is the recommended option according to all rules for decision-making with imprecise probabilities we have been considering. Simply put, practical rationality requires one to choose W_G. Nevertheless, this is not true if we move to the second scenario of uncertainty. In order to achieve this remarkable result, one must endorse the MaxiMax principle. Ultimately, this is the only principle that unequivocally leads to Pascal's conclusion about what agents should do. As we have just seen, the MaxiMin principle alone does not tell us which of those two options, if any, is the unique optimal choice. While $E-$Admissibility, Maximality, and Interval Dominance all recommend wagering for God as a rational course of action, it is not altogether clear whether wagering against God has the same standing as the former.[34] Neither of these rules provides a more definite verdict on its permissibility.

[33]Duncan (2003) also considers this principle, but rejects it as inappropriate due to its somewhat counter-intuitive implications for decision-making. For more discussion around this rule, see Seidenfeld (2004).

[34]Duncan (2003) shows that the second-order versions of both the MiniMax Regret rule and Hurwicz criterion recommend the act of wagering for God as the optimal decision, even when one's credence in G is vague over $[0, b]$. By extending the supervaluationist interpretation of $\mathcal{P}r$ to decision-making, Rinard (2018) also argues in support of the imprecise version of Pascal's Wager.

5 Rational Constraints on Imprecise Probabilities

So much for the principles of practical rationality. At this stage of our discussion, some might argue that rational agents actually do not have a vague credence in God's existence that includes 0 as its lower probability. Perhaps rationality requires agents to avoid having extreme values as the endpoints of their imprecise credence in G. If this claim is true, there will be no distribution in a rational agent's representor that gives probability 0 to the proposition that God exists. Consequently, the second scenario suggested above—viz. Uncertainty 2—would be ruled out as inadmissible, and Pascal's argument in support of betting on God could work in the imprecise setting as well. As far as I am concerned, there are two independent reasons that one can use to justify this claim. The first strategy invokes the regularity principle, while the second move appeals to the Bayesian updating rule known as conditionalization. Let us take a closer look at these two approaches.

Regularity has been thought of as a rational constraint on an agent's precise credences:

> (**The Regularity Principle**)
>
> No logically contingent proposition has rational credence 0.

By extension, regularity would demand that there is no pr in a rational agent's representor assigning credence 0 to any contingent proposition $P \in \mathcal{L}$. Assuming that G is a logically contingent proposition—and, presumably, this claim is controversial by itself— regularity would not allow one's imprecise credence in G to include 0 as one of its endpoints. Therefore, there could not be some distribution pr in $\mathcal{P}r$ such that $pr(G) = 0$. Together, the axioms of probability calculus and the regularity principle entail the following constraints:

> Only tautologies of $\mathcal{L}$ receive rational credence 1.
>
> Only contradictions of $\mathcal{L}$ receive rational credence 0.

In recent work, Susanna Rinard (2018) has argued that logically contingent propositions cannot receive extreme probability values.[35] Given that it is rationally permitted to have a vague credence in a contingent proposition P, rationality precludes agents from having some pr in their representor such that $pr(P) = 1$ or $pr(P) = 0$. The same goes for G if

[35] See also Rinard (2013).

God's existence turns out to be a contingent matter. She thinks that no rational agent has a representor that includes some pr according to which $pr(G) = 0$. By the same token, a rational agent's representor will rule out distributions that assign probability 1 to G. It is easy to see that Pascal's argument will continue to lead to the conclusion that the act of wagering for God is rationally mandatory. Just as in the case of Rinard's supervaluationist decision principle, every rule we have been discussing here vindicates this act as the uniquely optimal choice when we turn (exclusively) to Uncertainty 1.

Rinard's claim is motivated by a handful of intuitive, ordinary examples. But, more fundamentally, there is a norm of rationality for imprecise probabilities underlying her line of reasoning that supports that view. For Rinard (2018: 285), we should constrain a rational agent's representor by the following requirement:

> If P is a contingent proposition, then $pr(\top) > pr(P) > pr(\bot)$ for any $pr \in \mathcal{P}r$, where $\bot$ is a contradiction and $\top$ is a tautology of $\mathcal{L}$.

Although Rinard does not explicitly defend regularity as a rational constraint on imprecise probabilities, the requirement above entails such a principle. Since the probability calculus requires one to assign credence 1 to all tautologies of $\mathcal{L}$, every contradiction of $\mathcal{L}$ will receive credence 0. Thus, every distribution in $\mathcal{P}r$ will individually assign a real number ε to a contingent proposition so that $1 > \varepsilon > 0$. It does not matter whether you are more confident in G than $\neg G$ or the other way around. If each distribution in your representor obeys Rinard's requirement, then your credence in G will range over a numerical interval that does not encompass 0 and 1 as its endpoints. Though I recognize the force of Rinard's points, regularity has some troubling features we should take into account.

