Introduction

Paul Bartha and Lawrence Pasternack

1 Pascal's Wager: An Argument with Many Audiences

Pascal's Wager is one of the great classic arguments for belief in God. The other great theological arguments – ontological, cosmological, and teleological – aim to establish that God's existence is necessary or probable. Pascal's Wager, by contrast, is a prudential or "pragmatic" argument. The conclusion is a recommended action: it is in your interest to believe (or to strive to believe) in God, even if you insist that the probability of God's existence may be very low. In its most familiar contemporary formulation, the basic argument runs as follows: so long as the probability of an infinite afterlife reward for belief in God is greater than zero, "wagering for God" has infinite expected utility and is therefore superior to "wagering against God" (which has at best finite expected utility). In short, wagering is a good gamble.

The argument is embedded deep within the *Pensées*. That work was addressed to Pascal's first audience, worldly seventeenth-century Parisians for whom religion was largely an object of indifference. For Pascal, by contrast, nothing could be more serious. Influenced by Jansenist theology, Pascal believed that human nature was thoroughly corrupt. God's grace, freely accepted, could put us on the path to salvation.[1] In Pascal's libertine audience, however, a hard shell of skeptical hostility, reinforced by secular reasoning and habits, formed a barrier to divine grace. The Wager was just one in a series of manoeuvers designed by Pascal to chip away at that barrier. A distinctive feature of his argument is its appeal to practical reason (and gambling instincts), in recognition of the limitations of theoretical reason when it

[1] Here we set aside Pascal's belief in the existence of a predetermined class of *elect individuals* who will be saved. Several of the chapters in this collection (Wood, Moser, Franklin) address the complex theological background to the *Pensées* and its seeming conflict with an argument designed to promote belief in God.

comes to faith. Pascal explains that those who lack faith can only hope to achieve it through (practical) reasoning, “until God gives it by moving their heart” (L110/S142). The conclusion of the Wager is not that one should instantly believe in God – neither a realistic nor a possible act given Pascal’s understanding of grace – but rather that one should *take steps* toward a “cure,” measures that eliminate obstacles to faith. Pascal writes: by “taking holy water, having masses said, and so on,” you can “diminish the passions which are your great obstacles” (L418/S680).

The argument has been vigorously debated ever since its earliest appearance. We might identify a second, distinctively philosophical audience for the Wager beginning with Diderot, and extending through Kant, Kierkegaard, Nietzsche, and William James.[2] This group of readers recognized both the virtues and the weaknesses of Pascal’s argument, and in some cases developed descendant versions. Increasingly, the Wager came to be considered as a stand-alone argument, removed from the context of the *Pensées*.

This volume is dedicated to a third audience: contemporary philosophers and philosophy students. As is clear from its enduring influence and appeal, Pascal’s Wager is a remarkable argument in many ways. It is remarkable for bringing together *big ideas*: infinity, God, salvation, and prudential and evidential reasoning. It is remarkable for its invention of a new style of argument, ultimately formalized as *decision theory*. Decision theory provides both insight into the original argument of the *Pensées* and the tools to develop elaborate contemporary variations. The Wager is also remarkable for its influence on later philosophers. Finally, it is remarkable for its enduring relevance to many areas of philosophy: philosophy of religion, decision theory, formal epistemology, and broader currents of thought. Pascal’s Wager continues to have many audiences.

This collection is designed for contemporary students and philosophers. It includes chapters that explore the historical context of Pascal’s argument and its influence on later philosophers. It includes discussion of some of the central objections and debates about the Wager. Additional chapters show how ideas in the philosophy of probability and decision theory – imprecise and infinitesimal credences, non-standard decision theory, infinite utility – shape current debates about the argument. These chapters also reveal the many ways in which Pascal’s Wager, in turn, has had and continues to have an important influence on different areas of philosophy.

[2] Pascal’s influence on these later philosophers is discussed in this collection’s chapters by Buben and Jordan.

