P. BARTHA

# TAKING STOCK OF INFINITE VALUE: PASCAL'S WAGER AND RELATIVE UTILITIES

ABSTRACT. Among recent objections to Pascal's Wager, two are especially compelling. The first is that decision theory, and specifically the requirement of maximizing expected utility, is incompatible with infinite utility values. The second is that even if infinite utility values are admitted, the argument of the Wager is invalid provided that we allow mixed strategies. Furthermore, Hájek (Philosophical Review 112, 2003) has shown that reformulations of Pascal's Wager that address these criticisms inevitably lead to arguments that are philosophically unsatisfying and historically unfaithful. Both the objections and Hájek's philosophical worries disappear, however, if we represent our preferences using relative utilities (generalized utility ratios) rather than a one-place utility function. Relative utilities provide a conservative way to make sense of infinite value that preserves the familiar equation of rationality with the maximization of expected utility. They also provide a means of investigating a broader class of problems related to the Wager.

## 1. INTRODUCTION

It is striking that the earliest example of an argument in decision theory, Pascal's Wager, invokes the notion of an infinite reward. As McClennen (1994) has observed, the intellectual descendants of Pascal built decision theory around a set of assumptions that "logically excludes infinite utilities." At worst, the concept of infinite utility is incoherent; at best, its admission into decision theory leads to numerous paradoxes and puzzles that are acknowledged even by its advocates.[1] Many people have given up on infinite utility (Jeffrey 1983; Duff 1986; McClennen 1994; Jordan 1998). There have been attempts to reformulate the Wager in ways that invoke only finite utilities (Mougin and Sober 1994; Sobel 1996; Jordan 1998; Hájek 2003).

But the Wager's enduring fascination derives from its appeal to an infinite reward that, even if discounted by allowing only a tiny probability that God exists, swamps any expectation that can be

Synthese (2007) 154: 5–52 
DOI 10.1007/s11229-005-8006-z

derived from mere worldly pleasures. Following Hájek (2003), we may think of the Wager as depending upon three premises.

*Premise 1*: The decision matrix in Table I describes the choice facing the wagerer.

Here, $f_2$, $f_3$ and $f_4$ are finite utility values. Wagering for God involves taking steps that, one hopes, will lead to full belief. Wagering against involves taking steps that shore up an agnostic or atheistic stance, or simply choosing to ignore the whole business.[2]

*Premise 2*: God's existence has positive, finite probability $q$.[3]

*Premise 3*: We are rationally required to perform the act that has maximum expected utility.

The expected utilities are as follows:

$$\text{EU(Wager for God)} = q \cdot \infty + (1-q) f_2 = \infty$$
$$\text{EU(Wager against God)} = q \cdot f_3 + (1-q) f_4 = \text{finite}$$

For we are supposing that $q \cdot \infty = \infty$ for any positive $q$, and that $\infty + k = \infty$ for any finite number $k$. It follows that we are rationally required to wager for God.[4]

I shall be concerned primarily with two recent objections to this argument. The first has to do with *Premise 3*: decision theory, and specifically the requirement of maximizing expected utility, is incompatible with infinite utility values. My response, developed in Sections 4 and 5, is that we can accommodate infinite utility *ratios* (I call them *relative utilities*) and we can justify a modified version of *Premise 3*. These changes provide a satisfactory response to this first objection.

The second objection begins with an observation made by Jeffrey (1983), Duff (1986) and most recently Hájek (2003): even if infinite utility values are admitted and each of the three premises above is granted, the argument is invalid provided that we allow mixed

TABLE I

Pascal's Wager (standard form)

| | God exists | God does not exist |
|---|---|---|
| Wager for God | $\infty$ | $f_2$ |
| Wager against God | $f_3$ | $f_4$ |

strategies between wagering-for and wagering-against. Given the existence of mixed strategies, the premises do not entail that we are rationally required to wager for God. Any mixed strategy that gives a positive probability to coming to wager for God also yields infinite expectation. A strategy such as flipping a coin to decide whether or not to wager for God has infinite expectation and appears to have as great a claim on our rationality as the pure wager.

This second objection is compelling and has led people to develop elaborate reformulations of the Wager. These reformulations all yield valid arguments that bear some resemblance to Pascal's original argument. At this point, the 'invalidity objection' takes a more sophisticated form. One can question the reformulations both for their fidelity to Pascal and for their philosophical cogency. Hájek (2003) maintains that existing reformulations of the Wager fail on both counts, and presents an intriguing argument that no valid reformulation will ever meet these two criteria. His argument takes the form of a dilemma: our representation of the utilities in the Wager inevitably produces either an invalid argument or one that is philosophically unsatisfying and historically unfaithful.

My contention is that both objections to the Wager fail. We can find a representation of our preferences in the Wager that escapes both Hájek's dilemma and the incompatibility objection. The strategy is to represent our preferences using *relative utilities* rather than a one-place utility function.

The two objections are explained in Section 2. Section 3 reviews Hájek's dilemma. Section 4 defines relative utilities. Section 5 employs relative utilities in a reformulation of the Wager that is not vulnerable to any of the objections. The final two sections explore how relative utilities might illuminate a broader class of problems.

## 2. TWO CRITICISMS OF THE WAGER

### 2.1. *The Incompatibility Objection*

McClennen has noted a tension between *Premise 3* and *Premise 1*. *Premise 3*, that we ought always to maximize expected utility, is justified by a result referred to as the *Expected Utility Theorem* (Resnik 1987), or more simply as *linearity*. Yet as McClennen observes, the admission of outcomes with infinite utility sanctioned by *Premise 1* is incompatible with some of the very axioms that play a role in deriving linearity. This undercuts the Wager, because if *Premise 1* and *Premise 3* are not co-tenable, then the argument as stated must be unsound.

The *Expected Utility Theorem* is a result about how our preferences can be represented if they satisfy certain standard axioms. To appreciate McClennen's objection, we need to understand what the theorem says, the axioms upon which it depends, and its connection to *Premise 3*. We begin with the axioms, here formulated following Resnik (1987).

Let $A$, $B$ and $C$ stand for outcomes, and $p$ and $q$ for probability values. Let $[pA, (1-p)B]$ stand for a gamble with probability $p$ of yielding outcome $A$ and probability $(1-p)$ of yielding outcome $B$. Any pure outcome $A$ is equivalent to the gamble $[pA, (1-p)A]$ for any $p$ with $0 \leq p \leq 1$, or to $[1A, 0B]$ for any $B$. Accordingly, the terms 'gamble', 'outcome' and so forth will be used interchangeably, although we shall occasionally distinguish between pure outcomes and gambles.

Next, we assume a weak preference ordering $\preccurlyeq$ among outcomes. Write $A \preccurlyeq B$ if we weakly prefer $B$ to $A$ (meaning that $B$ is as good as or better than $A$), and $A \sim B$ if both $A \preccurlyeq B$ and $B \preccurlyeq A$ (i.e., we are indifferent between $A$ and $B$). Write $A \prec B$ if we strictly prefer $B$ to $A$ (i.e., $A \preccurlyeq B$ but not $A \sim B$). With this notation in hand, we can state the relevant axioms.

**(A1) Ordering conditions.**
The preference relation $\preccurlyeq$ is a *total* ordering (for any $A$ and $B$, either $A \preccurlyeq B$ or $B \preccurlyeq A$) that is *reflexive* ($A \preccurlyeq A$) and *transitive* ($A \preccurlyeq B$ and $B \preccurlyeq C$ imply $A \preccurlyeq C$).

**(A2) Better-prizes condition.**
$A \preccurlyeq B$ iff for any $0 \leq p \leq 1$ and any $C$, $[pC, (1-p)A] \preccurlyeq [pC, (1-p)B]$ and $[pA, (1-p)C] \preccurlyeq [pB, (1-p)C]$. If $A \prec B$ (and $0 < p < 1$), the other preferences are also strict.
[Keeping the probabilities fixed, prefer gambles that substitute better prizes.]

**(A3) Better-chances condition.**[5]
If $A \preccurlyeq B$, then for any $0 \leq p, q \leq 1$,

$$p \geq q \text{ iff } [pA, (1-p)B] \preccurlyeq [qA, (1-q)B].$$

If $A \prec B$ and $p > q$, the final preference is strict.
[Keeping prizes fixed, prefer a gamble just in case it offers a better chance for the better prize.]

**(A4) Reduction-of-compound-gambles condition.**
For any $A$ and $B$ and any $0 \leq p, q, r \leq 1$,
$$[p[qA, (1-q)B], (1-p)[rA, (1-r)B]] \sim [tA, (1-t)B]$$
where $t = pq + (1-p)r$.
[Evaluate compound gambles in accordance with the probability calculus.]

**(A5) Continuity (Archimedean) condition.**
If $A \preccurlyeq B$ and $B \preccurlyeq C$, then there is some number $p$ such that $0 \leq p \leq 1$ and $B \sim [pA, (1-p)C]$.[6]
[Whenever an outcome $B$ is ranked between two others $A$ and $C$, there is some gamble between $A$ and $C$ such that the agent is indifferent between it and $B$.]

The *Expected Utility Theorem* states that if our preferences satisfy these axioms, then they can be represented by a real-valued utility function $u$ with two important properties:

(1) $A \preccurlyeq B$ just in case $u(A) \leq u(B)$. The ranking of outcomes according to our preferences agrees with the ranking assigned by $u$.
(2) $u([pA, (1-p)B]) = pu(A) + (1-p)u(B)$, so that $u$ is linear in gambles. The utility of any gamble is exactly its expected utility.

This function $u$ is unique up to a positive linear transformation.

This result provides the justification for *Premise 3*. Following Luce and Raiffa (1957), McClennen reminds us that rationality is properly characterized as the maximization of preference satisfaction. If our preferences conform to axioms (A1) – (A5), then there is a utility function such that maximizing preference satisfaction coincides with maximizing expected utility. If not, there is no guarantee that the optimal gamble is the one with the highest expected utility. The actual utility of a gamble might exceed or fall short of its expected utility.

It is a mistake, then, to divorce the principle of maximizing expected utility from a highly structured set of preferences, as reflected in the conditions (A1)–(A5). Those conditions are needed to derive linearity, which in turn is required to justify *Premise 3* of the Wager. Only with *Premise 3*, it appears, can the argument of the Wager succeed.

Now suppose that our utility function $u$ can take on infinite values. We at once run into trouble with the axioms identified above.

*Failure of (A2)*: Suppose $u(A)=0$ and $u(B)=1$, so that $A \preccurlyeq B$. If we let $C$ be an outcome with $u(C)=\infty$, then $[pC, (1-p)A] \sim [pC, (1-p)B]$ provided $p>0$, because by the *Expected Utility Theorem*, each of these gambles has infinite utility.

*Failure of (A3)*: Suppose $u(A)=\infty$ and $u(B)=k$ for any finite $k$. Then $[pA, (1-p)B] \sim [qA, (1-q)B]$ so long as $p \geq q > 0$, because each of these gambles has infinite utility.

*Failure of (A5)*: Suppose $u(A)=\infty$, while $u(B)=1$ and $u(C)=0$. We should be indifferent between $B$ and some gamble between $A$ and $C$, but $u[pA, (1-p)C] = pu(A)+(1-p)u(C)$. This will be $\infty$ if $p>0$ and 0 if $p=0$, but never 1.

In each case, failure arises from a property of $\infty$ that Hájek terms *reflexivity under multiplication*:

(Ref $\cdot$) For any $p>0$, $p \cdot \infty = \infty$

It follows from reflexivity that all gambles that offer any positive chance for an infinite prize turn out to be equally good. As McLennen notes, this is a somewhat counter-intuitive position (and one that we shall eventually reject), but one cannot give it up without also giving up the idea that an agent's preferences can be represented by a utility function with the feature that the agent must be indifferent between acts that have the same expected utility. Furthermore, as Hájek points out, (Ref $\cdot$) is the very property that allows the Wager to work independently of the probability $q$ attached to God's existence, so long as this number is non-zero.

In summary, insofar as the axioms supporting *Premise 3* are incompatible with infinite utility values, the Wager's appeal to that premise is undercut. McLennen leaves open the possibility of identifying, and justifying, some alternative principle in a framework that admits of infinitely valued outcomes.

## 2.2. *The Invalidity Objection*

Following Jeffrey (1983) and Duff (1986), Hájek (2003) argues convincingly that Pascal's argument as presented in section 1 is invalid. Any mixed strategy that results in a positive probability for your eventually coming to wager for God yields maximal (infinite) expected utility. The requirement to maximize expected utility does not distinguish between wagering for God and flipping a coin to decide whether or not to wager for God, since both acts have infinite expected utility. Hájek goes further: you can do your very best to wager *against* God. So long as there is still a positive probability you might wind up wagering *for* God, your expected utility is still infinite. The source of trouble here is once again the reflexivity of infinity under multiplication. Just as the Wager is insensitive to the probability of God's existence, it is insensitive to the probability that one ends up coming to wager in favour of God.

In response to this objection, Schlesinger (1994) has proposed the following principle:

> In cases where two acts yield distinct probabilities for the same prize (or prizes of equal value), we ought to prefer the act associated with the higher probability.[7]

The principle is clearly correct for prizes of finite utility – in fact, it is just the Better-chances condition (A3). Schlesinger's plausible suggestion is that this should be extended to infinite prizes. No need for much deliberation to choose between an act that gives us a 99.9999% chance to win an infinite prize and one that gives us a 0.0001% chance for the same prize! Although both possess infinite expected utility, the choice is obvious. If Schlesinger's Principle is granted, the Wager becomes valid.