To be sure, regularity comes in two different varieties: one that applies only to an agent's prior distributions and the other that constrains an agent's credences at any stage of her investigation, which includes her posterior distributions. As I see it, the first version of regularity presents an acceptable requirement for credences (either precise or imprecise), particularly when the domain under consideration is a finite probability space.[36] It tells us

[36] Difficulties with the first and restricted version of regularity arise when we deal with infinite probability spaces, and the possibilities in question are equiprobable. See Weisberg (2011) for an overview, where he calls the first variety Initial Regularity and the second Continuing Regularity. It should be noted that philosophers often define the domain of regularity as the entire set of logically possible worlds. An alternative proposal would be to take the set of epistemically (or doxastically) possible worlds as

to judge anything other than a contradiction as having some initial plausibility. Roughly speaking, it recommends that we should adopt an open-minded attitude towards almost everything at the beginning of an investigation. Trouble arises when one endorses a widespread form of regularity. Despite being *prima facie* attractive, it is well-known that the second version of regularity conflicts with a standard diachronic norm of rationality for credences: the conditionalization rule. Agents who abide by conditionalization assign credence 1 to a new piece of evidence E after learning E from experience between a given time t and a later time t'. Since $pr_{t'}(E) = 1$ so that $pr_{t'}(P) = pr_t(P \mid E)$ for any $P \in \mathcal{L}$, their credences in E go to the maximum value. For any logically contingent proposition E, if E receives credence 1, then $\neg E$ receives credence 0. Recall that the negation of any contingent proposition is itself a logically contingent proposition. For this reason, assigning credence 0 to $\neg E$ clearly contradicts the second construal of regularity. By similar reasoning, one can show that the imprecise variant of regularity conflicts with the extended version of the conditionalization rule for imprecise probabilities.

A better approach against having 0 as the lower probability of God's existence can be found in Duncan (2003). Among other things, Duncan suggests that assigning credence 0 to G counts as sheer dogmatism. If this is so, the second type of case we have been exploring—viz. Uncertainty 2—represents an utterly wrong attitude from an epistemic standpoint. As I have mentioned, conditionalization has been widely accepted as an epistemic norm for precise credences. It would be natural to expect that this rule can be equally applied to the imprecise case, which involves updating every single distribution in an agent's representor: for any $P, E \in \mathcal{L}$, if E is learned from experience between t and t', then $pr_{t'}(P) = pr_t(P \mid E)$ for every $pr \in \mathcal{P}r$. Notice that if 0 were the lower probability of P, then this value would never increase, no matter what E one might get. The problem is that such rigidity would make the agent incapable of updating her G's lower probability, committing her to a form of dogmatism regarding G.

There are two closely related points that I want to add to Duncan's contention. First, I am inclined to think that G is not the kind of proposition whose falsity one can be certain of because I do not see the notion of God understood as an omnipotent, omniscient, and

the domain of those principles. On that account, rationality would require the agent to have positive credence in P if P were possible relative to what she knows (or believes).

morally perfect being as implying something incoherent.[37] In line with Duncan's claim, it seems that even an atheist or an agnostic leaning towards atheism, if rational, should concede that there is some initial plausibility attached to G, no matter how small. To be more precise, it seems wrong to take an extreme attitude towards the plausibility of God's existence. In a similar fashion, rational agents could not be certain of G's truth, either. Notice that this constraint does not prevent an agent from changing her set of distributions whenever she gathers more evidence about G. In fact, it is plausible to expect that rational agents should be open to reconsidering their original position whenever she acquires new relevant evidence. And this is precisely what a scenario like Uncertainty 2 rules out by fiat. Second, holding fixed the lower probability of G at zero, the difference between the lower and upper probabilities of G could be even larger after updating by conditioning on new information. That is to say, if the upper probability of G increased, the interval size would be larger than before, making one more uncertain about G's truth. This interval could also shrink depending on what kind of evidence comes in. Of course, whether the interval shrinks or becomes larger will hinge on the level of specificity of one's new evidence relative to G. It is interesting to note, however, that an agent cannot assign any other value besides 0 to G if her imprecise credence in G goes down to a precise one. Because one's imprecise probability in G will always have 0 at the bottom of the interval even after subsequent updates over time, moving each distribution of one's representor closer and closer to a precise value means ultimately assigning 0 to G.[38] In case of arriving at a precise distribution, why should we eliminate any other value of the unit interval besides 0 from the very beginning? I see no good reason for doing that.

To sum up, I see motivation in Duncan's line of thought against the claim that it is rational to assign 0 to the lower probability of God's existence, while I would resist accepting the general version of regularity as a requirement for imprecise probabilities. Now, in the final part of this paper, I explore an important difficulty that threatens Pascal's Wager with imprecise probabilities even if we restrict the domain of admissible distributions to Uncertainty 1. The mixed strategies objection poses a potential challenge

[37] A number of philosophers have argued that the concept of God is incoherent or conceptually impossible. Although I do recognize that my claim is contentious, I lack the space to discuss it in detail here. See Kenny (1979), Martin (2000), and Lovering (2013, ch. 6) for relevant critical discussion.