This introduction begins by setting out the technical background for contemporary discussions of Pascal's Wager. This consists of the basic elements of decision theory (§2) and assumptions needed for reasoning about infinite utility (§3). We then review (§4) three versions of Pascal's Wager that appear to be present in the *Pensées*, and provide a brief review of some of the classic objections to the most familiar version. Finally (§5), we provide an overview of the contributed chapters in this book.

2 Standard Decision Theory: Dominance, Expected Value, and Expected Utility

2.1 Dominance

Consider the following gamble. A fair coin will be tossed. On a result of *Heads* you win $100, while *Tails* pays you nothing. If you reject the gamble, you gain nothing. Assuming that you prefer $100 to $0, you should take the bet!

A *decision table* helps to illuminate the reasoning that leads to this conclusion. A decision table has one row for each possible *act* and one column for each possible *state*. Each act–state combination results in an *outcome*, identified here with a monetary sum.

One act *dominates* another if it does *at least as well* as the other act in every possible state, and *strictly better* in at least one state.[3] An act *strictly dominates* another if it does *strictly better* in every state. In Table I.1, *Bet* dominates *Don't bet* (although not with strict dominance). A good decision rule for this example is the *Dominance Principle*.

> *Dominance Principle*: Choose an act that dominates all other available actions, if such an act is available.

We have to be a bit careful with this rule, but it works well here.[4]

Table I.1 Dominance Reasoning

	Heads	*Tails*
Bet	$100	$0
Don't bet	$0	$0

3 This is sometimes called *weak dominance*, but we will simply call it *dominance*.
4 The main restriction is that the probabilities of the possible states should *not vary* depending on what action is selected; see Hájek's chapter in this volume for discussion.

Table I.2 EMV Reasoning

	(0.5) *Heads*	(0.5) *Tails*	
Bet	\$100	\$0	EMV(*Bet*) = (0.5)(\$100) + (0.5)(\$0) = \$50
Don't bet	\$40	\$40	EMV(*Don't bet*) = \$40

2.2 Expected Monetary Value (EMV)

The *expected monetary value* (EMV) if you take the bet is \$50, since the probability of *Heads* is 0.5. This number represents your average winnings: you expect to win \$100 half the time and \$0 half the time.

Suppose you learn that the gamble is not free. The price is \$40. No longer do we have a dominant act, but it still looks like we have a good bet (see Table I.2). We justify this conclusion with a decision table that is a slight variation from the previous one. Each outcome is now your *final amount of money* (assuming that you start with \$40 and keep it if you *Don't bet*).

The EMV for each act is calculated as a weighted average of the possible outcomes: you multiply the monetary outcome of that act under each state by the probability of the state, and sum over all possible states. (We have indicated the probability above each state. In any decision table, the states must be mutually exclusive and the probabilities must sum to 1.) A good decision rule for this case is *EMV Maximization*.

EMV Maximization: Choose an act that maximizes EMV.

In this case, that act is *Bet*. Notice that EMV maximization is consistent with the *Dominance Principle*, and would also recommend *Bet* in Table I.1.

2.3 Expected Utility (EU)

In many cases, we need a more general decision rule than either the *Dominance Principle* or *EMV Maximization*. That will certainly be true in decisions where there may be *non-monetary* outcomes: a trip to the beach, an enjoyable movie, or a miserable cold. Even with purely monetary outcomes, EMV maximization may not be your best guide. Although you prefer \$100 to \$40, suppose that you desperately need the \$40 for a cab fare to get home. You don't wish to risk it in a gamble even though the odds are in your favor. You represent your preferences by replacing the dollar values in the decision table

Table I.3 EU Reasoning

	(0.5) *Heads*	(0.5) *Tails*	
Bet	10	1	EU(*Bet*) = (0.5)(10) + (0.5)(1) = 5.5
Don't bet	9	9	EU(*Don't bet*) = 9

with *utilities*: $u(\$100) = 10$ in place of \$100, $u(\$0) = 1$ in place of \$0, and $u(\$40) = 9$ in place of \$40. (Don't worry too much about how we assign numerical values to the utilities; we'll get to this point shortly.) This yields the data laid out in Table I.3.