Following both Sorensen (1994) and McClennen (1994), however, Schlesinger's amendment appears *ad hoc* unless it can be embedded within a systematic framework of assumptions about preferences. Furthermore, as we have seen, it is not clear how we can distinguish preferentially among outcomes with equal (infinite) expected utility without giving up the virtues of traditional utility theory.

## 3. TWO REFORMULATIONS AND HÁJEK'S DILEMMA

A natural response to the objection of Section 2.2 is to try to reformulate the Wager in some way that makes it valid. There are a number of ways to go about doing this. Before considering them, it is helpful to take a closer look at the way we have represented the value of infinity in the Wager.

The setting that seems best to reflect Pascal's intuitions about infinity is the *extended real numbers*.[8] Typically, this set is defined by adding two elements to the system of real numbers: $\infty$ and $-\infty$. We extend the ordering relation by postulating

$$-\infty < x < \infty \quad \text{for each real number } x.$$

We extend the usual arithmetic operations by postulating that for each real number $x$,

$$
\begin{array}{ll}
(\text{Ref } +) & x + \infty = \infty \\
(\text{Ref } \cdot) & x \cdot \infty = \infty \text{ if } x > 0 \\
 & x \cdot -\infty = -\infty \text{ if } x > 0
\end{array}
$$

and

$$
\begin{array}{cc}
\infty + \infty = \infty, & -\infty + -\infty = -\infty \\
\infty \cdot (\pm\infty) = \pm\infty, & -\infty \cdot (\pm\infty) = \mp\infty
\end{array}
$$

We leave $\infty - \infty$ undefined, but set

$$0 \cdot \infty = 0.$$

We have already followed Hájek in calling the second postulate *reflexivity under multiplication*. Again adopting Hájek's terminology, we call the first postulate *reflexivity under addition*.

These postulates provide a simple and natural way to model the relationship of the infinite to the finite, and one that is congenial to Pascal's way of thinking.[9] Of course, we know now that there is no way to model the Wager directly using this number system. As we saw in Section 2, that approach is bound to make the Wager invalid. This observation has led to a number of reformulations that restore validity. Hájek (2003) proposes two criteria that any such reformulation should satisfy.

*Overriding Utility*. The utility of salvation must override any other utility that enters into the expected utility calculations (rendering irrelevant the exact value of the probability $q$ assigned to God's existence).[10]

*Distinguishable Expectations*. The smaller the probability of wagering for God associated with a strategy, the smaller should be the expectation for that strategy.

The first requirement ensures fidelity to Pascal's original argument. So long as the probability of God's existence is positive, the expected utility of salvation will swamp any of the other utilities in the decision matrix. The second requirement incorporates the idea behind Schlesinger's Principle, and involves a partial relaxation of (Ref $\cdot$ ): we cannot simply equate $p \cdot \infty$ with $q \cdot \infty$ where $p$ and $q$ are distinct values between 0 and 1.

Hájek goes on to discuss four reformulations of the Wager. Let us briefly examine two of them. The first makes use of Conway's (1976) *surreal numbers*, which bear some resemblance to non-standard or hyperreal numbers. Without delving into Conway's construction, the idea is to replace $\infty$ with $\omega$ in Table I, where $\omega$ is a particular infinite number. Addition, multiplication and the order relation are well defined in this system, and if we define expected utility by the usual formula, the expected utility of wagering for God is

$$q \cdot \omega + (1-q) f_2$$

which is infinite. This exceeds not only the expected utility of wagering against God, but also the expected utility of any mixed strategy, which will have the form

$$p \cdot \omega + k$$

where $k$ is finite and $p < q$. In Conway's system, $p < q$ implies $p \cdot \omega < q \cdot \omega$ for any infinite number, and hence the expected utility is lower with the mixed strategy. Both requirements, *Overriding Utility* and *Distinguishable Expectations*, are satisfied.[11] Similar reformulations can be given by replacing $\infty$ with $\aleph_0$, or with an infinite hyperreal number (Sobel 1996).

The second approach discussed by Hájek replaces one-dimensional utilities with two-dimensional vector-valued utilities of the form (earthly value, heavenly value). Outcomes thus have a utility $(x, y)$, where $x$ represents earthly goods and $y$ represents heavenly goods. While $x$ can be unbounded, Hájek assumes $0 \leq y \leq 1$, with the maximal value 1 reflecting salvation. The ordering on these utilities is lexicographic, with heavenly value being the dominant dimension:

$$(x, y) \leq (x', y') \leftrightarrow y < y' \quad \text{or} \quad (y = y' \;\&\; x \leq x')$$

Any small increment in heavenly value is better than an arbitrarily large increment in earthly value.

If we modify the decision matrix of Table I by replacing $\infty$ with $(f_1, 1)$ and $f_i$ by $(f_i, 0)$ for $i = 2$, 3 and 4, and if we define expected utility by the usual formula, the expected utility of wagering for God (where $q$ is the probability that God exists) is

$$(k, q)$$

for some finite $k$. Not only does this exceed the expected utility of wagering against God, but also the expected utility of any mixed strategy, which will be

$$(l, pq)$$

for some finite $l$.[12] Since $p < 1$, we have $pq < q$. Once again, both *Overriding Utility* and *Distinguishable Expectations* are satisfied.[13]

Both reformulations result in a valid argument for wagering in favour of God. Yet Hájek maintains that both are unsatisfactory, both as attempts to represent Pascal's reasoning and on independent philosophical grounds. Pascal believed that the utility of salvation was *absolutely infinite*, in no way to be surpassed. This is reflected in the postulate of reflexivity under addition: for any $p > 0$, $\infty + p = \infty$. Each reformulation discussed by Hájek fails in some way to respect this Pascalian intuition. If one takes $\omega$ for the utility of salvation in Conway's system, $\omega + p$ is greater.[14] For the two-dimensional model, the problem is less clear because the value of salvation is represented as $(f_1, 1)$, and $(f_1, 1) + p$ is undefined. Nevertheless, analogues of (Ref +) fail. For instance, why should the heavenly reward be $(f_1, 1)$ rather than $(f_1+1, 1)$? More importantly (in Hájek's view), why should there be only two dimensions of utility rather than three or more? The general problem is as follows. Technical devices permit a value for the utility of salvation that will swamp all other terms in the decision matrix (to satisfy *Overriding Utility*) and ensure that wagering dominates mixed strategies (to satisfy *Distinguishable Expectations*), but these devices also ensure that ever larger utility values are conceivable (violating (Ref +)).

It is tempting to respond by distinguishing between the *practical* and *logical* unsurpassability of the utility of salvation. In the reformulations of the Wager, the utility value attached to salvation fails to be logically unsurpassable because larger numbers exist. That value may nevertheless be practically unsurpassable because no other possible outcome has a higher utility value.

I side with Hájek in rejecting this response as unfaithful to Pascal's intuitions about the absolute infinity of the value of salvation. Furthermore, as Hájek notes, there are philosophical objections. Why should human utility be capped at some arbitrary number (such as $\omega$) when there is plainly something better available (such as $\omega + 1$)? Even if introducing surreal or vector utility values allows us to formulate a valid version of the Wager, what do such values mean? We can interpret real-valued utilities in terms of preferences for gambles, but such a straightforward approach is not available for vector-valued or surreal utilities. In addition, considerations of simplicity suggest that if we can defend the Wager without introducing non-standard utility values, we should do so.

Hájek concludes by posing a dilemma for any attempted reformulation of the Wager:

> If the utility of salvation is *both* reflexive under addition and under multiplication by positive, finite probabilities, as $\infty$ is, then the argument is invalid. If the utility of salvation is *neither* reflexive under addition nor under multiplication by positive, finite probabilities, as is the case with the reformulations, then salvation is so far from being the best thing possible that its utility is swamped by something that is swamped by something that is swamped … infinitely many times over. What is wanted, then, is the seemingly impossible: a representation of the reward of salvation that is reflexive under addition (so that it cannot be bettered), but *not* reflexive under multiplication by positive, finite probabilities (so that the mixed strategies can be distinguished in expectation from outright wagering for God). (47–8)

Hájek's reasoning is perfectly correct if we limit ourselves to representing preferences with a one-place utility function. The remainder of this paper is devoted to showing that we need not restrict ourselves in this way. With a three-place utility function, we can meet all of Hájek's desiderata. In effect, we do not assign utility values, but only utility ratios. We can then construct a representation of the Wager that uses nothing more complicated than the extended real numbers. The result is, I hope, simple enough to count as faithful to Pascal's intentions. After developing the apparatus and applying it to the Wager, I shall return (in section 5) to the question of whether we have escaped Hájek's dilemma.

## 4. RELATIVE UTILITIES

### 4.1. *Introducing Relative Utilities*

Pascal's Wager involves a contrast between infinite and finite values. Inevitably, that means non-Archimedean preferences, i.e., a preference ordering that violates axiom (A5). We have seen two ways of representing non-Archimedean preferences: a unidimensional representation using surreal or non-standard numbers, and a multidimensional representation employing a lexicographic ordering on ordered pairs. Utility functions that represent non-Archimedean preferences are generally of these two types.[15]

Hájek's dilemma can be expressed in the following way. The simplest non-Archimedean representation, using the extended real numbers, fails because of the invalidity objection. Any representation using a richer set of utility values fails because no matter what value we designate as the utility of salvation, it can be surpassed.

The strategy in this section and the next is to escape this dilemma by moving from direct representation of preferences with a utility value to indirect representation with utility ratios. These ratios contain all the information needed to present the argument of the Wager using nothing more complicated than the non-negative extended real numbers. To appreciate how this works, let us begin by considering how utility ratios behave in the ordinary case of Archimedean preferences.

Suppose, then, that our preferences satisfy all the axioms of Section 2.1 including (A5). By the *Expected Utility Theorem*, we can represent them with a real-valued utility function $u$, unique up to a positive linear transformation. Any other utility function that preserves our preference ranking and has the linearity property is of the form $au+b$, where $a>0$. In addition to the standard axioms, we suppose that among the relevant outcomes there happens to be a worst outcome, $W$. (This simplifying assumption will be dropped shortly.) For any two outcomes $A$ and $B$ (including gambles), define the *relative utility* of $A$ to $B$ by

$$\mathcal{U}(A,B)=\begin{cases}[u(A)-u(W)]/[u(B)-u(W)] & \text{if } u(B)\neq u(W)\\ 1 & \text{if } u(A)=u(B)=u(W)\\ \infty & \text{if } u(A)\neq u(W) \text{ and}\\ & \quad u(B)=u(W).\end{cases}$$

The value $\mathcal{U}(A,B)$ is independent of the choice of utility function. If $u'=au+b$, then the constants $a$ and $b$ cancel out. So we can simplify matters by assuming $u(W)=0$.

Since $u$ represents our preferences, it follows that $A\preccurlyeq B$ iff $u(A)\leq u(B)$, with strict preference corresponding to a strict inequality. Thus we have

$$B\preccurlyeq A\leftrightarrow\mathcal{U}(A,B)\geq 1,$$

again with strict preference corresponding to strict inequality. Furthermore, from the linearity of $u$, we have for any outcomes $A$, $A'$, $B$:

$$\mathcal{U}([pA,(1-p)A'],B]=p\mathcal{U}(A,B)+(1-p)\mathcal{U}(A',B).$$

This holds even if $u(B) = u(W) = 0$. So relative utility is linear (in the first component) over gambles.

It follows that maximizing expected utility corresponds to choosing an action whose *expected relative utility* with respect to each other available action is maximal. This criterion extends to mixed strategies. Note also that we can recover a utility function $u$ from $\mathcal{U}$ by arbitrarily setting $u(A) = 1$ for any outcome $A$ strictly preferred to the worst outcome $W$, and then setting $u(B) = \mathcal{U}(B, A)$ for all $B$. Nothing is lost, then, by representing our preferences in terms of relative utilities instead of utilities. Of course, there is no gain either.

The two approaches to representing preferences come apart when we move to a non-Archimedean setting. One-place utility functions give us no simple way to represent non-Archimedean preferences. We have to resort to constructions such as those outlined in Section 3. Relative utilities offer a more conservative way to extend decision theory to accommodate the notion of infinite value.

As a first step, we need to remove the restrictive assumption that there is a fixed worst outcome $W$. Relative utility is defined as a three-place function $\mathcal{U}(A, B; Z)$: the utility of $A$ relative to $B$ with base-point $Z$. By analogy with the ratios above, we shall refer to the argument $A$ as the *numerator* and $B$ as the *denominator*. The *base-point* $Z$ may be thought of as the designated origin or zero point, the location where we plant our 'measuring stick' to determine the ratio of the 'distance' to $A$ over the 'distance' to $B$. The only constraint will be that $Z$ must rank no higher in our preferences than either $A$ or $B$. So $\mathcal{U}(A, B; Z)$ will always be non-negative.

The shift from a one-place to a three-place utility function is merely a shift in representation. It has no metaphysical significance; preference remains a two-place relation. All of the objections to the Wager discussed in Section 2 are problems about representation. We are already in a position to see why a move to three-place utility functions might be useful for thinking about the non-Archimedean preferences that define Pascal's Wager.