[38] Jordan (2006: 137) makes a similar point against the permissibility of assigning $[0, b]$ to G.

to Pascal's Wager with imprecise probabilities, as with its canonical version.

6 The Mixed Strategies Objection

As briefly mentioned, Pascal's Canonical Wager has been criticized by many philosophers throughout the last century. More particularly, authors such as Duff (1986) and Hájek (2003, 2018*a*) contend that Pascal's original argument is, in truth, logically invalid: even if we accept all its premises (see section 2), the act of wagering for God is not rationally mandatory for agents who are facing a Pascalian decision problem. According to them, there is more than one option that enjoys maximal expectation, not only the pure strategy of wagering for God but also any mixed strategy that has both of those prospects of our simplest decision matrix (viz. W_G and $W_{\neg G}$) as possible outcomes.

As it stands, Pascal's Wager with imprecise probabilities is equally plagued with the mixed strategies objection. Every mixed strategy combining both W_G and $W_{\neg G}$ into a single act is rationally permissible for all five principles examined in this paper. For concreteness, let us suppose that instead of choosing (directly) either W_G or $W_{\neg G}$, you decide to flip a fair coin to determine which of these two acts you will perform:

> You choose W_G just in case the coin lands TAILS,
>
> You choose $W_{\neg G}$ just in case the coin lands HEADS.

To obtain the expectation of this entire option, we first multiply the probability of the coin landing on TAILS with the expected utility of W_G and the probability of the coin landing on HEADS with the expected utility given to $W_{\neg G}$. Then, we sum these products:

$$0.5 \times [\infty \times p + u_1 \times (1-p)] \quad + \\ 0.5 \times [u_2 \times p + u_3 \times (1-p)] = \infty.$$

In addition to the pure strategy of betting on God, choosing the mixed strategy above will also make you maximize the expected value as long as $p > 0$. As a consequence, in scenarios where the relevant probabilities are point-valued, the pure act of W_G is no longer the course of action (among all options) that uniquely maximizes the expectation. Adopting a mixed strategy, such as the one just presented, gives you an infinite reward.

With this result in hand, it is easy to see what follows from each of our rules for decision-making with imprecise probabilities when we allow one to go for a mixed strategy. Let us assume that 'MS' stands for a mixed strategy (such as that of tossing a fair coin). If an agent's credence in G ranges over $[a, b]$ (where $a \neq 0$), then the set of expected utilities of MS will be $\{\infty\}$. The expectation for MS is equal to ∞ for any $p \in [a, b]$. Given that MS $\in \mathcal{O}$, MaxiMin, MaxiMax, and Interval Dominance all count the option of choosing MS as rationally permissible (on the condition that X is either W_G or $W_{\neg G}$):[39]

$$\min\{EU_{pr}(\text{MS}) : pr \in \mathcal{P}r\} \geq \min\{EU_{pr}(X) : pr \in \mathcal{P}r\} \text{ for any other } X \in \mathcal{O},$$

$$\max\{EU_{pr}(\text{MS}) : pr \in \mathcal{P}r\} \geq \max\{EU_{pr}(X) : pr \in \mathcal{P}r\} \text{ for any other } X \in \mathcal{O},$$

$$\max\{EU_{pr}(\text{MS}) : pr \in \mathcal{P}r\} \geq \min\{EU_{pr}(X) : pr \in \mathcal{P}r\} \text{ for any other } X \in \mathcal{O}.$$

Two main general facts ensure that the act of choosing MS is also rationally permissible for the two remaining principles, E-Admissibility and Maximality:

$$\forall pr \in \mathcal{P}r \; EU_{pr}(\text{MS}) = EU_{pr}(W_G),$$

$$\forall pr \in \mathcal{P}r \; EU_{pr}(\text{MS}) > EU_{pr}(W_{\neg G}).$$

A potential reward of infinite utility has become somewhat of a double-edged sword. What at first sight seemed to favor the imprecise version of Pascal's Wager—specifically in scenarios covered by Uncertainty 1—now prevents the preference for the pure act of betting on God from being uniquely optimal. Accordingly, ∞'s swamping effect spreads wider. Any course of action (which can be either a pure or mixed strategy) that includes at least one outcome with an infinite value will be permissible for all rules considered here, as long as the range of probabilities does not include 0. More crucially, the trouble is that it can be generalized to any mixed strategy by which there is a chance of you adopting a religious way of life—and, ultimately, that you will come to believe in God—no matter how small it is. In a way, anything you might decide to do that has a chance of you coming to wager for God will be pragmatically justified according to each of our principles for decision-making under (severe) uncertainty. Any sort of randomization involving W_G and $W_{\neg G}$ will be capable of doing the job by virtue of the infinite utility in the matrix.

