Waging War on Pascal's Wager

Alan Hájek

1. Introduction

Pascal's Wager is simply too good to be true—or better, too good to be sound. There must be *something* wrong with Pascal's argument that decision-theoretic reasoning shows that one must (resolve to) believe in God, if one is rational. No surprise, then, that critics of the argument are easily found, or that they have attacked it on many fronts. For Pascal has given them no dearth of targets.

Virtually all of the Wager's critics have directed their campaigns against its premises. Other authors have rallied to its defense, buttressing those premises. I will argue that they are fighting a lost cause: developing arguments by Jeffrey (1983) and Duff (1986), I will contend that the Wager is simply invalid. This motivates a search for reformulations of the original argument that are valid, while upholding its spirit. I will offer four such reformulations, each of which finesses the decision matrix of the Wager, and in particular its problematic invocation of "infinite utility." Yet these reformulations fall too, albeit for a different reason. This, in turn, might prompt advocates of the Wager to conduct another search for still further reformulations. However, I will argue that such a search is likely to be futile. When we examine what is at the root of the failure of the original Wager, and of the reformulations that I offer, we realize that their failures are symptomatic of a deep problem that *any* variant of the Wager must overcome. I will present a dilemma for all such variants, and conclude that their prospects for success are dim.

2. The Wager, and Some Objections to Its Premises

We will think of Pascal's Wager as having three premises: the first concerns the probability that you should give to God's existence, the second offers a decision matrix, and the third is a standard decision-theoretic assumption about rational action.[1]

Pascal's Wager

1. Rationality requires you to give positive probability to God's existence.

2. The decision matrix is as follows:

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

Here f_1, f_2, and f_3 are finite utility values that need not be specified any further.

3. Rationality requires you to perform the act of maximum expected utility (when there is one).

Conclusion: Rationality requires you to wager for God.[2]

Some clarification is in order. According to Pascal, "wagering for God" and "wagering against God" are contradictories, as there is no avoiding wagering one way or another: "you must wager. It is not optional." (Unless otherwise stated, all quotations of Pascal are from his 1948, Trotter translation, §233.) The decision to wager for or against God is one that you make at a time—at t, say. But of course Pascal does not think that you would be infinitely rewarded for wagering for God momentarily, then wagering against God thereafter; nor that you would be infinitely rewarded for wagering for God sporadically—only on every other Thursday afternoon, for example. What Pascal intends by "wagering for God" is an *ongoing* action—indeed, one that continues until your death—that involves your adopting a certain set of practices and living the kind of life that fosters belief in God. The decision problem for you at t, then, is whether you should embark on this course of action; to fail to do so is to wager against God at t.[3]

I understand Pascal as regarding your salvation, with its infinite utility ("an infinity of infinitely happy life"), as *the best thing possible for you.* I take this to be in keeping with Catholic tradition,[4] but more importantly it is supported by Pascal's own text. In his preamble to the Wager, he writes: "Unity joined to infinity adds nothing to it … the addition of a unit can make no change in its nature." In particular, infinite utility is not augmented by the addition of a unit of utility. The point is made even more emphatically in the Wager itself, and here it will be useful to consult the original French text: "si vous gagnez vous gagnez tout." Trotter translates this accurately as: "If you gain, you gain all." We might just as accurately translate it as: "If you gain, you gain every-

thing." So according to Pascal, nothing could be better for you than your salvation. There is simply nothing more to be had.[5]