The main advantage is that the shift allows us to formulate a perfectly clear conception of infinite utility – or rather, of infinite *relative* utility. While infinite utility considered absolutely is a vexing notion, there is a straightforward definition of what it means for relative utility to be infinite. Indeed, Pascal shows us the way. My utility for $A$ is infinite relative to $B$ (with base-point $Z$) just in case I prefer *any* non-trivial gamble between $A$ and $Z$ to $B$. I am willing to sacrifice $B$ for any chance of obtaining $A$ over $Z$, no matter how slight. Formally:

DEFINITION 1 (INFINITE RELATIVE UTILITY). Suppose $Z \prec B$ and $B \prec A$. Then

$$\mathcal{U}(A, B; Z) = \infty \quad \leftrightarrow \quad B \prec [pA, (1-p)Z] \text{ for all } 0 < p \leq 1.$$

No notions of infinity are needed in the *definiens*. *Definition 1* simply employs the familiar idea of defining utilities in terms of preferences between gambles.[16]

The full definition of relative utility is given in the next section, but it is worth pausing to consider *Definition 1* more closely. The definition takes a structured set of preferences, rather than a utility function capable of taking the value $\infty$, as the starting point. Adopting this approach to infinite value compels us to view Pascal's argument somewhat differently from its presentation in Section 1. There, the argument ran:

1. Because salvation has infinite utility, the act of wagering for God has infinite expected utility. Therefore, wagering dominates any outcome with finite utility.
2. Not wagering for God has merely finite expected utility.
3. We ought to maximize expected utility.

Therefore, we should wager for God.

On the proposed understanding of infinite utility, steps 2 and 3 remain the same but step 1 gets things backwards. To assign salvation infinite utility (relative to earthly outcomes) is to *pre-suppose* 'Pascalian' preferences that rate any gamble with a finite, positive chance for salvation above any finite prize.[17] That wagering for God dominates any worldly prize is not the penultimate step in the argument. It is the starting point. The challenge for Pascal is not to convince the rest of us to take the wager, but rather to *rationalize* his own preferences. Pascalians need to find a utility function that can represent their preferences without succumbing to the objections raised in Section 2. In particular, they need to preserve some version of the criterion of maximizing expected utility.

In Section 5, we show that we can represent Pascalian preferences using a relative utility function $\mathcal{U}$, taken as primitive rather than defined via utility ratios as above. This representation incorporates

Hájek's crucial requirements, *Overriding Utility* and *Distinguishable Expectations*. Any mixed strategy giving us a shot at salvation will have infinite utility relative to any mere worldly prize, but only the pure wager will have maximal expected relative utility. Our relative utility function will distinguish between the pure wager and the coin toss.

### 4.2. *Preferences and Relative Utility*

In this section, I define relative utility and state an important result (*Theorem 6*): the relative utility of a gamble is its expected relative utility. In the next section, I state two corollaries that yield tests analogous to the usual criterion of maximizing expected utility. All proofs are deferred to Appendix A.

Begin by supposing that our preference ordering $\preccurlyeq$ satisfies the basic assumptions (A1)–(A4) of Section 2.1, as well as a non-triviality condition that there are at least two outcomes $A$ and $B$ such that not $A \sim B$. For any $P$ and $R$, define the *preference interval* $PR = \{Q / P \preccurlyeq Q \preccurlyeq R\}$. $PR$ is the set of all outcomes intermediate between $P$ and $R$. Provided $P \preccurlyeq R$, this set is nonempty and includes its endpoints $P$ and $R$. We further restrict our attention to non-empty intervals $PR$ between whose endpoints we are not indifferent, i.e., $P \preccurlyeq R$ but not $P \sim R$. First, we have a preliminary result.[18]

LEMMA 2 ($\alpha$-VALUE LEMMA). If $Q \in PR$, then there is a unique $\alpha$, $0 \leq \alpha \leq 1$, such that one of the following three cases holds:

(a) $Q \sim [\alpha R, (1-\alpha)P]$;
(b) $[kR, (1-k)P] \preccurlyeq Q$ for $0 \leq k < \alpha$ and $Q \preccurlyeq [kR, (1-k)P]$ for $\alpha \leq k \leq 1$, with $\alpha > 0$;
(c) $[kR, (1-k)P] \preccurlyeq Q$ for $0 \leq k \leq \alpha$ and $Q \preccurlyeq [kR, (1-k)P]$ for $\alpha < k \leq 1$, with $\alpha < 1$.

If we fix $PR$, Lemma 2 tells us that there we can assign a unique $\alpha$ to each $Q \in PR$, indicating how far $Q$ lies along the interval between $P$ and $R$.

Making use of the $\alpha$-value lemma, it is helpful to define a relation of *indifference relative to base-point* $Z$. An agent is *relatively indifferent* between $A$ and $B$ with base-point $Z$, written $A \approx_Z B$, if one of the following conditions holds:

(a) $B \preccurlyeq A$ and the $\alpha$-value of $B$ in $ZA$ is 1; or
(b) $A \preccurlyeq B$ and the $\alpha$-value of $A$ in $ZB$ is 1.

$\approx_Z$ is an equivalence relation. Observe that in case (a), we might strictly prefer $A$ to $B$ but still prefer $B$ to any non-trivial gamble between $A$ and $Z$. So $A \sim B$ implies $A \approx_Z B$ for any choice of base-point $Z$, but the reverse implication does not hold.

The following result is a straightforward consequence of Lemma 2.

LEMMA 3. If $Z \preccurlyeq A \preccurlyeq B$, then exactly one of three cases holds:

(i) The $\alpha$-value of $B$ in $ZA$ is 1: then $A \approx_Z B$.
(ii) The $\alpha$-value of $B$ in $ZA$ is strictly between 0 and 1: then $B \approx_Z [\alpha A, (1-\alpha)Z]$.
(iii) The $\alpha$-value of $B$ in $ZA$ is 0: then $B \preccurlyeq [pA, (1-p)Z]$ for all $0 < p \leq 1$.

A final preliminary fact: substituting relatively indifferent outcomes into gambles preserves the relationship of relative indifference.

LEMMA 4 (SUBSTITUTION LEMMA). If $B \approx_Z B'$, then for any $0 \leq p \leq 1$ and any $A$, $C$,

$$[pB, (1-p)C] \approx_Z [pB', (1-p)C] \text{ and}$$
$$[pA, (1-p)B] \approx_Z [pA, (1-p)B'].$$

We can now give the definition of relative utility and state the main result of this section.

DEFINITION 5 (RELATIVE UTILITY, $\mathcal{U}(A, B; Z)$). Recall that by (A1), for any $A$ and $B$, either $A \sim B$, $A \prec B$ or $B \prec A$.

If $A \sim B$, set $\mathcal{U}(A, B; Z) = 1$.
If $A \prec B$, set $\mathcal{U}(A, B; Z) = \alpha$ where $\alpha$ is the $\alpha$-value of $A$ in $ZB$.
If $B \prec A$, set $\mathcal{U}(A, B; Z) = 1/\alpha$ where $\alpha$ is the $\alpha$-value of $B$ in $ZA$ (taking $1/0 = \infty$).

THOEREM 6. If (A1)–(A4) are satisfied, then there exists a unique three-place relative utility function $\mathcal{U}$ with the following properties (where $A$, $B$ and $C$ are any outcomes, and $Z$ is always dominated by all other relevant outcomes):

(R1) $\mathcal{U}(A, B; Z)$ is a non-negative extended real number.
(R2) If $\mathcal{U}(A, B; Z) > 1$, then $B \prec A$.
If $\mathcal{U}(A, B; Z) < 1$, then $A \prec B$.
$\mathcal{U}(A, B; Z) = 1$ iff $A \approx_Z B$.
(R3) $\mathcal{U}([pA, (1-p)A'], B; Z) = p\mathcal{U}(A, B; Z) + (1-p)\mathcal{U}(A', B; Z)$ for $0 \leq p \leq 1$.
(R4) $\mathcal{U}(Z, A; Z) = 0$ unless $Z \sim A$ (in which case $\mathcal{U}(Z, A; Z) = 1$).
(R5) $\mathcal{U}(A, B; Z) = 1 / \mathcal{U}(B, A; Z)$ (where $1/0 = \infty$ and $1/\infty = 0$).
(R6) If $B \preccurlyeq A$, then for any $C$, $\mathcal{U}(A, C; Z) \geq \mathcal{U}(B, C; Z)$.
(R7) If $B \preccurlyeq A$ and $C \preccurlyeq A$ and $\mathcal{U}(C, A; Z) > 0$, then
$\mathcal{U}(C, B; Z) = \mathcal{U}(C, A; Z) \cdot \mathcal{U}(A, B; Z)$.

(R1) tells us that we can represent relative utilities using the non-negative extended real numbers. The other properties tell us that $\mathcal{U}(A, B; Z)$ behaves just like the utility ratio function described in Section 4.1. In particular, (R3) states that $\mathcal{U}(A, B; Z)$ is linear in the first component. So we have a version of the expected utility theorem, even though our set of allowed values now includes $\infty$.

To illustrate how relative utilities are evaluated, let us return to Pascal's Wager. Our original representation was the two-by-two table reproduced in Table II.

We have added labels for the four different outcomes, indexed by row and column. The first of these, $O_{11}$, represents salvation, to which we attached the value $\infty$. The other three outcomes are worldly prizes with finite utility values. We assume here that the finite values are assigned on an interval scale – i.e., preferences among worldly prizes satisfy all of axioms (A1)–(A5) and the utility function is linear among these worldly prizes. Further, we may assume that all utility values in the table are positive and that there is an outcome $Z$ assigned zero utility, which we take as our base-point.[19]

TABLE II

Pascal's Wager

| | God exists | God does not exist |
|---|---|---|
| Wager for God | $O_{11}(\infty)$ | $O_{12}(f_2)$ |
| Wager against God | $O_{21}$ $(f_3)$ | $O_{22}$ $(f_4)$ |

Relative utilities among the finite prizes can be calculated as ratios, just as in Section 4.1.

$$\mathcal{U}(O_{12}, O_{22}; Z) = f_2/f_4$$
$$\mathcal{U}(O_{21}, O_{12}; Z) = f_3/f_2$$

and so forth. This is a straightforward application of Definition 5. For instance, supposing $f_2 < f_4$, linearity implies that we are indifferent between $O_{12}$ and the gamble $[(f_2/f_4)O_{22}, (1 - f_2/f_4)Z]$.

As we shall see in the next section, relative utilities whose denominator is an optimal outcome for the decision problem are especially significant. The unique optimal outcome here is $O_{11}$. Using Definition 5 again, we obtain:

$$\mathcal{U}(O_{11}, O_{11}; Z) = 1$$
$$\mathcal{U}(O_{12}, O_{11}; Z) = 0 \quad (\text{and} \quad \mathcal{U}(O_{11}, O_{12}; Z) = \infty)$$
$$\mathcal{U}(O_{21}, O_{11}; Z) = 0 \quad (\text{and} \quad \mathcal{U}(O_{11}, O_{21}; Z) = \infty)$$
$$\mathcal{U}(O_{22}, O_{11}; Z) = 0 \quad (\text{and} \quad \mathcal{U}(O_{11}, O_{22}; Z) = \infty)$$

Finally, consider a gamble such as $[qO_{11}, (1 - q)O_{12}]$, which represents the act of wagering for God (where $q$ is the probability that God exists). We might expect that its utility relative to $O_{11}$ is $q$, and this is indeed the case. By the linearity property (R3), we have

$$\mathcal{U}([qO_{11}, (1-q)O_{12}], O_{11}; Z)$$
$$= q\mathcal{U}(O_{11}, O_{11}; Z) + (1-q)\mathcal{U}(O_{12}, O_{11}; Z) = q.$$

### 4.3. *Expected Relative Utility and Optimal Strategies*

Just as standard linearity leads to the criterion of expected utility maximization, the linearity property (R3) leads to an analogous requirement that we maximize expected relative utility. In standard decision theory, maximization gives us a simple way to identify the optimal action or strategy, but as yet we have no equally simple test for relative utilities. The main result of this section is that there is such a test (with certain limitations), provided our decision problem is finite (i.e., there are finitely many pure actions and finitely many possible states). In what follows, we regard an action or strategy (pure or mixed) as a gamble over outcomes and identify its utility as the utility of this gamble.

Consider a decision problem involving pure actions (or strategies) $A_1, \ldots, A_n$ and mutually exclusive possible states $S_1, \ldots, S_m$. Write $O_{ij}$ for the outcome associated with action $A_i$ and state $S_j$. There are $m \cdot n$ of these possible outcomes in our decision problem. Let $q_j$ be the probability of $S_j$, and assume each $q_j > 0$.[20] Let us begin by recalling some basic facts from ordinary decision theory. Assume the utility function $u$ is linear. The expected utility of action $A$ is the weighted sum of the utilities it yields in each of the $m$ possible states, i.e., $\Sigma q_j u(A\&S_j)$. If $A$ is a mixed strategy of the form $[p_1 A_1, \ldots, p_n A_n]$ where $\Sigma p_i = 1$, then its expected utility is $\Sigma q_j \Sigma p_i u(O_{ij})$. Any action that maximizes this quantity is optimal.