The *expected utility* (EU) of each act is calculated by multiplying the utility of that act under each state by the probability of the state,[5] and summing over all possible states. (This quantity is also referred to as the *expectation* of the act.) The decision rule for this case is *EU Maximization*.

EU Maximization: Choose an act that maximizes EU.[6]

In this case, that act is *Don't bet*. Notice that *EU Maximization* can differ from *EMV Maximization*.

In this analysis, a *utility function u* represents your preferences. The function u assigns a numerical value (a real number) to each possible outcome. The most basic requirement for faithful representation is that higher numbers correspond to preferred outcomes: $u(\text{Outcome 1}) \geq u(\text{Outcome 2})$ exactly when Outcome 1 is as good as or better than Outcome 2.

EU Maximization is more generally useful than *EMV Maximization*, and it is the most fundamental principle in standard decision theory. For this decision rule to make sense, however, you have to be careful about how you represent your preferences. Suppose that you change the numbers in Table I.3 to give a different utility function u^*: $u^*(\$100) = 10$, $u^*(\$40) = 2$, and $u^*(\$0) = 1$. Then EU*(*Bet*) = 5.5 as before, but EU*(*Don't bet*) = 2. The new utility function u^* accurately represents your preference ordering, but now *EU Maximization* seems to recommend *Bet*! To prevent this type of

[5] EU calculations use subjective probabilities, also called *credences*. In this example, we assume that the subjective probabilities coincide with the objective probabilities of 0.5 for *Heads* or *Tails*.

[6] In decision problems where several acts are tied for maximum EU, there is no clear recommendation. One other subtlety: in many decision problems, the probabilities of one or more states may vary depending on the act. For these problems, the expected utility of each act is calculated using *conditional probabilities* rather than fixed probabilities for each possible state. In the case of *Pascal's Wager*, we can reasonably suppose that the agent's choice has no influence on the probability that God exists.

contradiction, the *intervals* between utilities must also reflect your preferences accurately. In Table I.3, the interval between having \$40 and having \$0 is eight times as large as the interval between having \$100 and having \$40. If this reflects the importance of keeping \$40 for cab fare, then the second utility function u^* is not a faithful representation of your preferences. In order to be useful in expected utility calculations, the utility function u should represent not just the preference ordering but also the *structure* of your preferences.

Standard decision theory proves a remarkable result that justifies the rule of EU maximization and ensures that there will be no conflicting recommendations, but it requires some strong assumptions about rational agents. The theory assumes, first, that rational agents have a well-defined *preference ordering* among all possible outcomes. This means that for any two outcomes O_1 and O_2, the agent either prefers O_1 to O_2 (written $O_1 \succ O_2$), prefers O_2 to O_1 (written $O_2 \succ O_1$), or is indifferent between O_1 and O_2 (written $O_1 \sim O_2$).[7] Second, standard decision theory identifies a set of *preference axioms* which are supposed to be satisfied by rational agents. For instance, preferences should be *transitive*: if the agent prefers O_1 to O_2 and prefers O_2 to O_3, then the agent prefers O_1 to O_3.[8] Given these assumptions, the fundamental result of standard decision theory is the *Expected Utility Theorem*: there is a way to represent the agent's preferences with a utility function u so that the utility of any gamble is identical to its expected utility and the best act is always an act that maximizes expected utility. The function u that represents the fine structure of your preferences in this way is almost unique.[9]

Mixed strategies are also an important part of decision theory. In our decision tables, each row corresponds to a *pure act*, something that is within the agent's control. An example of a mixed strategy is to flip a fair coin and then *Bet* if the result is *Heads*, but *Don't bet* if the result is *Tails*. We can represent this as [0.5 *Bet*, 0.5 *Don't bet*], which indicates a probability 0.5 attached to each possible pure act. Of course, decision tables can have more than two rows. In general, a mixed strategy is *any* assignment of probabilities to a set of possible pure acts, so long as the probabilities sum to 1. The *Expected Utility Theorem* shows that the utility

[7] Standard decision theory assumes that you have these clear preferences even when the outcomes are gambles, rather than simple outcomes.

[8] For a full list of the axioms of standard decision theory, see Resnik (1987).