Some critics question Pascal's assumption that a rational agent should assign positive probability to God's existence. After all, a thoroughgoing atheist may insist on the rationality of an assignment of 0, as Rescher (1985) points out. Others attack the decision matrix. Various critics argue that Pascal conflates outcomes whose utilities should be distinguished. According to some, the states are not individuated finely enough. Perhaps there is more than one God to consider, as Diderot (1875–77) pointed out long ago, inaugurating the flourishing industry that has come to be known as "the many Gods objection." According to others, the acts are not individuated finely enough. Perhaps there is more than one way to wager for God—for instance, God might not reward those who strive to believe in him only for the very mercenary reasons that the Wager gives, as James (1956) has observed. Maybe the matrix is different for different people—as it might be, a predestined infinite reward for the Chosen, whatever they do, and finite utility for the rest, a possibility raised by Mackie (1982). And even granting Pascal his assumption of a single 2 x 2 matrix for all people, one could dispute the utilities that enter into it. Jeffrey (1983) and McClennen (1994) find the very notion of infinite utility suspect. Then there are the critics who, far from objecting to infinite utilities, want to see *more* of them in the matrix. For example, it might be thought that a forgiving God would bestow infinite utility upon wagerers-for and wagerers-against alike (Rescher 1985); or, more pessimistically, that wagering against an existent God yields infinitely awful damnation.[6]

Finally, there have been various salvos aimed at the third premise. The Allais (1953) and Ellsberg (1961) paradoxes, for example, are said to show that maximizing expectation can lead one to perform intuitively sub-optimal actions. So too the St. Petersburg paradox, in which it is supposedly absurd that one should be prepared to pay any finite amount to play a particular game with infinite expectation. Or one could insist that rational choices must be ratifiable (à la Jeffrey 1983 or Sobel 1996), and that the act of maximal expectation might not be. Moreover, while the expectation of wagering for God is infinite if we accept Pascal's earlier assumptions, as we will see, *so is the variance.* Expectation may not be a good guide to choice-worthiness when the variance is large, especially in a one-shot decision problem such as this, let alone when the variance is infinite—see Weirich 1984 and Sorensen 1994.

So Pascal's premises have come under heavy fire. Nonetheless, it has been generally assumed that his argument is *valid*—indeed, a number of critics have made a point of explicitly conceding this to Pascal (for example, Mackie (1982), Brown (1984), Rescher (1985), Mougin and Sober (1994), and most emphatically, Hacking (1994)). Recall how Pascal's reasoning, fleshed out in modern parlance, goes. Let p be your positive subjective probability for God's existence. Your expected utility for wagering for God is

$$\infty.p + f_2.(1 - p) = \infty.$$

As Pascal puts it, "our proposition is of infinite force." On the other hand, your expected utility for wagering against God is

$$f_1.p + f_3.(1 - p).$$

This is finite.[7] By the third premise, you should perform the act of maximum expected utility. Therefore, you should wager for God, concludes Pascal.

In section 3, developing and refining points made first by Jeffrey (1983) and Duff (1986) (see footnotes 8 and 11), I will argue that Pascal's reasoning is invalid. Even waiving the problems with his premises, his argument simply does not go through. This prompts the search for a more satisfactory reformulation of the argument that is valid. In section 4 I undertake this task, offering four such reformulations. I hope that in the process I will provide some illumination of the notion of infinite utility in general (and that my proposals for analyzing or replacing it will be of wider interest, with possible applications beyond the philosophy of religion). However, as I will argue in section 5, each of the reformulations is open to a new objection. Combining these results, I will argue that there is a fundamental dilemma that any version of the Wager must face.

3. Pascal's Argument Is Invalid

3.1 Mixed Strategies

Grant Pascal every premise of his argument. It is still not the case that wagering for God is rationally mandated. This will be the thrust of my attack on Pascal's Wager; arguing for it, and teasing out some further embarrassing results for Pascal, are my main purposes in this section.

Pascal's specious step is to assume that *only* the action of wagering for God gets the infinite expected utility. To see that this is not the case,

consider the following strategy: you toss a fair coin, and wager for God if the coin lands heads (probability ½); otherwise, you wager against God. By Pascal's lights, this strategy's expected value is the average of infinity and something finite:

$$½ (\infty) + ½ (f_1.p + f_3.(1 - p)).$$

This is infinite: $½ (\infty) = \infty$, and the second summand is finite. So we have found another way to get infinite expected value.[8] Now that we see the trick, we can run it again and again. Wager for God if and only if a die lands 6 (a sixth times infinity equals infinity …); if and only if your lottery ticket wins next week; if and only if you see a meteor quantum-tunnel its way through the side of a mountain and come out the other side … Pascal has ignored all these *mixed strategies*— probabilistic mixtures of the "pure actions" of wagering for and wagering against God—and infinitely many more besides.[9] And all of them have maximal expectation. Nothing in his argument favors wagering for God over all of these alternative strategies.[10]