[39] Recall from section 4 that the expected utilities of $W_{\neg G}$ range over an interval whose upper and lower bounds are finite numbers in our two kinds of scenarios (Uncertainty 1 and 2). Just like with MS, the set of expected utilities of W_G consists of $\{\infty\}$ in Uncertainty 1.

In what follows, I examine three different responses to the mixed strategies objection. The first attempt employs Schlesinger's criterion as a way of breaking the tie between W_G and MS. Unfortunately, this principle does not work for agents whose probabilities are imprecise. The last two approaches—one that uses well-behaved infinities and the other that adopts finite utilities only—rely on alternative ways of interpreting the outcome of salvation. When combined with imprecise probabilities, they may successfully justify (under certain conditions) the pure act of wagering for God as rationally mandatory.

6.1 Schlesinger's Principle and Imprecise Probabilities

In response to this objection, one might appeal to Schlesinger's decision principle to try to break the tie between MS and W_G. Assume for one moment that G's probability is precise. Since there are different strategies that enjoy maximal expectation, Schlesinger (1994: 90) proposes the following: "the criterion for choosing the outcome to bet on is its probability." Here, the proponent of Pascal's Canonical Wager may invoke Schlesinger's criterion to argue that the probability of getting the best possible outcome—viz. salvation—increases if one chooses the pure strategy of W_G rather than any mixed strategy. Though non-theists might dispute this claim, there is a straightforward route for those interested in defending the traditional Wager. If there is more than one option that maximizes expected utility, one should choose the act that maximizes one's probability of getting the outcome of salvation. Presumably, the probability of getting salvation by choosing W_G will generally be greater than the probability of achieving salvation by adopting some MS.[40]

Considering now that it is rationally permissible for an agent to assign an imprecise probability to G, how can we put this criterion into effect? Arguably, imprecise probabilities reflect the uncertainty in an agent's evidence or, more precisely, the level of specificity of her evidence. As we have seen, the typical way of representing an agent's entire doxastic state whenever her evidence does not justify having a single distribution is by means of a set of distributions. At this point, the pressing question to be asked is: why should one favor the upper probability of G to break the tie? Why not go for the lower probability of G or, alternatively, the midpoint value of $[a, b]$? Unless one argues that agents should

[40]See Sorensen (1994) and Bartha (2007, 2018) for more discussion around Schlesinger's criterion in the context of precise probabilities.

choose a particular value of $[a, b]$—or, to be more precise, a certain $pr \in \mathcal{P}r$—one cannot compare it to the probability of achieving the best outcome via adopting a mixed strategy between W_G and $W_{\neg G}$. Besides, I have been assuming that the consequences of adopting a mixed strategy, such as tossing a fair coin as a way of determining what one has to do, have precise probabilities. But it is worth stressing that the problem will be more acute if we admit that those outcomes can harbor imprecise probabilities, too. If the probability of achieving salvation by adopting some mixed strategy were imprecise, we would need a principled criterion under which we could compare different sets of probabilities. Either way, providing a non-ad hoc tie-breaking procedure for such cases may be challenging.

Though Schlesinger's criterion provides no help to the imprecise version of the Wager, there are other attractive approaches that may save it. I will look into some of them now.

6.2 Orders of Infinity

So far, I have been assuming a naïve understanding of salvation. And, as I have said, Pascal himself suggested that God will reward the wagerer with "an infinitely happy life to gain." That is to say, Pascal believed that the outcome of salvation has infinite value. However, as already noted, he did not make use of the lemniscate symbol (viz. ∞) in any passage of his Pensées. So why should we stick to this naïve approach to salvation?

One reason is that such an approach is both simple and handy. It is probably the more direct way of construing Pascal's idea of salvation. However, it is important to point out that there are more rigorous formulations of infinite utilities that also allow us to describe this particular outcome. What is more, those approaches have some advantages over the naïve view. For one thing, they can discriminate between different orders or magnitudes of infinity. By contrast, the naïve view cannot accommodate the following property:

$$\infty > \infty \, . \, x, \text{ where } x \in (0, 1).$$

There are several proposals that satisfy this property, but for the sake of brevity I will focus on Hájek's suggestion (2003, 2018*a*) of representing salvation by a surreal infinite number. His exposition is based on Conway's (1976) surreal numbers. Using basically two rules for constructing distinct orders of infinity, Conway's system of numbers generates ordinals such as ω, $\omega - 1$, $\omega + 1$, $\frac{\omega}{2}$, 2ω, $\sqrt{\omega}$, ω^2, ω^ω, etc.[41] As we can clearly see, his system is appealing because it includes not only large orders of infinity but also smaller ones. Since it forms a totally ordered field, we can compare those numbers in size.