We would like to find an analogous principle for relative utilities. Just as in the ordinary case, we want to define a function that assigns a number to each possible action $A$, and we want maximization of this function to correspond to the optimality of $A$. Relative utility is a three-place function $\mathcal{U}(A, B; Z)$, however, so we shall need to fix a denominator $B$ and a base-point $Z$. The choice of $Z$ is discussed in later sections; here, we simply assume that $Z$ is a fixed lower bound for all possible outcomes. Indeed, the choice that makes the most sense is to take the largest possible such lower bound, which is well defined since our decision problem is finite. We then have two interesting questions. First, can we find a fixed denominator $B$ such that when we let $A$ vary, a maximal value for $\mathcal{U}(A, B; Z)$ indicates that $A$ is optimal? Second, is there an easy way to compute the values $\mathcal{U}(A, B; Z)$? A positive answer to both questions would provide a test for optimality analogous to what we have in ordinary decision theory. Corollaries 8 and 9, respectively, answer these two questions in the affirmative, though with a caveat.

To begin with, we define the *expected utility* of $A$ *relative* to $B$ (with base-point $Z$) as the weighted sum of relative utilities, $\Sigma q_j \mathcal{U}(A\&S_j,\ B;\ Z)$. (R3) tells us that this is equal to $\mathcal{U}(A, B; Z)$: the relative utility of a gamble is its expected relative utility. If $A$ is one of the pure strategies $A_i$, then this sum is equal to $\Sigma q_j \mathcal{U}(O_{ij}, B; Z)$. It can be obtained by forming a *relative decision matrix*, the matrix whose entry in row $i$ and column $j$ is $\mathcal{U}(O_{ij}, B; Z)$, and then computing the weighted sum of these values along row $i$, just as with normal expected utility calculations. In the more general case where $A$ is a mixed strategy of the form $[p_1 A_1, \ldots, p_n A_n]$ where $\Sigma p_i = 1$, we have

$$\begin{aligned}\mathcal{U}(A, B; Z) &= \Sigma q_j \mathcal{U}(A\&S_j, B; Z) && \text{by } (R3)\\ &= \Sigma q_j \Sigma p_i \mathcal{U}(A_i\&S_j, B; Z) && \text{by } (R3)\\ &= \Sigma q_j \Sigma p_i \mathcal{U}(O_{ij}, B; Z) && \text{since } O_{ij} = A_i\&S_j.\end{aligned}$$

All of this is directly analogous to the case of a one-place utility function.

Our first result is a simple consequence of this equality.

LEMMA 7. Suppose $\mathcal{U}(A, O_{ij}; Z) = \mathcal{U}(B, O_{ij}; Z)$ for all $i, j$. Then $\mathcal{U}(A, B; Z) = 1$.

Relative to $Z$, we are indifferent between any two actions that have the same set of utilities relative to the possible outcomes $O_{ij}$.

Call an action $A$ *Z-optimal* if for any $B$, either $B \preccurlyeq A$ or $A \approx_Z B$. The following corollaries show that we can find a $Z$-optimal action by maximizing expected utility relative to (1) all possible outcomes $O_{ij}$, or more simply, relative to (2) any single optimal outcome in this set.

COROLLARY 8.

(8.1) An action $A$ is $Z$-optimal if and only if $\mathcal{U}(A, O_{ij}; Z)$ is maximal for each outcome $O_{ij}$: $\mathcal{U}(A, O_{ij}; Z) \geq \mathcal{U}(B, O_{ij}; Z)$ for all actions $B$.

(8.2) $A$ is $Z$-optimal if and only if $\mathcal{U}(A, O; Z)$ is maximal for any outcome $O$ that is optimal among the $O_{ij}$: $\mathcal{U}(A, O; Z) \geq \mathcal{U}(B, O; Z)$ for all $B$.

This result immediately yields a test for finding $Z$-optimal actions. The key idea, expressed in (9.2), is to compute the appropriate relative decision matrix. From that point on, the calculations are just the same as in ordinary decision theory.

COROLLARY 9. (Z-OPTIMALITY = MAXIMAL EXPECTED RELATIVE UTILITY).

(9.1) For each possible outcome $O_{kl}$, form the relative decision matrix whose entry in row $i$ and column $j$ is $\mathcal{U}(O_{ij}, O_{kl}; Z)$. Calculate expected relative utilities. $A$ is $Z$-optimal if and only if in every case (i.e., for each $O_{kl}$), $A$ has maximal expected relative utility.

(9.2) Let $O$ be an optimal outcome, and form the relative decision matrix with entries $\mathcal{U}(O_{ij}, O; Z)$. Calculate expected relative utilities. $A$ is $Z$-optimal if and only if $A$ has maximal expected relative utility.

In each of these two corollaries, part 1 shows that maximizing utility relative to each of the $m \cdot n$ possible outcomes is good enough to secure $Z$-optimality. Part 2 shows that by maximizing utility relative to just *one* of these outcomes, so long as it is optimal, we automatically maximize it relative to all of them. Corollary 9.2 enables us to do everything with just one relative decision matrix in which all values will be between 0 and 1.[21] We use 9.2 exclusively in the remainder of this paper. A good slogan might be: maximize expected utility relative to the best outcome. The next section illustrates the principle with several examples.

We have the criterion we sought, the analogue for the familiar principle of maximizing expected utility. It reduces to that principle if our preferences are Archimedean: the relative decision matrix is what we get when we scale our utility function to assign 0 to $Z$ and 1 to the best possible outcome. But we now have a more general principle that lets us handle non-Archimedean preferences.

Here is the caveat: our test identifies not optimality, but $Z$-optimality. We pick out a class of outcomes or actions between which we are relatively indifferent from the perspective of the base-point $Z$. This is not a sufficient condition for optimality. $Z$-optimality (for any choice of $Z$) is, however, a necessary condition for optimality. Furthermore, when we apply the test to Pascal's Wager and other puzzles below, there is frequently just one $Z$-optimal action, which is therefore truly optimal. In Section 6, we refine this test to handle cases where we encounter more than one $Z$-optimal action.

## 5. RELATIVE UTILITY AND PASCAL'S WAGER

In our original two-by-two decision matrix for the Wager (Tables I and II), we saw that the infinite value attached to salvation was problematic. The reformulation in terms of relative utilities, however, is quite simple. First, we know that the unique optimal outcome is $O_{11}$ in the top-left corner: salvation. So we can apply Corollary 9.2. As in Section 4.2, we may assume a base-point $Z$ with utility 0.[22] As we showed in Section 4.2, we have $\mathcal{U}(O_{ij}, O_{11}; Z) = 0$ for all outcomes $O_{ij}$ except salvation, and $\mathcal{U}(O_{11}, O_{11}; Z) = 1$. So the relative decision matrix is just this:

Since we have assumed $q > 0$, the expected relative utility of Wagering-for is $q$, while that of Wagering-against is 0. Hence, given only the choice between the two pure strategies, we ought to wager for God. No other action is $Z$-optimal.

The same table can be used to handle mixed strategies of the form $W_p = [p \text{ (Wager-for)}, (1 - p) \text{ (Wager-against)}]$. By linearity, the expected relative utility of $W_p$ is $pq$, which is less than that of the pure wager unless $p = 1$. So even when we allow mixed strategies, the pure wager is the unique $Z$-optimal choice and hence the optimal choice. This answers the invalidity objection of Section 2.2. Schlesinger's response is essentially correct. We have simply provided a systematic way to incorporate his principle into decision theory. We ought to prefer the gamble giving us the higher probability for a fixed, infinite prize, and we can make sense of this by means of expected relative utility calculations no more difficult than ordinary expected utility calculations.

It can easily seem, however, that Table III is a gross misrepresentation of the Wager. In one sense, the table nicely captures Pascal's dictum: 'the finite is annihilated in the presence of the infinite and becomes pure nothingness.'[23]. This is accomplished by the assignment of 0 to every outcome other than salvation. What is troublesome, though, is that Table III fails to preserve distinctions among the earthly rewards. The table is a kind of projection of data about our preferences onto the highest plateau; inevitably, information about our lower-order preferences gets lost. In discussing a somewhat similar proposal by Swinburne (1969) that assigns utilities 1 and 0 on row 1 and −1 and 0 along row 2, Sobel (1996, 19) writes:

Pascal's comparative evaluations cannot be measured on a bounded scale. It is sufficient to note that Pascal differentiates amongst ordinary worldly valuables, holds that the infinite quite annihilates the finite, and says that highest eternal valuables are infinite. This means that on a bounded scale his ordinary valuables, in order to fix their relations to highest valuables, would all have measure 0, which, however, would misrepresent their relations amongst themselves.

TABLE III

Pascal's Wager (utilities relative to best outcome)

| | $q$<br>God exists | $1-q$<br>God does not exist |
|---|---|---|
| Wager for God | 1 | 0 |
| Wager against God | 0 | 0 |

TABLE IV

Pascal's Wager (relative to happy secular outcome)

| | God exists | God does not exist |
|---|---|---|
| Wager for God | $\infty$ | $f_2/f_4$ |
| Wager against God | $f_3/f_4$ | 1 |

When worldly prizes are stacked up against the standard of infinity, it becomes impossible to differentiate between them. Sobel's point is valid if our intention is to locate all the outcomes in the Wager along a single, unidimensional utility scale. This difficulty vanishes, though, when we move to relative utilities. All the information we might desire about the ranking of earthly prizes can be recovered by looking at relative utilities among those prizes.

Consider, for instance, the relative decision matrix with entries $\mathcal{U}(O_{ij}, O_{22}; Z)$ relative to the happy secular outcome (wager against, no God), as given in Table IV.

Once again, we can calculate expected relative utilities. In this case, the calculations look exactly like the standard analysis of Pascal's Wager, as presented in Sections 1 and 2. Just as in that discussion, this table demonstrates the superiority of Wagering-for to Wagering-against but fails to show why the pure Wager is better than any mixed strategy. By considering relative utilities with respect to a non-optimal outcome, we come up with a less informative answer about the optimal choice – even though it is more informative about the ranking of the finite outcomes. Corollary 9.2 assures us that Table III, not Table IV, is the appropriate choice when our objective is to identify an optimal action.

We have seen how to respond to the invalidity objection of Section 2.2. What of McClennen's incompatibility objection, raised in Section 2.1? We have already dealt with his main concern. We must give up *Premise 3*, the familiar principle of maximizing expected utility, but in its place we have an analogous principle: maximize expected utility relative to the best outcome. This principle is compatible with assigning some prizes infinite utility relative to others. Furthermore, we can retain all the standard axioms (A1)–(A4), giving up only (A5), *Continuity*. The violations of the *Better-prizes condition* and *Better-chances condition* discussed in Section 2.1 disappear.[24]

Before claiming that we have an acceptable reformulation of Pascal's argument, let us return to Hájek's dilemma. It is clear that our version of the Wager satisfies *Overriding Utility*. Tables III and IV above preserve the idea that the utility of salvation relative to any other pure outcome is infinite. Our version of the argument also satisfies *Distinguishable Expectations*: the smaller the probability of wagering for God associated with a strategy, the smaller the (relative) expectation for that strategy. We have the resources to distinguish between the pure wager and mixed strategies.

What of the reflexivity conditions? Recall Hájek's desideratum:

> What is wanted, then, is the seemingly impossible: a representation of the reward of salvation that is reflexive under addition (so that it cannot be bettered), but *not* reflexive under multiplication by positive, finite probabilities (so that the mixed strategies can be distinguished in expectation from outright wagering for God).

With relative utilities, the reward of salvation is not reflexive under multiplication by positive, finite probabilities. Reflexivity under addition merits careful discussion, though, since the reformulations considered by Hájek founder on this point.

In one sense, reflexivity under addition is trivially satisfied. If $O_{11}$ represents salvation and $E$ is any earthly reward, we have $\mathcal{U}(O_{11}, E; Z) = \infty$. There can be no outcome $O$ for which $\mathcal{U}(O, E; Z) = \infty + 1$ because we are using the extended real numbers. In fact, $\mathcal{U}(O_{11}, O; Z) \geq 1$ for every possible outcome $O$: salvation cannot be bettered. So there is no obvious failure of reflexivity under addition.

Going beyond this purely technical answer to Hájek, we need to consider whether there is some more general philosophical failing in our formulation of the Wager. Hájek's real problem with the two reformulations of Section 3 is that they fail to represent salvation as absolutely maximal: "salvation is so far from being the best thing possible that it is swamped by something that is swamped by something … infinitely many times over." This problem appears to be inevitable if we use an unbounded unidimensional utility scale (such as the hyperreal or surreal numbers), or a finite-dimensional lexicographic utility function. We have to make an arbitrary choice for the utility of salvation – an arbitrary value on the one-dimensional scale, or an arbitrary number of dimensions. We have no response to the question, "why stop there?"

Against the present proposal, this sort of objection might take the form: why can't there be something whose utility relative to salvation is $\infty$? We have earthly prizes and we have a heavenly reward, but what prevents an unending 'staircase' of higher and higher levels of relative utility? Just as with the alternative formulations, salvation would be swamped by any reward on one of these higher steps. So the present reformulation is just as vulnerable to the problem of arbitrary cut-off as any other version of the Wager.

This objection can be put in a different way. Suppose that our preferences regarding the Wager can actually be modelled by a one-place utility function. That is, suppose that some such model – either a unidimensional or lexicographic ordering – is equivalent to our relative utility function. We know from Section 3 that in this equivalent model, the value of salvation must be fixed at an arbitrary point along an unbounded scale. We cannot duck the problem of arbitrariness simply by refusing to settle on one of these one-place utility functions.