But this still understates Pascal's troubles. For isn't *anything* that an agent might choose to do really a mixed strategy between wagering for and wagering against God, for some appropriate (rational subjective) probability weights? For whatever one does, one should assign some positive probability to winding up wagering for God. Even if you are currently an atheist, dear reader, you should assign positive probability to your wagering for God by the time you reach the end of this sentence (a probability greater by many orders of magnitude, I would hazard to say, than the probability of the meteor tunneling). In fact, I would hazard to say (and the next section will furnish an argument for my saying) that *every* rational agent's life is a constant series of such "gambles," with wagering for God as one of the outcomes. The probability of ending up wagering for God should be positive even for those who single-mindedly do all they can to wager *against* God—by practicing devil worship, say. The point generalizes to any course of action. By Pascal's lights, for every rational agent, every action has maximal expected utility.[11] It seems that we have here a "proof" that Leibniz was on the right track after all: in an important sense, this really is the best of all possible worlds!

3.2 Regularity

Let us pursue this line of attack still further. Call a probability function *regular* if and only if it assigns probability 1 only to logical truths (and

0 only to contradictions). 'Cautious' or 'undogmatic' would be more evocative words for this property, but I follow standard terminology here. Think of regularity as the converse of the usual requirement, honored as an axiom of the probability calculus, that all logical truths receive probability 1. Many authors assume, and for present purposes let us join them in assuming, if only for the sake of the argument, that a rational agent's probability function is always regular. Otherwise the agent displays a certain sort of dogmatism—total belief in some proposition that could be false, as far as logic is concerned—that would remain in the face of any future evidence.[12] In the words of Edwards et al. (1963, 211) in their own defense of regularity: "Keep the mind open, or at least ajar," and similar sentiments are endorsed by various writers from Jeffreys (1961) to Jeffrey (1983). Shimony (1970) shows that regularity is required in order for you to avoid susceptibility to a semi–Dutch Book: a series of acceptable bets for which there is no possible circumstance in which you enjoy a net gain, and some possible circumstance in which you suffer a net loss. Further proponents of (close relatives to) regularity include Kemeny (1955), Carnap (1963), Stalnaker (1970), Appiah (1985), and Lewis (1980, 1986) (at least for initial credence functions, and for less-than-perfectly-rational agents).

Note that if God's existence is not contradictory—and it had better not be if the Wager is to have a point—then Pascal should *welcome* regularity: it provides a snappy defense of his hitherto unargued-for premise that rationality requires you to assign positive probability to God's existence. Be that as it may, regularity forces the rational agent to regard *any* choice as a genuine gamble, with eventual wagering for God as one of the outcomes—for that eventuality is surely not a contradiction, and thus by regularity cannot be assigned probability 0. If you don't want to think about the Wager, go have a beer. By regularity, you should assign positive probability that you will wind up wagering for God nonetheless. And so it goes for *any* action that you might undertake. So the problem for the Wager is only intensified: regularity *requires* you to stay open-minded to your wagering for God, whatever you decide to do now; and according to many authors, rationality *requires* regularity.

Still, my appeal to regularity is really overkill. All I need is for you to assign positive probability, for whatever reasons you might have, to the prospect of eventually wagering for God by some non-Pascalian route—if regularity is not your reason, that doesn't matter. And my main point was made *before* I made my appeal to regularity: the invalid-

ity of Pascal's Wager had already been exposed. The coin toss strategy alone sufficed to make that point.[13] What followed was just further skirmishing.