Applying Conway's surreal numbers to Pascal's Wager—particularly to the case where the relevant probabilities are precise—Hájek shows that the expected utility of W_G will be strictly greater than the expected utility of any given mixed strategy MS:

$$EU(W_G) > EU(\text{MS}).$$

We may replace ∞ with any infinite surreal number we want. By representing salvation as having some infinite surreal utility—viz. ω, $\omega + 1$ or any other one—the above inequality will hold because Conway's system is able to distinguish between different magnitudes of infinity. Unlike the naïve view, anything added to infinity will make a difference now. For instance, if we represent the utility of salvation with the first surreal number ω, then the fair coin-toss strategy will have a smaller infinite expected utility than wagering for God. Surreal numbers provide a successful response to the mixed strategies objection, as long as one assigns a positive and finite probability to God's existence.[42] In a word, this reconstruction of the traditional Wager vindicates W_G as the uniquely optimal choice.[43]

Back to our original point, what if one shifts to imprecise probabilities? What follows from those five principles for decision-making under severe uncertainty? Since I have ruled out Uncertainty 2 as irrational, $\mathcal{S}$'s credence in G will range over $[a, b]$, where $a \neq 0$. In this case, both MaxiMin and MaxiMax count W_G as rationally mandatory:

$$\min\{EU_{pr}(W_G) : pr \in \mathcal{P}r\} > \min\{EU_{pr}(\text{MS}) : pr \in \mathcal{P}r\},$$

$$\max\{EU_{pr}(W_G) : pr \in \mathcal{P}r\} > \max\{EU_{pr}(\text{MS}) : pr \in \mathcal{P}r\}.$$

[41] Conway first identifies each number with a form (a pair of sets of surreal numbers)—written as $\{L \mid R\}$—where every member of the right set R is strictly greater than any member of the left set L. The second rule says that, for any numbers a and b, $a \geq b$ just in case there is no member of a's right set that is less or equal to b and there is no member of b's left set that is strictly greater than a.

[42] That is, G's probability must be positive and finite, so this value cannot be an infinitesimal number.

[43] However, Hájek (2003: 45) raises a problem with using an infinite surreal number to represent salvation. Surreal numbers are not reflexive under addition: adding any positive number to, say, ω in fact increases this quantity. Thus, there exists a surreal utility that is greater than that associated with salvation. What is worse, a rival God could reward its followers with this higher surreal utility. If so, wagering for Pascal's Christian God would not necessarily be the optimal choice. Although extending the partition of theistic hypotheses might be a potential worry for the imprecise version of the Wager, I will not investigate this further here. Thanks to an anonymous referee for pointing out this problem.

Even if we admit more room for uncertainty, the smallest (greatest) expected utility of W_G will be strictly greater than that of MS. Because Conway's system gives rise to well-behaved infinities, it is true that those infinities have many sizes and, consequently, can generate (infinite) expected utilities of different magnitudes. The crux of the matter is that we will always have to add more numbers to infinity whenever we calculate the expected utilities associated with MS (compared to those of W_G): viz. $\omega > \omega \,.\, x$ for any $x \in (0, 1)$.

In general, the expected utility of any mixed strategy is less than that of wagering for God if salvation has a surreal utility and G's probability is positive and finite. Likewise, the following consequence holds for the case of imprecise probabilities:

$$\forall pr \in \mathcal{P}r \;\; EU_{pr}(W_G) > EU_{pr}(\text{MS}).$$

It is the condition by which E-Admissibility makes W_G rationally mandatory, considering that the interval of expectations for $W_{\neg G}$ harbors only finite values. There is a simple recipe for checking if this is the case. First, select a surreal number to represent the utility of salvation. Second, given a certain representor $\mathcal{P}r$, compare the expected utilities of W_G with those of MS by the lights of each $pr \in \mathcal{P}r$. If there is no $pr \in \mathcal{P}r$ according to which $EU_{pr}(\text{MS}) \geq EU_{pr}(W_G)$, then W_G will be once again the uniquely optimal choice.

To see how it works, let us assume that ω is the utility of salvation. A mixed strategy, such as the coin-toss strategy, has probability q for W_G and $(1-q)$ for $W_{\neg G}$, provided that $1 > q > 0$. By slightly abusing the notation, for every $pr \in \mathcal{P}r$, it follows that:[44]

$$\omega \,.\, pr \;+\; u_1 \,.(1 - pr) >$$
$$q \,.\, [\omega \,.\, pr \;+\; u_1 \,.(1 - pr)] \;+\; (1-q) \,.\, [u_2 \,.\, pr \;+\; u_3 \,.(1 - pr)],$$

which means that moving to surreal utilities justifies W_G as mandatory if we endorse E-Admissibility as our favored decision rule. This is similar to the way that, in the case of precise probabilities, EU maximization justifies W_G as the single permissible choice. It does not matter which pr is selected. As long as every pr is such that $pr(G)$ is neither 0

[44] Although I abuse the notation here, it is important to point out that $\mathcal{P}r^*(G) = 1 - \mathcal{P}r_*(\neg G)$ and $\mathcal{P}r_*(G) = 1 - \mathcal{P}r^*(\neg G)$ whenever one assigns imprecise probabilities to G and $\neg G$.

nor 1 (together with the conditions above), E-Admissibility will favor W_G over any MS.