To the contrary, I maintain that a significant advantage of the relative utilities approach is that it remains neutral between the many alternative representations employing a one-place utility function. Because we are describing only relations of utility, there is no need to pre-suppose a definite number of dimensions or to attach a definite value to the utility of salvation. It is only once we assign such a value that we violate reflexivity under addition because greater rewards become conceivable. With the present proposal, that outcome is avoided. There might be many (even infinitely many) jumps in utility, but it is perfectly consistent to suppose that salvation is absolutely maximal.

As we noted earlier, the problem concerns the representation of our preferences, not our preferences themselves. We want a utility function that preserves the idea that salvation is absolutely the best possible outcome. We cannot accomplish this with a one-place utility function without sacrificing the validity of the Wager. Our alternative is to use a three-place function. We represent the reward of salvation not with any chosen value, but rather as infinite relative to any other possible outcome (with the exception of gambles that give us a shot at salvation).[25]

## 6. VARIATIONS ON A THEME

Relative utilities help to shed light on a traditional objection to the Wager, and to analyse two different versions of the argument.

### 6.1. *The Many-gods Objection*

In its simplest form, this famous objection runs: if Pascal's argument for infinite expectation succeeds for the Christian god, then it succeeds for any rival deity who offers an infinite reward to believers and has non-zero probability of existing. The argument proves too much! Confronted with a choice between incompatible options, each of them offering infinite expectation, there is no basis for a decision. This is the problem of Buridan's ass on a large scale: the symmetry of the situation induces paralysis.

This difficulty closely resembles Hájek's problem of mixed strategies. When competing options all have infinite expectation, there is no way to choose between them. The response we gave to that objection works just as well here. We can break the symmetry by using relative decision tables. A decision based on relative utilities recommends selecting the god with the highest probability. Table V is the relative decision table analogous to Table III, for the simplest case of two gods, $A$ and $B$.

We assume here, and throughout this section, that the infinite rewards offered by rival deities are identical (so that their relative utilities are clearly 1); this restriction will be lifted in the next section. For this example, it is clear that the expected relative utilities are $q_1$ if you wager for $A$, $q_2$ if you wager for $B$, and 0 if you wager against both. You should make your decision on the basis of your subjective probabilities.

In this way, we get beyond the paralysing symmetry of options in the standard presentation of the many-gods objection. That

TABLE V

Pascal's Wager with two gods

| | $q_1$ <br> $A$ exists | $q_2$ <br> $B$ exists | $1-(q_1+q_2)$ <br> No god exists |
|---|---|---|---|
| Wager for $A$ | 1 | 0 | 0 |
| Wager for $B$ | 0 | 1 | 0 |
| Against all | 0 | 0 | 0 |

symmetry only remains if the subjective probabilities are equal. Even in this case, atheism ('wager against all') is ruled out so long as that common probability is non-zero.

Of course, subjective probabilities differ among agents, but that is a point readily conceded both by advocates and critics of Pascal's argument. If we change our subjective probabilities enough, atheism can become the best wager even if God's existence has positive probability. We might assign high subjective probability to a god who rewards atheists. Or, following Mougin and Sober (1994), we might worry about "X-theology," which asserts that atheists go to heaven and theists go to hell regardless of whether God exists. If you assign sufficiently high positive probability to X-theology, atheism will be prudent. If the point of such stories is that some subjective probability assignments justify a decision to wager against a particular god or even against all gods, the insight is legitimate but not very troublesome for the Pascalian. Such an objection does not challenge the appropriateness of the wager for somebody who does possess the requisite subjective probabilities. Peaceful co-existence of theists and atheists is possible.

It is more interesting to see these alternative scenarios as signalling the need to justify certain subjective probabilities tacitly presupposed by the simple Wager. The beauty of that argument is that it does not depend on the probability of God's existence, so long as that is non-zero. The many-gods objection exposes the hidden assumption that any other deity, and indeed any other state of affairs that might lead to an infinite reward, has zero probability. That assumption is both unjustified and implausible. If we are prepared to assign positive probability to at least one god, why stop there?

The subjective probabilities upon which the Wager rests have not been our concern in this paper. Nevertheless, once we employ relative utilities to remove the stumbling block of "infinite expectations all-around," we open up new avenues for investigating these background probabilities and our decision of how to wager. For one thing, we can immediately lighten the burden on the Pascalian. The analysis above already shows that there is no need to defend a dogmatic assignment of zero probability to other states of affairs that offer an infinite reward. In Table V, Pascal's argument justifies wagering for the deity with highest subjective probability, even though the other deity has non-zero probability. This point clearly generalizes to decisions involving finitely many gods. The same result can be achieved by replacing the infinite reward of salvation with a large finite reward (Mougin and Sober 1994; Jordan 1998), but that drastic measure is not necessary if we employ relative probabilities.

The general point is that by representing all information using finite values, relative decision matrices allow us to apply familiar decision-theoretic tools to many-gods versions of the Wager. This opens the way to more sophisticated analyses of the many-gods objection.

### 6.2. *A Super-Pascalian Wager*

As a variation on the many-gods objection, a *super-Pascalian Wager* occurs when we have to make a decision involving more than one level of infinite relative utility. Argle offers believers *regular salvation*, an unending life whose every moment is filled with large but finite happiness. Meanwhile, Bargle offers believers *deluxe salvation*, an unending life whose every moment is filled with infinite happiness, or in Pascal's words, "an infinity of infinitely happy life". As before, assume a finite, non-zero subjective probability for the existence of each of these two deities. How are we to make our choice?

There are at least three plausible answers (apart from denying the coherence of gradations of infinite happiness). First, on the naïve view of infinite utility exemplified in the early part of this paper, we have no basis for distinguishing between the two wagers: both have infinite expectation. Second, it might strike us as sensible to wager for Bargle: the chance for infinite gain that drives the Pascalian to make the ordinary wager seems to favour the switch from Argle to Bargle, no matter how much less likely Bargle's existence might be. Third, following the pattern of Section 6.1, we might be inclined to use subjective probability as the sole criterion for our decision.

As it turns out, the problem as stated is poorly formulated. Relative utilities help both to clear up the question and to provide answers. Let $S_A$ stand for the regular salvation that Argle offers, and $S_B$ for Bargle's deluxe salvation. Our decision table is a $3 \times 3$ matrix similar to Table V:

| | Argle exists | Bargle exists | No god exists |
|---|---|---|---|
| Wager for *A* | $S_A$ | – | – |
| Wager for *B* | – | $S_B$ | – |
| Against all | – | – | – |

All outcomes other than $S_A$ and $S_B$ consist of mere earthly goods. It looks like $S_B$ is infinitely better than $S_A$, which in turn is infinitely better than any of the worldly outcomes. But things are not so simple. The choice of, and assumptions about, the base-point $Z$ turn out to be crucial.

The base-point $Z$ is the outcome with respect to which utility ratios are computed. It represents the default outcome in all gambles used for determining relative utilities. The only compulsory requirement in designating a base-point, given that all relative utilities are non-negative extended real numbers, is that it must be a lower bound for all possible outcomes in the decision at hand. Beyond this constraint, the appropriate choice of $Z$ is governed by pragmatic considerations. We have to determine the fallback position in our assessments of relative utility. As it turns out, though, the precise designation of $Z$ is usually unimportant. All that matters is the 'level' of the base-point. This is certainly the case in an Archimedean setting, where a shift in origin makes no essential difference. Analogously, in a non-Archimedean context, any two base-points that occupy the same level (and are lower bounds for all possible outcomes) will lead to the same ranking of outcomes by relative utility. So the choice of $Z$ is not particularly problematic.

Typically (but not always), $Z$ will be some finite worldly outcome that is a lower bound for all those under consideration. But in the examples below, we will also consider cases where the base-point is infinitely good or infinitely bad.

Suppose first that as in the ordinary Wager, $Z$ denotes some earthly reward dominated by all those appearing in the table. For any two worldly outcomes F and G in the table, $\mathcal{U}(\mathrm{F}, \mathrm{G}; Z)$ is finite. It is also plain that $\mathcal{U}(\mathrm{S_A}, \mathrm{F}; Z) = \mathcal{U}(\mathrm{S_B}, \mathrm{F}; Z) = \infty$, by the same reasoning as in the ordinary Wager. The key question is what value to assign to $\mathcal{U}(\mathrm{S_B}, \mathrm{S_A}; Z)$. It is tempting to let this be $\infty$ as well, but we should consider carefully what that would mean. By Definition 1,

$$\mathcal{U}(\mathrm{S_B}, \mathrm{S_A}; Z) = \infty \quad \text{if and only if} \quad S_A \preccurlyeq [pS_B, (1-p)Z].$$

For the relative utility to be infinite, you must prefer any gamble that offers a slight chance for deluxe salvation to regular salvation – even when that gamble entails a very high probability for ordinary finite existence! You must be willing to sacrifice the certainty of a comfortable, unending life for the virtual certainty of a finite existence, given only a tiny chance of the deluxe reward. Although it is certainly possible to have such preferences, they are not on par with those of the ordinary Pascalian who, after all, is making only a finite sacrifice for a shot at salvation.

In my view, $\mathcal{U}(S_B, S_A; Z) = 1$ is the most reasonable value. Some people may want to assign a value higher than 1, signalling a kind of high-level Pascalian preference structure. I won't present an argument for my position – *de gustibus non est demonstrandum* – but I do offer a formal postulate and an analogy.

POSTULATE 10 (PARALLAX POSTULATE). Let $A$ and $B$ be pure outcomes. Suppose that for some $C$ strictly preferred to $Z$ we have $\mathcal{U}(A, C; Z) = \infty$. Then $\mathcal{U}(B, A; Z) = 1$ whenever $A \preccurlyeq B$.

Roughly: if $A$ already looks infinitely good from base-point $Z$, then you are relatively indifferent between $A$ and any superior outcome $B$. You are unwilling to risk losing $A$ to upgrade to $B$ when the fall-back is $Z$.[26] From an earthly base-point, Bargle's paradise does look much better than Argle's, but no non-trivial gamble for the former justifies the sacrifice of the latter.

Of course, either your preferences conform to this principle or they do not. The postulate is not compulsory, but rather an interesting test of intuition. I regard it as a prudent piece of advice; others may see it as overly restrictive.

Now for the analogy. Stellar parallax refers to the angular displacement of a nearby star, relative to distant stars, as the earth orbits the sun. This displacement can be used to determine the nearby star's distance from earth. For all practical purposes, we can treat the background stars as infinitely remote. If we look at any two of these 'infinitely' distant stars, $A$ and $B$, we see no angular displacement. They remain fixed in the same position relative to each other as the earth traces its orbit, and yield the same angular displacement for any nearby star $C$. Analogously, two infinitely good prizes $A$ and $B$ have a fixed relation to each other, and to any worldly prize $C$, from our worldly perspective. It shouldn't matter that $B$ is actually 'further' (better) than $A$.

The Parallax Postulate gives us one way to resolve the super-Pascalian Wager. If we embrace the postulate, then $\mathcal{U}(S_A, S_B; Z) = 1$. Hence, if we form the decision matrix relative to the optimal outcome $S_B$ as required by Corollary 9, the result is identical to Table IV. It follows that the correct choice is to go with the god who has higher subjective probability. Indeed, if the Parallax Postulate is granted, the resolution of the many-gods argument in Section 5.1 becomes perfectly general. Pairwise comparison of salvation under different divinities always yields a relative utility of 1, provided that $Z$ denotes a mere earthly reward. Subjective probability is then the sole guide to selecting a divinity.

If we reject the Parallax Postulate, we get a different solution to the super-Pascalian Wager. The only real alternative is to set $\mathcal{U}(S_B, S_A; Z) = \infty$. Then the decision matrix relative to $S_B$ is the same as Table IV except that we replace the top-left 1 with a 0. Wagering for Bargle becomes the optimal choice, regardless of the probability values. If you have these preferences, the super-Pascalian Wager is essentially the same as the ordinary Wager. Argle's paradise is a trifling distraction more or less on par with earthly prizes, since it cannot compete with Bargle's paradise.

There is yet a third solution if the agent making the choice is an angelic being who already enjoys low-grade immortality. The base-point $Z$ represents this current state, from which both $S_A$ and $S_B$ are upgrades. The assumption $\mathcal{U}(S_B, S_A; Z) = \infty$ is reasonable here, because the default prize $Z$ is not so much worse than $S_A$. The solution here is the same as in the second case: we should wager for Bargle.

The foregoing discussion highlights the importance of designating the base-point $Z$. Different solutions can be justified, depending upon our specification of $Z$ and the crucial value $\mathcal{U}(S_B, S_A; Z)$. It is worth comparing this analysis to what we might obtain using a standard representation of non-Archimedean preferences. One approach is to modify the vector-valued utility function of Section 3 and employ a three-dimensional representation. Outcomes have utility $(x, y, z)$, where $x$ represents earthly goods, $y$ represents regular heavenly goods and $z$ represents deluxe heavenly goods. We have a lexicographic ordering, with priority given to the third and then the second component. By reasoning as we did in Section 3, it is easy to see that the optimal choice is wagering for Bargle, regardless of the probabilities (so long as both are positive). By contrast, I argued for a choice based on subjective probability. The important point, though, is that relative utilities bring options into view that one does not have with the lexicographic approach. With relative utilities, we can formulate criteria that might lead to quite distinct decisions. The representation of preferences by relative utilities with a variable base-point offers greater flexibility in analysing such decisions than does a one-place utility function, whether unidimensional or multi-dimensional.