Let me summarize. The strength of Pascal's argument is that it is insensitive to the exact value of the positive probability that figures in the expected value calculation: whatever it happens to be, the multiplication by infinity swamps it, thus yielding maximal expected utility for wagering for God. The undoing of Pascal's argument is that it is insensitive to the exact value of the positive probability that figures in the expected value calculation: whatever it happens to be, the multiplication by infinity swamps it, thus yielding maximal expected utility for any act whatsoever. What Pascal overlooked was that in opening the door to all the various positive probabilities for God's existence, he also let in all the various mixed strategies, with their various probability weights for wagering for and against God. That is, he let in everything.

3.3 Tie-Breaking

With each of infinitely many actions equally sanctioned by decision theory, it seems we have the predicament of Buridan's ass in spades. You might choose one of the actions arbitrarily, but *that* was hardly Pascal's advice! (Indeed, even if you happened to choose arbitrarily to wager for God, that would still not count as following his advice.) In the face of this multiplicity of acts that maximize your expectation, can you appeal to some other tie-breaking criterion?

Schlesinger (1994) offers one: "try and increase the probability of obtaining the prospective prize" (97). Of course, "the prospective prize" here is salvation. Schlesinger is suggesting that decision theory should be supplemented with a new principle. In our present case, it amounts to this: rationality requires you to perform the action that maximizes your probability of salvation. This clearly rules out the coin-tossing strategy, the die-tossing strategy, and all the other mixed strategies, since these have lower probabilities of your achieving salvation than outright wagering for God does. The principle is prima facie plausible, and Pascal might have done well to adopt it (though, see Sorensen 1994 for dissent). Note, however, that Schlesinger in no way undermines our objection: Pascal's Wager, as it stands, is invalid—period. For nowhere does Pascal appeal to Schlesinger's principle in his argument. The fact that there are *other* arguments in the neighborhood that are valid does not change that. Indeed, I will offer four such arguments in section 4. Moreover, we will see there how the expecta-

tions of the various mixed strategies can be distinguished, so that there will be no need for a further tie-breaking principle. Expected utilities can still carry the day.

In the meantime, it seems that Pascal's Wager as it stands not only fails, but fails in the worst possible way. All of a rational agent's actions apparently have exactly the same expected utility, ∞. Thus, all decisions turn out to be equally good according to that agent, and all practical reasoning turns out to be useless. Since practical reasoning is surely *not* useless, this is a reductio either of decision theory, or of infinite utilities understood naively in this way—in any case, it is a reductio of the use of such infinite utilities *in* decision theory. The irony is that Pascal, that champion of infinite utility, is often touted as being the father of decision theory. It seems that Lakatos's adage that "every research program is born refuted" has a confirming instance right here.

4. Reformulating the Wager

So it is all over for the Wager as it stands. Can we do better? An adequate reformulation of the Wager must meet the following requirements:

Requirement of Overriding Utility
The utility of salvation must completely override any of the other utilities that enter into the expected utility calculations, thus rendering irrelevant the exact value of the probability one assigns to God's existence. (We impose this requirement in order to uphold the spirit of the original argument—for otherwise we would not have a reformulation of it, but rather some quite different argument.)

Requirement of Distinguishable Expectations
We must be able to distinguish in expectation outright wagering for God from the various mixed strategies (based on the coin toss, die toss, and so on) discussed in section 3. In particular, the smaller the probability of winding up wagering for God, the smaller should be the expectation, so that one is rationally compelled to make that probability as high as one can.

I will now consider four strategies that I think meet these requirements. In each case, a somewhat delicate balance is struck: on the one hand, reaping the benefit of attributing an infinite utility to salvation (namely, the swamping effect in the expectation calculation that ren-

ders irrelevant the exact probability value); on the other hand, reaping the benefit of attributing a finite utility to it (namely, immunity to the objection that all gambles at wagering for God are equally attractive). To put it succinctly, infinitude will be given in each case a finite-looking gloss. The reformulations also provide another function, I hope: precisifying Pascal's argument where previously it was somewhat imprecise and ambiguous (for the term 'infinite' is both). Nonetheless, I will argue in section 5 that the reformulations face other difficulties.