Maximality says that MS is rationally permissible just in case, for any other option X, there is a $pr \in \mathcal{P}r$ such that the expectation for MS is greater than that of X by the lights of pr. By comparing the expectations for MS and $W_{\neg G}$, it is not hard to see that the values associated with the former are, by and large, greater than those relative to the latter. Mixed strategies will always have some infinite surreal number in the calculation, whereas only finite values figure in the outcomes associated with $W_{\neg G}$. Nevertheless, as just said, it is not the case that there is a distribution pr according to which $EU_{pr}(\text{MS}) \geq EU_{pr}(W_G)$, which would be necessary to make MS equally acceptable. Just as E-Admissibility, the inequality above prevents a mixed strategy from being a permissible choice for Maximality.

Now we turn to the weakest principle we have been examining. Theoretically speaking, it is possible that Interval Dominance vindicates some MS as a permissible option, which would amount to: $\max\{EU_{pr}(\text{MS}) : pr \in \mathcal{P}r\} \geq \min\{EU_{pr}(W_G) : pr \in \mathcal{P}r\}$. However, it is worth highlighting that there are three major factors upon which the truth of this inequality depends. The first and more obvious constraint is the size of $[a, b]$—namely, how distant the lower probability of G is from its upper probability. The second is q, the probability of ending up with W_G—and, obviously, $(1 - q)$. These probabilities reflect the kind of randomization involved between W_G and $W_{\neg G}$. More obviously, they have an impact on the calculation of the expectations for the mixed strategy itself. The last are the quantities associated with u_1, u_2, and u_3. Those factors (taken together) determine the set of expected utilities—which surreal numbers are included in it—of any given mixed strategy. We already know that if salvation is represented by an infinite surreal number, then the expected utility of W_G will be strictly greater than that of MS for each $pr \in \mathcal{P}r$. But since Interval Dominance is a very lenient principle, which compares the greatest expected utility of a given MS with the lowest expected utility of W_G, it remains open whether MS is impermissible until one is able to determine the values of those factors.

Allowing different orders of infinity, such as those from Conway's system, justifies the act of wagering for God as rationally mandatory for all imprecise rules except Interval Dominance, whose verdict remains unclear. At this point, one may reasonably claim that Interval Dominance is too permissive to be considered seriously as a plausible criterion for decision-making. After all, Interval Dominance is (by far) the criterion that delivers the largest set of permissible options.[45] Whether or not we accept it as a genuine candidate principle, one thing is still true. Depending on what values one gives to those factors above, Interval Dominance could make a mixed strategy rationally permissible as well.

6.3 Finite Utilities

Another popular response is to say that the idea of infinite utility typically used to represent the outcome of salvation is not even well-motivated in the first place. Maybe we as humans cannot have a glimpse of what would be an infinite amount of utility, much less a genuine experience of, as Pascal said, "an infinitely happy life." For instance, Duncan (2007, 2018) endorses using an arbitrarily large finite number to represent the outcome of salvation. For him, salvation has incomparable goodness, so all earthly goods fall short of this heavenly standard. Pushing this idea a bit further, what if salvation—understood as the best-case scenario—leads to a large but finite amount of happiness instead of ∞?

As already stressed, all the differences among our five criteria for decision-making vanish if we incorporate ∞ into the calculation of expected utilities, as long as the range of probabilities includes only positive and finite numbers. The first thing to emphasize, however, is that the conditions for what makes an option rationally permissible—and, ultimately, rationally mandatory—vary wildly across those criteria when we shift our focus to finite values of utility only. On the assumption that all outcomes are finitely valued, the permissibility of an option will hinge on two distinct factors: first, the differences between the utilities and, second, the size of the ranges of probabilities will matter.

Generally speaking, the utility associated with salvation must be sufficiently large in comparison to each u_i in the matrix, where $i \in \{1, 2, 3\}$, to vindicate Pascal's Wager in the imprecise setting. And what one exactly means by 'sufficiently large' depends on which principle is being applied to the decision problem. In addition, the size of the ranges of probabilities will have a huge impact on the results we get from each of those principles. For example, the larger the size of $[a, b]$, the larger the set of permissible acts for principles such as E-Admissibility, Maximality, and Interval Dominance. On the other hand, both MaxiMin and MaxiMax will usually include fewer elements in the set of permissible acts

[45]See Troffaes (2007) for more discussion on Interval Dominance.

than those three criteria if there is no outcome with infinite utility in the matrix.