[9] If u and v are two utility functions that can be used in EU calculations, then there are constant numbers a and b with $a > 0$ such that $v = au + b$.

of any mixed strategy is its expected utility. For example, based on Table I.3,

$$\begin{aligned} EU([0.5\ Bet,\ 0.5\ Don't\ bet]) &= 0.5\ EU(Bet)\ +\ 0.5\ EU(Don't\ bet) \\ &= 0.5(5.5)\ +\ 0.5(9) \\ &= 7.25 \end{aligned}$$

Since EU(*Don't bet*) = 9, you are better off with the pure act *Don't bet* than with the mixed strategy of the coin toss.

3 Infinite Utility

Pascal's Wager is commonly represented as a decision-theoretic argument. Pascal seems to suggest several different versions of the argument, but all of them share the same basic decision table (see Table I.4).[10]

The numbers f_1, f_2, and f_3 are all *finite* utilities. If God does not exist, then the rewards and penalties for wagering or not wagering are finite. Pascal seems to think that if God does exist, then the penalty for not wagering is also finite. In the top-left corner, however, we have the possibility of salvation, which Pascal characterizes as "an infinity of infinitely happy life." Since this is better than any finite reward, we represent the utility as ∞.[11]

Infinity in mathematics is an extremely useful concept for talking about certain kinds of limits. For instance:

$$1 + 2 + 4 + 8 + \ldots = \infty.$$

On the left, we have an *infinite series*. What the equation means is that the finite partial sums (first 10 terms, first 100 terms, and so forth) increase without bound. No matter what number we pick, we'll eventually pass it by adding enough terms in the series. If we replace the terms with −1, −2, and so on, then the finite partial sums decrease without bound.

Table I.4 Pascal's Wager

	God exists	*God does not exist*
Wager for God	∞	f_1
Wager against God	f_2	f_3

10 Here we follow Hájek (2012b).

11 Note that even if wagering only produces a *positive chance* of salvation if God exists, the correct value for this outcome is still ∞, as will be explained shortly.

In order to talk about such limits, it is convenient to add the two elements $+\infty$ (usually written ∞) and $-\infty$ to the real numbers. The new set is called the *extended real numbers*.[12] We extend the definition of *less-than*, <, by the assumption

(1) For all real numbers x: $-\infty < x < \infty$.

Note that real numbers are all finite; the new elements ∞ and $-\infty$ are *not* real numbers. We extend the definition of basic arithmetical operations by making the following assumptions for operations where one argument is finite and one is infinite:

(2a) For all real numbers x: $x + \infty = \infty$ and $x - \infty = -\infty$.
(2b) For all real numbers x: $x \cdot \infty = \infty$ if $x > 0$ and $x \cdot \infty = -\infty$ if $x < 0$.
(2c) For all real numbers x: $x \cdot -\infty = -\infty$ if $x > 0$ and $x \cdot -\infty = \infty$ if $x < 0$.
(2d) $0 \cdot \infty = 0$ and $0 \cdot -\infty = 0$.

Finally, we make some assumptions about operations where both arguments are infinite:

(3a) $\infty + \infty = \infty$ and $-\infty + -\infty = -\infty$. Note that $\infty - \infty$ is *undefined*.
(3b) $\infty \cdot \infty = \infty$ and $\infty \cdot -\infty = -\infty$.
(3c) $-\infty \cdot \infty = -\infty$ and $-\infty \cdot -\infty = \infty$.

These assumptions correspond to results about limits. For example, (3a) corresponds to the fact that if we take two infinite series, each of which sums to ∞, and form a new series by adding them together term by term, the new series will also sum to ∞. But if we subtract them term by term, the new series might sum to ∞, $-\infty$, or any finite number; hence, in general, $\infty - \infty$ is undefined.