4.1 Salvation Has Surreal Infinite Utility

Is there a way of telling apart the various infinite expectations that previously came out the same? And indeed, can we make precise what is even meant by the term 'infinite expectation'? I think that there are several such ways. For example, one can appeal to *nonstandard analysis*, and there are several ways of formulating it. Robinson (1966), Nelson (1987), and Lindstrøm (1988) are just some of the mathematicians who have given sound foundations to such a theory. The key idea is that there are nonstandard models of a first-order theory of the real numbers, containing so-called "hyperreals," with elements that behave like infinitesimals, and others that behave like infinite numbers. Skalia (1975), for example, shows how nonstandard models of the real numbers can be used in a "non-Archimedean" decision theory. Sobel (1996) also argues that we should remain open to the employment of hyperreals in decision theory. I applaud these approaches. However, since I find Conway's (1976) construction of what have come to be called the *surreal numbers* especially ingenious and user-friendly, and since it offers similar benefits, I will focus upon it instead.[14]

I begin with some brief expository remarks. Conway constructs new surreal numbers out of previously constructed surreal numbers according to two rules. First, every number is identified with two sets of previously constructed numbers, a "left" set and a "right" set, such that no member of the left set is greater than or equal to any member of the right set. The newly constructed number lies between the members of the left set and the members of the right set. (The gist of this idea will be familiar to those who know Dedekind's construction of the reals from the rationals.) Second, one number x is greater than or equal to another number y if and only if no member of x's right set is less than or equal to y, and x is less than or equal to no member of y's left set.

To get the construction off the ground, Conway begins at stage 0 with the number whose left and right sets are both empty, <∅,∅>. This

number is called '0', and by considering some of its properties one can easily show that it deserves the name. At stage 1 is constructed the number $<\{0\},\varnothing>$ with left set consisting of 0, and empty right set, called '1'; and another number with these sets reversed, called '–1'. Again, these names are well chosen. The construction continues along these lines, with new numbers being formed at each stage as appropriate left and right sets of numbers formed at previous stages.

Here we come to the crucial point. After infinitely many stages, we can define among other things the number whose left set is $\{0, 1, 2, \ldots\}$, and whose right set is empty. This is ω, the first infinite number to be constructed. Also at this stage comes $1/\omega$, an infinitesimal greater than 0 but less than any positive real number, namely $<\{0\},\{1, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \ldots\}>$. At the next stage comes $\omega - 1 = <\{0, 1, 2, \ldots\}, \{\omega\}>$, and $\omega + 1 = <\{0, 1, 2, \ldots, \omega\}, \varnothing>$ among others.[15] And so on. Repeating this process ad infinitum again, we eventually construct numbers such as $\omega/2$, 2ω, $\sqrt{\omega}$, ω^2, and ω^ω. The system of numbers that is finally produced is a totally ordered field—thus, it is closed under all the usual operations (addition, subtraction, multiplication, division, exponentiation, extracting roots, and so on), it is commutative under addition and multiplication, and all numbers can be compared in size.[16]

Now let Conway meet Pascal. Let the decision matrix be as before, except now identify the utility of salvation as an *infinite number in Conway's system.* But which one? It doesn't really matter: (with a qualification to be mentioned shortly) Pascal's argument will go through whichever one you pick. Indeed, we could leave it as a variable ranging over all infinite values. However, for definiteness, let's pick ω (on a given utility scale). The decision matrix is now:

	God exists	God does not exist
Wager for God	ω	f_2
Wager against God	f_1	f_3

Assume p is some positive, finite (as opposed to infinitesimal) probability[17] for God's existence. The expectation of wagering for God is

$$\omega.p + f_2.(1 - p),$$

which is infinite. The expected utility of wagering against God is

$$f_1.p + f_3.(1 - p),$$

which is finite. Wagering for God wins! Since we have assumed nothing about p other than its positivity and finiteness, we see that its exact value is irrelevant (modulo those assumptions), swamped as it is by the multiplication with ω. So we have met the Requirement of Overriding Utility. Moreover, we can now distinguish infinite expectations of various sizes: for example, $\omega/6$ is smaller than $\omega/2$, which is smaller than ω; and in general, a mixed strategy, with probability $q > 0$ for wagering-for and probability $(1 - q) > 0$ for wagering-against, has expectation

$$q(\omega.p + f_2.(1 - p)) + (1 - q)(f_1.p + f_3.(1 - p)),$$

which is less than the expectation of wagering-for, and indeed is a strictly increasing function of q. Thus, the Requirement of Distinguishable Expectations has been met. Notice also that we no longer need a "tie-breaking" principle, such as Schlesinger's (discussed in section 3.3), since the expectations are not tied. Rational choice can once again be a matter solely of maximizing expectation.

Our reformulation of Pascal's Wager is valid, and it concludes that wagering for God is rationally mandated. If you want to resist the conclusion, you must resist one of the premises: you must discredit either decision theory (which I will not do), or the revised decision matrix (which I will do in section 5), or the assumption that the probability of God's existence is positive and finite. Let's consider the last of these courses now.

There are two ways to go here. The first we have already seen: simply challenge the assumption that this probability is positive, with a reminder that this is not true of the atheist. The second way is more interesting (and this brings us to the qualification that I warned you of just before the decision matrix). Oppy (1990) suggests that this probability might not be finite, but *infinitesimal* instead. In support of his point, we surely should be prepared to countenance infinitesimal probability given that we are prepared to countenance infinite utility. After all, the infinitesimals turn out to be simply reciprocals of the infinite numbers in Conway's system. Indeed, infinitesimal probabilities seem to be necessarily connected to infinite utilities: for instance, you assign an infinitesimal probability $1/\omega$ to X if and only if you consider 1 unit of utility the fair price to pay for a bet that pays the infinite utility ω if X. Furthermore, one might even argue in favor of just such an assignment. For example, it is easy to generate infinitely many incompatible hypotheses about exactly which God you must wager for in order to achieve salvation (I mentioned the "many Gods objection" in

section 2). Lacking reasons that support some of these hypotheses over others, a flat probability distribution over them might seem to be in order—which is to say an assignment of infinitesimal probability to each one of them.

Once infinitesimal probabilities are allowed, the reformulated argument no longer goes through automatically: the infinitesimal probability can "cancel" the infinite utility so as to yield a finite expectation for wagering for God; and this may be exceeded by the expectation of wagering against God. For example, multiplying the infinitesimal probability $1/\omega$ by the infinite utility ω yields the finite value 1. And there is no guarantee that

$$1 + f_2.(1 - 1/\omega) \approx 1 + f_2, \text{ exceeds}$$

$$f_1.1/\omega + f_3.(1 - 1/\omega) \approx f_3.$$

since we did not assume that f_2 exceeds $f_3 - 1$.

So the reformulated argument does not catch in its net all agents who assign positive probability to God's existence: some agents who assign infinitesimal probability slip through. On the other hand, it does catch *some* agents who assign infinitesimal probability. After all, the infinite numbers come in a hierarchical ordering: roughly, an infinite number of a certain order is infinitely large compared to an infinite number of lower order, in the sense that the ratio of the former to the latter is infinite. And since they are reciprocals of infinitesimals, the infinitesimals display a similar hierarchy. For example, ω is infinitely large compared to $\sqrt{\omega}$ (their ratio is $\sqrt{\omega}$, which is infinite); and the infinitesimal $1/\sqrt{\omega}$ is infinitely large compared to $1/\omega$ (similarly). Now an infinite number multiplied by an infinitesimal of the same order does indeed yield a finite number (the case we considered in the previous paragraph). However, an infinite number of higher order multiplied by that same infinitesimal yields another infinite number. So a theological skeptic whose infinitesimal probability for God's existence (for instance, $1/\sqrt{\omega}$) happens to be "mismatched" with the utility of salvation (ω) like this may still feel the pull of our reformulated Wager. And given that the hierarchy of infinite numbers is itself infinite (there are infinitely many orders), the opportunity for such a mismatch is ample. Still, the point remains that infinitesimal probabilities can undermine the Wager.