### 6.3. *Harsh versions of the Wager*

In our formulation of the Wager in Section 1, we represented the outcome in the bottom-left corner, where you wager against God yet God exists, with a finite utility value. Let us call this the humane version of the argument. Suppose that a harsh version of the argument is contemplated, in which this value is set at $-\infty$ to signify damnation, some form of ceaseless punishment. Such a decision table might seem to make the argument for wagering in favour of God even more compelling, but it complicates things when we consider mixed strategies and the many-gods argument. Some deities may mete out eternal torment to those who wager for their rivals. How are we to compute sums involving both $+\infty$ and $-\infty$? Just as relative utilities clarify our reasoning for the humane version of the Wager, they are helpful in discussing the harsh version.

The basic idea is to let $Z$ signify damnation, the worst outcome. In fact, this choice is practically forced upon us by the constraint that the base-point be a lower bound for all outcomes under consideration.[27] Note that $\mathcal{U}(O, Z; Z) = \infty$ for any of the other outcomes $O$ that might occur in the Wager: all are infinitely better than damnation. If the Parallax Postulate is granted, then we have $\mathcal{U}(O_1, O_2; Z) = 1$ for any two such outcomes. Even for the sake of salvation, nobody will take the risk of landing in Hell. Relative to the best outcome (salvation), then, the decision matrix looks like as given in Table VI.

In this case, maximizing expected relative utility justifies the pure wager over all mixed strategies. Without the Parallax Postulate, we can assume that Salvation has infinite utility relative to both earthly rewards in the right column. The resulting relative decision matrix looks the same as Table V except that both values in the right

TABLE VI

Pascal's Wager (harsh version)

| | God exists | God does not exist |
|---|---|---|
| Wager for God | 1 | 1 |
| Wager against God | 0 | 1 |

column change to 0. Again, the pure wager comes out ahead of all mixed strategies. The Parallax Postulate makes no difference in this case.

We can also employ relative utilities to find a solution when we combine the harsh wager with many gods. Suppose that our theology encompasses two possible deities, both of whom reward believers with salvation and punish non-believers with damnation. Table VII is the naïve decision matrix.

Our reasoning reaches a dead end at once: a naïve calculation of expected utility involves the quantity $\infty - \infty$, which is undefined. Letting $Z$ stand for damnation and relativizing everything to the best outcome (salvation), however, we have the relative decision matrix given in Table VIII.

This table assumes the Parallax Postulate, but without that assumption we would simply replace the right column of 1's with three 0's. Either way, the solution is what we might expect: to wager in favour of the deity whose existence is assigned the highest subjective probability.

For another variation, let Bargle be relatively benign. Replace the two occurrences of $-\infty$ in column two of Table VII with finite values. Perhaps Bargle rewards believers, but is not especially hard on non-believers or believers in other divinities. The corresponding change to Table VIII is that column 2 now consists entirely of 1's.

TABLE VII

Harsh wager, many-gods I

| | $q_1$ Argle exists | $q_2$ Bargle exists | $1-(q_1+q_2)$ no god exists |
|---|---|---|---|
| Wager for A | $\infty$ | $-\infty$ | $f_1$ |
| Wager for B | $-\infty$ | $\infty$ | $f_2$ |
| Against all | $-\infty$ | $-\infty$ | $f_3$ |

TABLE VIII

Harsh wager, many-gods II

| | $q_1$ Argle exists | $q_2$ Bargle exists | $1-(q_1+q_2)$ no god exists |
|---|---|---|---|
| Wager for A | 1 | 0 | 1 |
| Wager for B | 0 | 1 | 1 |
| Against all | 0 | 0 | 1 |

Not surprisingly, to maximize expected relative utility we must now wager for Argle. The same thing happens if Bargle is really nice and rewards everyone with salvation – once again, Argle wins our allegiance. Nice gods finish last!

Finally, suppose that Argle perversely metes out infinite punishment to those who wager for Argle and rewards everyone else with salvation, while Bargle is a more conventional deity who rewards only Barglites. When we relativize everything to the best outcome (with damnation as the base-point), we get the following table if we assume the Parallax Postulate:

| | $q_1$ Argle exists | $q_2$ Bargle exists | $1-(q_1+q_2)$ no god exists |
|---|---|---|---|
| Wager for A | 0 | 1 | 1 |
| Wager for B | 1 | 1 | 1 |
| Against all | 1 | 1 | 1 |

This lets us rule out wagering for Argle, but no more. That appears to spell trouble for our whole approach. Even though we know that wagering for Bargle should come out ahead of wagering against all gods, both options have equal expected relative utility.

There is a natural way to supplement Corollary 9 that lets us handle this and similar examples. Employ a mixture of admissibility reasoning and maximization of expected relative utility. Here is how it works. If two or more rows are tied when we calculate expected relative utilities, strike out the rows that are dominated. The initial base-point is now irrelevant, so we move to a base-point that is essentially the minimum of the remaining outcomes under consideration, and re-calculate utilities relative to the optimal outcome. If this does not single out a unique choice, repeat the process.

In the example at hand, we strike out the first row and move to a base-point that represents a finite worldly outcome no better than any in the table. The new table of utilities relative to the optimal outcome (salvation) looks like this:

| | $q_1$ Argle exists | $q_2$b Bargle exists | $1-(q_1+q_2)$ no god exists |
|---|---|---|---|
| Wager for B | 1 | 1 | 0 |
| Against all | 1 | 0 | 0 |

We now have a justification for wagering in favour of Bargle.

## 7. CONCLUSION: BEYOND PASCAL'S WAGER

Relative utilities give us a conservative way to do infinite decision theory. The proposal is a simple generalization of finite utility theory that lets us retain the idea that rationality is characterized by maximizing expected (relative) utility while avoiding the technical and philosophical difficulties associated with attempts to model infinite value using lexicographic orderings or non-standard number systems. The approach works well when applied to problems like Pascal's Wager, where relative utilities are either supplied or readily computed. Its potential for treating other puzzles about infinite value lies in its flexibility for modelling different intuitions about preferences among gambles.

While this paper has focused entirely on the Wager, relative utilities may have a valuable application in the analysis of moral reasoning. Moral preferences, even in mundane settings where the question of infinite value does not arise, are non-Archimedean. Consider a model of moral decision-making that combines Kantian and utilitarian intuitions. Let us suppose that we regard some outcomes as permissible and others as intolerable. We use utilitarian principles in making comparisons among the permissible outcomes and among the intolerable ones (a gentle murder is better than a cruel one), but we cannot employ a common scale to make cross-boundary comparisons. That is, we can never be indifferent between a permissible outcome and a gamble that gives a positive probability for an intolerable result. The Archimedean condition is violated. Furthermore, there might be downward hierarchies of increasingly intolerable actions. We may regard both massive fraud and murder as intolerable, but view murder as infinitely worse than fraud.

To model this sort of relationship is a challenge for utilitarians. Following the strategies of Section 3, we could make use of a vector-valued utility function with lexicographic ordering, or non-standard utility values. In my view, it is simpler to use relative utilities. This approach has a clear advantage when it comes to updating our preferences. For example, suppose we have employed a two-dimensional lexicographic ordering to represent earthly and heavenly rewards, but now find (as in Section 6.2) that we need to make decisions involving different grades of salvation. We have to add a third dimension, and this means that all of our utility assignments must be changed. By contrast, adding information about previously unknown relative utilities may be perfectly compatible with existing relative utility assignments.

If relative utilities are used to model our preferences, then we have a generalized utilitarian framework that lets us model decision-making even in cases where no trade-offs are possible. We can do this by using relative decision matrices and expected relative utilities. Here is a simple illustration.

Alex is being robbed by a highwayman. Escape, injury, murder, brutal murder: all are possible outcomes. His options are to flee or to co-operate with the thief. He believes that the outcome depends upon his assessment of the thief's temperament (nasty or not) and physical condition (fit or flabby), as summarized in the following table.

| | NASTY | | NOT NASTY | |
|---|---|---|---|---|
| | Fit | Flabby | Fit | Flabby |
| Flee | Brutal murder | Escape | Injury | Escape |
| Co-operate | Gentle murder | Gentle murder | No injury | No injury |

If he flees, he escapes unharmed provided the thief is slow, but risks a beating or brutal murder (plus theft) if caught. If he co-operates, the theft is assured and the thief will either kill him quickly or leave him unharmed.

How should Alex make his decision? Dominance reasoning gives no clear verdict, so he turns to naïve expected utility reasoning. Since murder is intolerable to Alex, he assigns $-\infty$ to either type of murder, a large negative value to theft with injury, a moderate negative value to theft without injury, and zero or a positive value to escape. But Alex soon discovers that both actions have an expected value of $-\infty$. Just as in naïve versions of Pascal's Wager, expected utility calculations are of no use. Nevertheless, it seems clear that the prudent thing to do is to flee, because flight offers some chance of escape in the more important case of a nasty assailant.

We can capture this reasoning using relative utilities. If $Z$ represents the violent murder (or something even worse) and $O$ represents the optimal outcome of escape, then the relative utility values $\mathcal{U}(-, O; Z)$ give us the following relative decision matrix:

| | NASTY | | NOT NASTY | |
|---|---|---|---|---|
| | Fit | Flabby | Fit | Flabby |
| Flee | 0 | 1 | 1 | 1 |
| Co-operate | 0 | 0 | 1 | 1 |

Here I have assumed the Parallax Postulate. Any outcome where Alex survives is more or less equally good relative to the horrific base-point.[28] From this table, it is clear that Alex should flee, regardless of the probabilities. Flight is transformed into a dominant choice. The relative decision matrix collapses distinctions between the two murder outcomes and distinctions among the three non-murder outcomes, but that is in accordance with intuition.

Many of the assumptions in this analysis are debatable. Alex might think that if he co-operates, there is some chance that a nasty criminal will let him live. Or perhaps flight might provoke even a mild-mannered thief to murder. Such modifications can be accommodated along lines similar to the treatment of the cases in Section 6.2. Decisions can still be made by maximizing expected relative utility.

Can we employ relative utilities to analyze even more complex examples involving infinitesimal probabilities, or decisions that must take into account infinitely many options or states? Despite some obvious difficulties, I'd wager that it's worth a try.

## APPENDIX A: PROOFS

### A.1. *Redundancy of (A3), the Better-chances Condition*

Assume $A \preccurlyeq B$. If $q \leq p$, then we have

$$A \sim [{}^{q}/_{p}A, (1-q/p)A] \preccurlyeq [{}^{q}/_{p}A, (1-q/p)B] \text{ by (A2).}$$

It follows that

$$\begin{aligned}[pA, (1-p)B] &\preccurlyeq [p[{}^{q}/_{p}A, (1-q/p)B], (1-p)B] \text{ by (A2)} \\ &\sim [qA, (1-q)B] \text{ by (A4).}\end{aligned}$$

Strictness follows from (A2) if both $A \prec B$ and $q < p$. If $q > p$, the reverse inequality holds by symmetry.

A.2. *Proof of Lemma 2*

The set of all $k$ such that $[kR, (1-k)P] \preccurlyeq Q$ is non-empty and bounded above by 1. Let $\alpha$ be the supremum of this set. If $\alpha = 0$ then case (c) obtains; if $\alpha = 1$, then we have case (b). If $0 < \alpha < 1$, then it is easy to show that one of cases (a)–(c) obtains.

A.3. *Proof of Lemma 4*

Lemma 4 states:

If $B \approx_Z B'$, then for any $0 \leq p \leq 1$ and any $A$, $C$,

(1) $[pB, (1-p)C] \approx_Z [pB', (1-p)C]$ and
(2) $[pA, (1-p)B] \approx_Z [pA, (1-p)B']$.

We show only (1), since the proof of (2) is similar. Further, if $B \sim B'$ then (1) follows at once from the Better-Prizes Condition. So we may assume $B \preccurlyeq B'$ where the preference is strict.

The assumption $B \approx_Z B'$ means that

(a) $[qB', (1-q)Z] \preccurlyeq B$, all $0 \leq q < 1$
(b) $[qB, (1-q)Z] \preccurlyeq B'$, all $0 \leq q < 1$.

The proof divides into two cases: $B' \preccurlyeq C$ and $C \preccurlyeq B'$, but we give the argument only for the case $B' \preccurlyeq C$ (since the arguments are similar for the second case). Assume, then, that $B' \preccurlyeq C$.

By Lemma 3, one of the following holds:

(i) $C \approx_Z B'$;
(ii) $B' \approx_Z [aC, (1-a)Z]$ for a unique $a$ with $0 < a < 1$;
(iii) $B' \preccurlyeq [aC, (1-a)Z]$ for all $0 < a < 1$.

For (i): it suffices to show that $[pB, (1-p)C] \approx_Z C$, since then, by parity of reasoning and transitivity of $\approx_Z$, we have $[pB, (1-p)C] \approx_Z C \approx_Z [pB', (1-p)C]$. So we need to show that

$$[kC, (1-k)Z] \preccurlyeq [pB, (1-p)C] \text{ if } 0 \leq k < 1.$$

But this follows because $[kC, (1-k)Z] \preccurlyeq B$ (since $B \approx_Z C$) and $[kC, (1-k)Z] \preccurlyeq C$ (by (A3)). (Note: the fact that $B' \preccurlyeq C$ is not used in this argument.)