Using the extended real number ∞ in the decision table is the easiest way to represent the infinite value of salvation. Our assumptions, especially (2a) and (2b), give us all the mathematics that we need to perform the expected utility calculations that we encounter in many discussions of Pascal's Wager. It is also plausible to argue that these assumptions faithfully represent Pascal's own statements about infinity (Hájek, 2003). But there is still a big problem. As McClennen (1994) notes, the introduction of ∞ as a possible utility value takes us beyond standard decision theory. Infinite utility is incompatible with one of the preference axioms needed in standard decision theory. Most discussions of Pascal's Wager simply ignore this problem and assume that

[12] The following assumptions are taken from Royden (1968).

decision theory, augmented with ∞ and −∞ as possible utility values, works in almost exactly the same way as standard decision theory.[13] We refer to this non-rigorous approach as *naïve infinite decision theory*. One way to avoid this problem is to use a *very large but finite* value for the utility of salvation. This is a sensible option adopted in several of the chapters below, but it has significant consequences for Pascal's argument. Finally, there are a variety of rigorous approaches, unavailable in Pascal's day, which allow us to represent the value of salvation as infinite. We refer to them as *non-standard decision theory*, but we leave their discussion to individual chapters of the book.

4 Pascal's Wager

Hacking (1994 [1972]) has identified three distinct decision-theoretic arguments in *Pensées* L418/S680. In this section, we present the core text (Krailsheimer translation), along with the three decision-theoretic arguments.[14]

Pascal begins by characterizing the decision about belief in God as one that reason cannot decide (on the basis of evidence), but also as one that cannot be avoided. He notes that two things are at stake, knowledge and happiness, but (at least initially) it seems that knowledge is not to be had, so that the case for wagering must rest on happiness.

> Let us then examine this point, and let us say: "Either God is or he is not." But to which view shall we be inclined? Reason cannot decide this question. Infinite chaos separates us. At the far end of this infinite distance a coin is being spun which will come down heads or tails. How will you wager? Reason cannot make you choose either, reason cannot prove either wrong.
>
> Do not then condemn as wrong those who have made a choice, for you know nothing about it. "No, but I will condemn them not for having made this particular choice, but any choice, for, although the one who calls heads and the other one are equally at fault, the fact is that they are both at fault: the right thing is not to wager at all."
>
> Yes, but you must wager. There is no choice, you are already committed. Which will you choose then? Let us see: since a choice must be made, let us see which offers you the least interest. You have two things to lose: the true

13 It won't be exactly the same because we might occasionally encounter a bet whose expected utility is undefined because it involves the calculation ∞ − ∞.

14 These three arguments, together with Hacking's analysis, receive detailed discussion in Hájek's chapter in this volume.

> and the good; and two things to stake: your reason and your will, your knowledge and your happiness; and your nature has two things to avoid: error and wretchedness. Since you must necessarily choose, your reason is no more affronted by choosing one rather than the other. That is one point cleared up.

Pascal now proceeds to the first version of the Wager, which Hacking calls the *argument from dominance*.

4.1 Argument from Dominance: $f_1 \geq f_3$

The text continues:

> But your happiness? Let us weigh up the gain and the loss involved in calling heads that God exists. Let us assess the two cases: if you win you win everything, if you lose you lose nothing. Do not hesitate then; wager that he does exist.

Consider the decision table (Table I.5). Pascal states, "if you lose you lose nothing." Hacking writes: "if God is not, then both courses of action are pretty much on a par. You will live your life and have no bad effects either way from supernatural intervention." If we interpret this as $f_1 = f_3$ in Table I.5, we have the argument from dominance. *Wager for God* weakly dominates *Wager against God*; the *Dominance Principle* tells us to accept the Wager.

Much later in the text, Pascal suggests that we might actually have what amounts to an argument from *strict dominance*, i.e., $f_1 > f_3$. He offers the following reassurance for one who opts to wager that God exists:

> Now what harm will come to you from choosing this course? You will be faithful, honest, humble, grateful, full of good works, a sincere, true friend . . . It is true you will not enjoy noxious pleasures, glory and good living, but will you not have others?
>
> I tell you that you will gain even in this life, and that at every step you take along this road you will see that your gain is so certain and your risk so

Table I.5 Pascal's Wager

	God exists	*God does not exist*
Wager for God	∞	f_1
Wager against God	f_2	f_3