Nevertheless, there is another crucial feature we should keep in mind when we examine the expectations for MS and W_G. It is the definition of the expected utility of a mixed strategy as a weighted sum of the expected utilities of both wagering for God and against God. Because a mixed strategy consists in randomizing an agent's choice, the weights are given by the probability of ending up with W_G and that of ending up with $W_{\neg G}$. As a result, the expected utilities of W_G will tend to be greater than those of a mixed strategy as long as the utility of salvation is sufficiently large, even when the interval $[a, b]$ has a considerable size. For any $p \in [a, b]$, there will be an expectation for the act of wagering for God and a smaller expectation for the mixed strategy. To have a concrete example, let us assume that salvation yields 20 utils and $\mathcal{S}$'s credence in G is vague over $[0.1, 0.7]$. If u_1 and u_3 has both 5 utils and u_2 has just 1 util, then the expected utilities of wagering for God will range over $[6.5, 15.5]$. The interval of expected utilities of a mixed strategy where there is a 50-50 chance of achieving salvation will be $[5.5, 8.85]$. Despite the fact that there is an overlap between these two intervals of expected utilities, the length of the former is greater than the latter. In other words, the distance between their upper boundaries is considerably larger than the area where they overlap.

Although it depends on those various factors, turning exclusively to finite utilities provides another interesting line of defense against the mixed strategies objection. In virtue of the definition of the expected utility of a mixed strategy, proponents of the Wager may argue in support of the pure act of betting on God by using our most strict rules—namely, MaxiMin and MaxiMax. However, as an alternative, they might appeal to a weaker criterion for decision-making, such as E-Admissibility. Of course, even granting that salvation has a large finite utility, some work will be needed for this approach to succeed, especially if one attempts to justify W_G as the unique optimal choice for all those five rules.

Though interpreting the outcome of salvation as having a finite utility fails to do justice to Pascal's original thought, the general lesson seems to be clear. One cannot draw any substantial conclusion from this case unless one provides more information about the extent of an agent's imprecise probabilities and the values of all utilities in the matrix. If the decision matrix does not vary from person to person, the decisive question would be whether salvation has a sufficiently large utility, one that makes W_G rationally mandatory.

7 Concluding Remarks

In this paper, I have surveyed how decision-making with imprecise probabilities affects Pascal's Wager. I have put forward five distinct principles, each understood as a natural extension of the EU maximization rule, showing the results of applying them to two scenarios of state uncertainty. Also, I have claimed that rationality requires ruling out 0 as the lower probability of God's existence. Even granting that this is the case, introducing mixed strategies into the set of available options poses complications to Pascal's Wager with imprecise probabilities. However, as I have argued, there are at least two theoretical approaches to salvation that can help Pascal's argument to get off the ground in the imprecise setting. When combined with imprecise probabilities, discriminating between orders of infinity provides a way out for the proponents of the Wager except for Interval Dominance, which might recommend a mixed strategy as a rationally permissible option. Moving to finite utilities is another potentially successful route for defending Pascal's Wager, though a substantial caveat must be made. Because E-Admissibility, Maximality, and Interval Dominance typically include larger sets of permissible acts, it is definitely less obvious that wagering for God will be rationally mandatory for such principles. For those permissive rules, depending on the range of probabilities, the utility of salvation must be enormously large for the act of betting on God to defeat any mixed strategy.

Acknowledgements: I am grateful to the John Templeton Foundation for funding during my year as a 2019 LATAM Bridges Mentoring Fellow at Purdue University. I am especially indebted to Paul Draper for his generous help and advice during my stay at Purdue. I would also like to thank Luis Rosa and two anonymous referees for extensive comments and valuable suggestions on an earlier version of this paper.

References

Bartha, P. 2007. "Taking Stock of Infinite Value: Pascal's Wager and Relative Utilities." *Synthese* 154(1):5–52.

Bartha, P. 2018. Pascal's Wager and the Dynamics of Rational Deliberation. In *Pas-*

cal's Wager, ed. Paul Bartha and Lawrence Pasternack. Cambridge, UK: Cambridge University Press pp. 236–259.

Bartha, P. and Pasternack L. 2018. Introduction. In *Pascal's Wager*, ed. Paul Bartha and Lawrence Pasternack. Cambridge, UK: Cambridge University Press pp. 1–23.

Bradley, R. and M Drechsler. 2014. "Types of Uncertainty." *Erkenntnis* 79(6):1225–1248.

Buchak, L. 2013. *Risk and Rationality.* Oxford: Oxford University Press.

Chandler, J. 2014. "Subjective probabilities need not be sharp." *Erkenntnis* 79(6):1273–1286.

Conway, J. A. 1976. *On Numbers and Games.* London: Academic Press.

Duff, Antony. 1986. "Pascal's Wager and Infinite Utilities." *Analysis* 46(2):107–109.

Duncan, Craig. 2003. "Do Vague Probabilities Really Scotch Pascal's Wager?" *Philosophical Studies* 112(3):279–290.

Duncan, Craig. 2007. "The Persecutor's Wager." *The Philosophical Review* 116(1):1–50.

Duncan, Craig. 2018. The Many-Gods Objection to Pascal's Wager: A Defeat, then a Resurrection. In *Pascal's Wager*, ed. Paul Bartha and Lawrence Pasternack. Cambridge, UK: Cambridge University Press pp. 148–167.