For (ii): To prove (1), it suffices to prove that

$$(1^*) \quad [\alpha[pB', (1-p)C], (1-\alpha)Z] \preccurlyeq [pB, (1-p)C] \quad \text{for all } 0 \leq \alpha < 1,$$

since given that $B \preccurlyeq B'$, the other half of (1) is an immediate consequence of the Better-Prizes Condition (A2) and the Better-Chances Condition (A3).

Since $\approx_Z$ is transitive, we have $B \approx_Z [aC, (1-a)Z]$ as well. Hence for any $0 < \beta < 1$ we have

$$[\beta[aC, (1-a)Z], (1-\beta)Z] \preccurlyeq B$$

and so by (A2),

$$(3) \qquad \begin{aligned} &[pB, (1-p)C] \\ &\succcurlyeq [p[\beta[aC, (1-a)Z], (1-\beta)Z], (1-p)C] \\ &\sim [p[\beta aC, (1-\beta a)Z], (1-p)C] \text{ by (A4)} \\ &\sim [tC, (1-t)Z] \text{ by (A4) ,} \end{aligned}$$

where $t = p\beta a + (1-p)$.

Also, since $B' \approx_Z [aC, (1-a)Z]$, we have

$$B' \leqslant [bC, (1-b)Z] \quad \text{if} \, a < b \leq 1,$$

and so for all such $b$,

$$(4) \qquad \begin{aligned} &[\alpha(pb + (1-p))C, (1-\alpha(pb + (1-p)))Z] \\ &\sim [\alpha[(pb+(1-p))C, (1-(pb+(1-p)))Z], (1-\alpha)Z] \\ &\quad \text{by (A4)} \\ &\sim [\alpha[p[bC, (1-b)Z], (1-p)C], (1-\alpha)Z] \\ &\quad \text{by (A4)} \\ &\succcurlyeq [\alpha[pB', (1-p)C], (1-\alpha)Z] \\ &\quad \text{by (A2).} \end{aligned}$$

But we can choose $b$ and $\beta$ so that $a < b \leq 1$ and $0 < \beta < 1$ and

$$p\beta a + (1-p) > \alpha[pb + (1-p)].$$

(1*) then follows from (3), (4) and (A3).

For (iii): it suffices to show that $[pZ, (1-p)C] \approx_Z [pB, (1-p)C]$, since by parity of reasoning and transitivity of $\approx_Z$ we will have $[pB, (1-p)C] \approx_Z [pZ, (1-p)C] \approx_Z [pB', (1-p)C]$. So we need only show that

$$[k[pB, (1-p)C], (1-k)Z] \preccurlyeq [pZ, (1-p)C] \quad \text{for } 0 \leq k < 1.$$

But $B \preccurlyeq [\alpha Z, (1-\alpha)C]$ for all $0 \leq \alpha < 1$, so that if $t = \alpha(1-k(1-p))$

$$\begin{aligned}
&[tZ, (1-t)C] \\
&\quad \sim [(1-k(1-p))[\alpha Z, (1-\alpha)C], k(1-p)C]] && \text{by (A4)} \\
&\quad \succcurlyeq [(1-k(1-p))B, k(1-p)C] && \text{by (A2)} \\
&\quad \sim [k[pB, (1-p)C], (1-k)B] && \text{by (A4)} \\
&\quad \succcurlyeq [k[pB, (1-p)C], (1-k)Z] && \text{by (A2).}
\end{aligned}$$

The required result now follows via (A3) provided that we can choose $\alpha$ so that $t > p$, i.e., $1 > \alpha > p/[1-k(1-p)]$. Since $p < 1$, we have $p(1-k) < (1-k)$, so that $p < 1-k(1-p)$, and we can find such an $\alpha$.

## A.4. *Proof of Theorem 6*

*Step 1*: Definition of $\mathcal{U}$. This is Definition 5.

*Step 2*: *Proof of simple properties of* $\mathcal{U}$.

(R1), (R2), (R4) and (R5) follow immediately from the definition and the fact that $\approx_Z$ is an equivalence relation. (For (R4), we need to observe that if $A \approx_Z Z$, then $A \sim Z$. For if $A \approx_Z Z$ and $Z \prec A$ (strict preference), then for $0 < p < 1$, $[pA, (1-p)Z]$ would be strictly preferred to $Z$ by (A2), contradicting $[pA, (1-p)Z] \preccurlyeq Z$.)

Property (R3) is proven below, at Step 3.

Proof of (R6):

If $\mathcal{U}(A, C; Z) = \infty$, the inequality is obviously true.

If $\mathcal{U}(A, C; Z) = 0$, then $A \preccurlyeq [pC, (1-p)Z]$ for all $0 < p < 1$ and so $B \preccurlyeq [pC, (1-p)Z]$ for all $0 < p < 1$, implying that $\mathcal{U}(B, C; Z) = 0$.

If $\mathcal{U}(A, C; Z) = k \leq 1$, then $A \approx_Z [kC, (1-k)Z]$. Applying Lemmas 3 and 4 together with (A4), either $B \approx_Z [\alpha kC, (1-\alpha k)Z]$ for some $0 < \alpha \leq 1$ or $B \preccurlyeq [\alpha kC, (1-\alpha k)Z]$ for all such $\alpha$. In either case, $\mathcal{U}(B, C; Z) \leq \mathcal{U}(A, C; Z)$.

Finally, if $\mathcal{U}(A, C; Z) = k > 1$, then $[(1/k)B, (1-(1/k))Z] \preccurlyeq [(1/k)A, (1-(1/k))Z] \approx_Z C$ by (A2) and it follows that $\mathcal{U}(B, C; Z) \leq k$.

*Proof of* (R7):

Let $\mathcal{U}(C, B; Z) = d$, $\mathcal{U}(C, A; Z) = e$ and $\mathcal{U}(A, B; Z) = f$. We know $0 < e \leq 1 \leq f$.

If $f=\infty$, then for any $p>0$ we have

$$\begin{aligned} B &\preccurlyeq [peA,(1-pe)Z] \\ &\approx_Z [p[eA,(1-e)Z],(1-p)Z] \\ &\approx_Z [pC,(1-p)Z] \text{ by the Substitution Lemma,} \end{aligned}$$

so that $R(C,B;Z)=\infty$ and $d=ef$ as required.

If $f$ is finite, then from (R5) and (R6) it follows that $e \leq d \leq f$, so $d$ is also positive and finite. If $d \leq 1$, then $C \approx_Z [dB,(1-d)Z]$ and $B \approx_Z [(1/f)A,(1-1/f)Z]$, so by the Substitution Lemma,

$$\begin{aligned} C &\approx_Z [d[(1/f)A,(1-1/f)Z],(1-d)Z] \\ &\approx_Z [(d/f)A,(1-d/f)Z] \text{ by (A4)} \end{aligned}$$

which proves that $e=d/f$ as required.

If $d>1$, then $B \approx_Z [(1/d)C,(1-1/d)Z]$ and $C \approx_Z [eA,(1-e)Z]$, so by the Substitution Lemma,

$$\begin{aligned} B &\approx_Z [(1/d)[eA,(1-e)Z],(1-1/d)Z] \\ &\approx_Z [(e/d)A,(1-e/d)Z] \text{ by (A4)} \end{aligned}$$

and we have $f=d/e$, as required.

*Step 3*: Proof of (R3)
(R3) states:

$$\mathcal{U}([pA,(1-p)A'],B;Z)=p\mathcal{U}(A,B;Z)+(1-p)\mathcal{U}(A',B;Z) \quad \text{for } 0 \leq p \leq 1.$$

Write L for $[pA,(1-p)A']$, and assume $0<p<1$ for non-triviality.

*Case 1*: At least one of $\mathcal{U}(A,B;Z)$ or $\mathcal{U}(A',B;Z)$ is $\infty$. Suppose $\mathcal{U}(A,B;Z)=\infty$; the argument is similar if $\mathcal{U}(A',B;Z)=\infty$. In this case, the right side of (R3) is $\infty$.
Let $L'=[pA,(1-p)Z]$. First note that $\mathcal{U}(L',B;Z)=\infty$ for $p>0$: for $B \preccurlyeq L'$ because $\mathcal{U}(A,B;Z)=\infty$, and if $[kL',(1-k)Z] \approx_Z B$ for $0<k<1$, then by (A4) we would have $[kpA,(1-kp)Z] \approx_Z B$ which contradicts $\mathcal{U}(A,B;Z)=\infty$.
It now follows that $\mathcal{U}(L,B;Z)=\infty$, from (R6) and $L' \preccurlyeq L$.

*Case 2*: Both $\mathcal{U}(A,B;Z)$ and $\mathcal{U}(A',B;Z)$ are finite.

We first prove the following result, a special case of (R3).

LEMMA 11. (PARTIAL LINEARITY). For all $A$,$B$ and $0 \leq p \leq 1$, $\mathcal{U}([pA, (1-p)Z], B; Z) = p\mathcal{U}(A, B; Z)$.

*Proof*: Let $\mathrm{L}' = [pA, (1-p)Z]$. We may assume $0 < p < 1$.

If $\mathcal{U}(A, B; Z) = \infty$, the result is entailed by *Case 1*.

If $\mathcal{U}(A, B; Z) = 0$, then from $\mathrm{L}' \preccurlyeq A$ and (R6), we have $\mathcal{U}(\mathrm{L}', B; Z) \leq \mathcal{U}(A, B; Z) = 0$ and both sides of the equality are 0.

If $\mathcal{U}(A, B; Z) = k \leq 1$, then $A \approx_Z [kB, (1-k)Z]$, which implies $\mathrm{L}' \approx_Z [pkB, (1 - pk)Z]$. Then $\mathcal{U}(\mathrm{L}', B; Z) = pk = p\mathcal{U}(A, B; Z)$, as required.

Finally, suppose $\mathcal{U}(A, B; Z) = k > 1$, $k$ finite. We have $B \preccurlyeq A$ and $\mathrm{L}' \preccurlyeq A$. By (R7),

$$\mathcal{U}(L', B; Z) = \mathcal{U}(L', A; Z)\mathcal{U}(A, B; Z) = p\mathcal{U}(A, B; Z).$$

We complete the argument for Case 2 by considering three sub-cases.

*Case 2a*: Both $\mathcal{U}(A, B; Z)$ and $\mathcal{U}(A', B; Z)$ are zero.
The right side of (R3) is 0. The left side is also 0, because either $A' \preccurlyeq A$ or $A \preccurlyeq A'$. By (A2), either $L \preccurlyeq A$ or $L \preccurlyeq A'$. It follows that $\mathcal{U}(L, B; Z) = 0$.

*Case 2b*: $\mathcal{U}(A, B; Z) = c > 0$ and $\mathcal{U}(A', B; Z) = 0$ (i.e., just one of the two terms is zero). The right side of (R3) is $pc$.
By (A2) and Lemma 11, $\mathcal{U}(L, B; Z) \geq \mathcal{U}([pA, (1-p)Z], B; Z) = pc$.
From (R7), $\mathcal{U}(A', A; Z) = 0$.
By (A2) and Lemma 11, for every $0 < k$ we have

$$\begin{aligned} \mathcal{U}(L, B; Z) &\leq \mathcal{U}([pA, (1-p)[kA, (1-k)Z]], B; Z) \\ &= \mathcal{U}([tA, (1-t)Z], B; Z) \quad \text{where } t = p + (1-p)k \\ &= (p + (1-p)k)c. \end{aligned}$$

As required, it follows that $\mathcal{U}(L, B; Z) = pc$ (since $k$ can be as small as we please).

*Case 2c*: Both $c = \mathcal{U}(A, B; Z)$ and $c' = \mathcal{U}(A', B; Z)$ are finite and non-zero.

If $A \preccurlyeq B$ and $A' \preccurlyeq B$, then $A \approx_Z [cB, (1-c)Z]$ and $A' \approx_Z [c'B, (1-c')Z]$. It follows by substitution that

$$\begin{aligned} L &\approx_Z [p[cB, (1-c)Z], (1-p)[c'B, (1-c')Z]] \\ &\approx_Z [tB, (1-t)Z] \text{ where } t = pc + (1-p)c', \\ &\qquad \text{by (A4)} \end{aligned}$$

and this gives the required result

$$\begin{aligned} \mathcal{U}(L, B; Z) &= pc + (1-p)c' \\ &= p\mathcal{U}(A, B; Z) + (1-p)\mathcal{U}(A', B; Z). \end{aligned}$$

The other possibility is that either $A$ or $A'$ is the dominant outcome; we may suppose $A' \preccurlyeq A$ and $B \preccurlyeq A$, and the argument will run similarly if $A'$ is dominant. Let $k = \mathcal{U}(A', A; Z)$; we know from (R7) that $c' = kc$. By substitution,

$$\begin{aligned} L &\approx_Z [pA, (1-p)[kA, (1-k)Z]] \\ &\approx_Z [tA, (1-t)Z] \quad \text{where } t = p + (1-p)k, \\ &\qquad \text{by (A4)} \end{aligned}$$

and by Lemma 11, this gives the required result

$$\begin{aligned} \mathcal{U}(L, B; Z) &= (p + (1-p)k)\mathcal{U}(A, B; Z) \\ &= p\mathcal{U}(A, B; Z) + (1-p)\mathcal{U}(A', B; Z). \end{aligned}$$

### A.5. *Uniqueness of U in Theorem 6*

We first prove a preliminary result.