Elga, A. 2010. "Subjective Probabilities Should be Sharp." *Philosopher's Imprint* 10(5):1–10.

Elga, A. 2012. "Errata for Subjective Probabilities should be Sharp." *Unpublished manuscript.* URL: http://www.princeton.edu/ adame//papers/sharp/sharp-errata.pdf. Last access: July 25, 2020.

Gärdenfors, P and N Sahlin. 1982. "Unreliable Probabilities, Risk Taking, and Decision Making." *Synthese* 53(3):361–386.

Gilboa, I and D Schmeidler. 1993*a*. "Maxmin Expected Utility with a Non-unique Prior." *Journal of Mathematical Economics* 18(2):141–153.

Gilboa, I and D Schmeidler. 1993*b*. "Updating Ambiguous Beliefs." *Journal of Economic Theory* 59(1):33–49.

Hájek, A. 2000. "Objecting Vaguely to Pascal's Wager." *Philosophical Studies* 98(1):1–16.

Hájek, A. 2003. "Waging War on Pascal's Wager." *The Philosophical Review* 112(1):27–56.

Hájek, A. 2018*a*. The (In)validity of Pascal's Wager. In *Pascal's Wager*, ed. Paul Bartha and Lawrence Pasternack. Cambridge, UK: Cambridge University Press pp. 123–147.

Hájek, A. 2018*b*. Pascal's Wager. In *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), ed. Edward N. Zalta. URL: https://plato.stanford.edu/archives/sum2018/entries/pascal-wager/.

Hammond, P. 1988. "Orderly decision theory: a comment on Professor Seidenfeld." *Economics & Philosophy* 4(2):292–297.

Hausner, M. 1954. Multidimensional Utilities. In *Decision Processes*, ed. R. M Thrall, C. H Coombs and R. L Davis. New York: John Wiley and Sons pp. 167–180.

Jordan, J. 2006. *Pascal's Wager: Pragmatic Arguments and Belief in God.* Oxford: Oxford University Press.

Joyce, J. 2005. "How Probabilities Reflect Evidence." *Philosophical Perspectives* 19(1):153–178.

Kenny, Anthony. 1979. *The God of Philosophers.* Oxford: Clarendon Press.

Levi, I. 1974. "On indeterminate probabilities." *The Journal of Philosophy* 71(13):391–418.

Levi, I. 1980. *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance.* Cambridge, MA: MIT Press.

Lovering, Rob. 2013. *God and Evidence: Problems for Theistic Philosophers.* New York: Bloomsbury.

Martin, Martin. 2000. Omniscience and Incoherence. In *Medieval Philosophy and Modern Times*, ed. G. Holmström-Hintikka. Dordrecht: Kluwer Academic Publishers pp. 17–34.

McClennen, E. 1994. Finite Decision Theory. In *Gambling on God: Essays on Pascal's Wager*, ed. J Jordna. Lanham, Maryland: Rowman and Littlefield pp. 115–138.

Pascal, B. 1960. *Pensées*. Trans. John Warrington. London: J. M. Dent & Sons.

Rinard, S. 2013. "Against Radical Credal Imprecision." *Thought* 2(2):157–165.

Rinard, S. 2018. Pascal's Wager and Imprecise Probability. In *Pascal's Wager*, ed. Paul Bartha and Lawrence Pasternack. Cambridge, UK: Cambridge University Press pp. 278–292.

Savage, Leonard Jimmie. 1954. *The Foundations of Statistics*. New York: Dover.

Schlesinger, G. 1994. A Central Theistic Argument. In *Gambling on God: Essays on Pascal's Wager*, ed. J Jordan. Lanham, MD: Roman and Littlefield pp. 83–100.

Seidenfeld, T. 2004. "A Contrast between Two Decision Rules for use with (convex) Sets of Probabilities: Γ-maximin versus E-admissibility." *Synthese* 140(1/2):69–88.

Sorensen, R. 1994. Infinite Decision Theory. In *Gambling on God: Essays on Pascal's Wager*, ed. J Jordan. Lanham, MD: Roman and Littlefield pp. 139–159.

Troffaes, M. 2007. "Decision making under uncertainty using imprecise probabilities." *International Journal of Approximate Reasoning* 45(1):17–29.

Van Fraassen, B. 1980. "Rational belief and probability kinematics." *Philosophy of Science* 47(2):165–187.

von Neumann, J. and O Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.

Weatherson, B. 1998. "Decision Theory with Imprecise Probabilities." *Unpublished manuscript.* URL: http://brian.weatherson.org/vdt.pdf. Last access: July 20, 2020.

Weisberg, Jonathan. 2011. Varieties of Bayesianism. In *Handbook of the History of Logic, vol. 10, Inductive Logic*, ed. Dov Gabbay, Stephan Hartmann and John Woods. Amsterdan: North-Holland pp. 477–551.