LEMMA 12. Suppose the ordering $\preccurlyeq$ satisfies (A1)–(A4), and $\mathcal{U}$ is as defined at Step 1. Then

(i) $A =_Z B \leftrightarrow 0 < \mathcal{U}(A, B; Z) < \infty$ defines an equivalence relation on the set $\{A/Z \preccurlyeq A\}$. Write $[A]_Z$ for the equivalence class of $A$.
(ii) If $[A]_Z \neq [B]_Z$ and $B \preccurlyeq A$, then $B' \preccurlyeq A'$ for any $A' =_Z A$ and $B' =_Z B$. In this case, we may write $[B]_Z \preccurlyeq [A]_Z$.

*Proof*: (i) is evident from (R2), (R5) and (R7).

For (ii), if $[A]_Z \neq [B]_Z$ and $B \preccurlyeq A$, it must be that $\mathcal{U}(B, A; Z) = 0$. So

$$B \preccurlyeq [pA, (1-p)Z] \text{ for every } 0 < p \leq 1.$$

By considering cases, it is easily shown that $\mathcal{U}(B', A'; Z) = 0$.

Once we fix the base-point $Z$, the equivalence classes of Lemma 12 are steps along which outcomes are comparable, i.e., Archimedean. Successive steps are infinitely preferable to their predecessors.

Now we can complete the uniqueness proof. Suppose that $\mathcal{U}'$ has the properties (R1)–(R7). We must show that for any $A$ and $B$, $\mathcal{U}'(A, B; Z) = \mathcal{U}(A, B; Z)$.

- If $[B]_Z \preccurlyeq [A]_Z$, then $\mathcal{U}(B, A; Z) = 0$ and $B \preccurlyeq [pA, (1-p)Z]$ for $0 < p \leq 1$. It follows that for all $0 < p \leq 1$,

$$\begin{aligned} \mathcal{U}'(B, A; Z) &\leq \mathcal{U}'([pA, (1-p)Z], A; Z) && (R6) \\ &= p\mathcal{U}'(A, A; Z) + (1-p)\mathcal{U}'(Z, A; Z) && \text{(R3)} \\ &= p && \text{(R2), (R4)} \end{aligned}$$

and hence $\mathcal{U}'(B, A; Z) = 0 = \mathcal{U}(B, A; Z)$. The argument is similar if $[A]_Z \preccurlyeq [B]_Z$.

- If $[A]_Z = [B]_Z$, suppose first that $B \preccurlyeq A$. Then if $c = \mathcal{U}(B, A; Z)$ we have

$$B \approx_Z [cA, (1-c)Z]$$

  and so we must have

$$\begin{aligned} \mathcal{U}'(B, A; Z) &= \mathcal{U}'([cA, (1-c)Z], A; Z) \\ &= c\mathcal{U}'(A, A; Z) + (1-c)\mathcal{U}'(Z, A; Z) && \text{(R3)} \\ &= c && \text{(R2), (R4)} \end{aligned}$$

  and again we have $\mathcal{U}'(B, A; Z) = \mathcal{U}(B, A; Z)$. The argument is similar if $A \preccurlyeq B$.

## A.6. *Proof of Lemma 7*

First, $\mathcal{U}(O_{ij}, A; Z) = \mathcal{U}(O_{ij}, B; Z)$ for all $i, j$ by property (R5). So we have

$$\begin{aligned} \mathcal{U}(A, B; Z) &= \Sigma q_j \Sigma p_i \mathcal{U}(O_{ij}, B; Z) && \text{since } O_{ij} = A_i \& S_j \\ &= \Sigma q_j \Sigma p_i \mathcal{U}(O_{ij}, A; Z) && \text{as noted above} \\ &= \mathcal{U}(A, A; Z) \\ &= 1 && \text{by (R2)} \end{aligned}$$

## A.7. *Proof of Corollary 8*

*Proof* of (1). The 'only if' claim follows from (R6) and (R7). To prove the 'if' claim, observe first that if we have $\mathcal{U}(A, O_{ij}; Z) > \mathcal{U}(B, O_{ij}; Z)$ for even a single $O_{ij}$, it must be (by (R6)) that $B \preccurlyeq A$. But if $\mathcal{U}(A, O_{ij}; Z) = \mathcal{U}(B, O_{ij}; Z)$ for all $O_{ij}$, then $A \approx_Z B$ by Lemma 7. Hence $A$ is $Z$-optimal.

*Proof* of (2). The 'only if' claim follows from part (1). To prove the 'if' claim, suppose $\mathcal{U}(A, O; Z) \geq \mathcal{U}(B, O; Z)$ for all available actions $B$. Fix $B$. We know that if $\mathcal{U}(A, O; Z) > \mathcal{U}(B, O; Z)$ we must have $B \preccurlyeq A$ by (R6), so we may suppose $\mathcal{U}(A, O; Z) = \mathcal{U}(B, O; Z)$. We show that $\mathcal{U}(A, O_{ij}; Z) = \mathcal{U}(B, O_{ij}; Z)$ for all i, j, so that $A \approx_Z B$. This result is derived as follows:

$$\begin{aligned} \mathcal{U}(A, O_{ij}; Z) &= \mathcal{U}(A, O; Z) \cdot \mathcal{U}(O, O_{ij}; Z) && \text{by (R7)} \\ &= \mathcal{U}(B, O; Z) \cdot \mathcal{U}(O, O_{ij}; Z) && \\ &= \mathcal{U}(B, O_{ij}; Z) && \text{by (R7).} \end{aligned}$$

For the application of (R7), however, we need $\mathcal{U}(A, O; Z) > 0$. Recall that $O = O_{kl}$ for some $k$ and $l$, and we know that $\mathcal{U}(A_k, O_{kl}; Z) \geq q_1 > 0$, where $q_1$ is the probability of state $S_1$. Since by assumption $\mathcal{U}(A, O; Z)$ is maximal, we have $\mathcal{U}(A, O; Z) \geq \mathcal{U}(A_k, O; Z) > 0$.

## NOTES

[1] In addition to McClennen, others who discuss infinite value include Sorensen (1994), Sobel (1996), Vallentyne (1993) and Vallentyne and Kagan (1997).
[2] We might doubt whether deliberate action can alter one's degree of belief in God, but in this paper I take it for granted that it can.
[3] I shall assume that all probabilities are real-valued. Infinitesimals are not allowed. This restriction is partly for the sake of simplicity and partly for fidelity to Pascal's original argument.
[4] As Hacking (1972) points out, Pascal provides several distinct arguments, though most discussion has focused on versions similar to the one just presented.
[5] Resnik adopts this as an independent axiom. In fact, it may be derived from (A1), (A2) and (A4) (see Appendix A.1 at the end of this paper). I include it because it is useful in discussion.
[6] Weaker versions of this axiom are common (see Fishburn 1971), but I shall not discuss them since we are going to dispense with all of them.
[7] Schlesinger (1994) puts the point this way: "In cases where the mathematical expectations are infinite, the criterion for choosing the outcome to bet on is its probability" (90).
[8] For a brief informal presentation, see Royden (1968).

[9] In support of attributing something like (Ref +) to Pascal, we have his remark in *Pensées* 233: "unity joined to infinity adds nothing to it". In support of attributing (Ref ·), we have the Wager itself.
[10] Here we continue to assume that the argument is as represented in Table I, where all other utilities reflect mere worldly prizes. In particular, no outcome is assigned negative infinite utility.
[11] One thing Hájek does not do, *pace* McClennen, is show that one can justify the *Expected Utility Theorem* in the setting of Conway's construction. Preferences are clearly non-Archimedean (i.e., (A5), the Continuity condition, fails), so a justification is required. Continuity can be restored if we allow not just the utilities but also the probabilities to take on surreal values, and in this case it may be that an analogue of the Expected Utility Theorem can be proven.
[12] Multiplication of a utility vector by a real number is just scalar multiplication.
[13] Once again, though, an independent justification should be provided for the requirement of maximizing expected utility.
[14] The same problem affects infinite hyperreals, though not $\aleph_0$; however, the Cantorian formulation is vulnerable to the broader objection to arbitrary cut-off points discussed below.
[15] See Fishburn (1974) for a survey and also Skala (1975).
[16] Jeffrey (1983, ch. 5) provides a good account of this approach, ultimately due to Ramsey.
[17] Putting things in this way also makes it starkly clear how hard it would be to adopt such an utterly detached attitude to worldly values – and how fanatical one's attachment to salvation might appear to be. On this analysis, it is clear that Pascal's argument would have no grip whatsoever on someone who does not have such preferences.
[18] The definition of preference intervals and Lemma 2 are both found in Fishburn (1971).
[19] It is possible to take one of the outcomes in the decision table – $O_{21}$ for example – as the base point. That would change some of the relative utilities, but it would make no difference to the analysis of the Wager developed below (Section 5).
[20] We assume independence, but there is no essential difficulty if we drop this assumption and move to conditional probabilities.
[21] Note that there is no difficulty if the optimal outcome is not unique. By (R7), we get the same relative decision matrix with respect to any optimal outcome.
[22] Alternatively, we may take $Z$ to be the worst of the four outcomes in the decision table, consistent with the practice we shall adopt later.
[23] *Pensées* 233.
[24] We might still be worried about the choice of the base-point $Z$. This concern will be addressed shortly.
[25] One could assert that God's happiness has to be represented as infinitely great relative to human salvation. This is the problem of interpersonal utility comparisons writ large! It is sufficient that salvation be absolutely maximal for humans. Pascal's original argument demands no more than this.
[26] The restriction of the Parallax Postulate to pure outcomes $A$ and $B$ is necessary. If $W_p$ is the gamble $[pS_A, (1-p)Z]$, then we have $\mathcal{U}(W_p, C; Z) = \infty$ for

any earthly reward $C$, but $\mathcal{U}(W_p, S_A; Z) = p$. For non-trivial gambles, we can use linearity to figure out relative utilities.

[27] A Dante-like series of infinitely worse hells would complicate matters, but might be handled in a way analogous to the super-Pascalian Wager of the preceding section. This scenario raises problems for the Parallax Postulate, though, since it dictates that from the lowest levels of the inferno, all higher levels (including paradise!) have approximately equal attraction. The approach developed in this paper seems to work best where the successive infinite jumps in utility move in the positive direction.

[28] It is important here that the two types of murder be seen as different only in degree.

## ACKNOWLEDGEMENTS

My thanks to Alan Hájek, Chris Stephens and two anonymous referees for extremely helpful comments and criticisms.

## REFERENCES

Conway, J.: 1976, *On Numbers and Games*, Academic Press, New York.

Duff, A.: 1986, 'Pascal's Wager and Infinite Utilities', *Analysis* **46**, 107–119.

Fishburn, P.: 1971, 'A Study of Lexicographic Expected Utility', *Management Science* **17**(11), 672–678.

Fishburn, P.: 1974, 'Lexicographic Orders, Utilities and Decision Rules: A Survey', *Management Science* **20**(11), 1442–1471.

Hacking, I.: 1972, 'The Logic of Pascal's Wager', *American Philosophical Quarterly* **9**(2), 186–192.

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

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

Jeffrey, R.: 1983, *The Logic of Decision, 2*nd ed., University of Chicago Press, Chicago.

Jordan, J. (ed.): 1994, *Gambling on God: Essays on Pascal's Wager*, Rowman & Littlefield, Savage MD.

Jordan, J.: 1998, 'Pascal's Wager Revisited', *Religious Studies* **34**, 419–431.

Luce, R. D. and H. Raiffa.: 1957, *Games and Decisions*, John Wiley & Sons, New York.

McClennen, E.: 1994, Finite Decision Theory, in Jordan 1994.

Mougin, G. and E. Sober.: 1994, 'Betting Against Pascal's Wager', *Nous* **28**, 382–395.

Pascal, B.: 1948, *Pensees*, trans. W. F. Trotter, J. M. Dent & Sons, London.

Resnik, M. D.: 1987, *Choices*, University of Minnesota Press, Minneapolis.

Royden, H. L.: 1968, *Real Analysis, 2*nd ed., Macmillan Publishing Co., London.

Schlesinger, G.: 1994, A Central Theistic Argument, in Jordan 1994.

Skala, H.: 1975, *Non-Archimedean Utility Theory*, D. Reidel, Dordrecht.
Sobel, H.: 1996, 'Pascalian Wagers', *Synthese* **108**, 11–61.
Sorensen, R.: 1994, Infinite Decision Theory, in Jordan 1994.
Swinburne, R.: 1969, 'The Christian Wager', *Religious Studies* **4**, 217–228.
Vallentyne, P.: 1993, 'Utilitarianism and Infinite Utility', *Australasian Journal of Philosophy* **71**, 212–217.
Vallentyne, P. and S. Kagan.: 1997, 'Infinite Value and Finitely Additive Value Theory', *Journal of Philosophy* **94**(1), 5–27.

Department of Philosophy
University of British Columbia
1866 Main Mall
E-370 Vancouver
Canada V6T 1Z1
Email: bartha@interchange.ubc.ca