Interestingly, Pascal seems to have taken care of this concern—and to my knowledge, this point has been overlooked. Indeed, I think he

deserves considerable credit for apparently *having a notion of infinitesimal probability years ahead of his time*—a somewhat indistinct notion, to be sure, but I would not want to fault him for lacking our modern-day rigorous formulation of it. He writes: "there is here … a chance of gain against a finite number of chances of loss … . wherever the infinite is and there is not an infinity of chances of loss against that of gain, there is no time to hesitate, you must give all." I take him to be ruling out infinitesimal values of p here, thus dispelling the concern. It does, however, make an earlier concern worse. It was questionable that rationality rules out a zero probability assignment to God's existence; all the more, it is questionable that rationality rules out all positive infinitesimal probability assignments as well.

There is a curious consequence of this reformulated version of the argument. Recalling my discussion in section 3, it seems to imply that we all get infinite expected utility whatever we do, as long as the probability of our winding up wagering for God is positive and finite. For any positive finite number multiplied by an infinite number yields an infinite number. So even if you are currently an atheist, dear reader, you should agree that you are nonetheless performing an act with infinite expectation in reading this sentence, since with positive finite probability you might wager for God before reaching the end of it—an expectation less than ω, of course, but infinite nonetheless. Still, a proponent of this reformulation might be prepared to bite this bullet, provided the *ordering* of expected utilities is right, and I have shown above that Pascal should be pleased with the ordering here.

4.2 Vector-Valued Value: Salvation Has Finite "Heavenly" Value

> *Infinite things alone—for example, eternity and salvation— cannot be equaled by any temporal advantage. We ought never to place them in the balance with any things of the world.*
> —Arnauld (1964, 357)

Suppose that there are two sorts of value: we might call them "earthly value" and "heavenly value" in order to have a handy way to refer to them. It is not their names that matter, but rather their structure. Suppose that the overall expected utility of one's life, rather than being a one-dimensional (scalar) quantity, is a two-dimensional (vector) quantity, of the form (x, y). And suppose that salvation has the maximal amount of heavenly value, which we will stipulate to be one unit—one eternal life in heaven, one might say, though the interpretation is not important. A probability p of salvation corresponds to p units of "heav-

enly expectation." We can thus picture one's overall expectation as a point in the plane, with the horizontal coordinate representing one's expectation in earthly value, and the vertical coordinate representing one's expectation in heavenly value—and thanks to our stipulation, the point lands in a horizontal strip of unit thickness. We can think of that overall expectation as a complex number of the form x + iy, with $0 \leq y \leq 1$ (without begging the question in favor of atheism with the usual reading of the horizontal component as "real," and the vertical component as "imaginary"!). We will not, however, be assuming anything of the structure of the complex numbers, beyond the way that they behave under addition and multiplication by real or infinitesimal constants.

Finally, suppose that any increase in heavenly expectation, however small, trumps any increase in earthly expectation, however large. The thought is that salvation is a good of such magnitude that any increase in the chance of its attainment is worth any earthly good. We have a so-called lexicographic ordering: when choosing between two actions, we compare firstly their heavenly expectations, preferring the action with greater heavenly expectation; if these are tied, we then prefer the action with the greater earthly expectation.[18] (Compare looking up words in a dictionary consisting exclusively of two-letter words.)

The decision matrix is now as follows, with e_1, e_2, e_3, and e_4 amounts of earthly value:

	God exists	God does not exist
Wager for God	$(e_1, 1)$	$(e_2, 0)$
Wager against God	$(e_3, 0)$	$(e_4, 0)$

In the second component we have weak dominance of wagering for God over wagering against God, and even *superdominance* (the worst payoff associated with wagering-for is at least as good as the best payoff associated with wagering-against, with at least one payoff genuinely better; cf. McClennen 1994). The expectation of any action can now be calculated by finding the earthly and heavenly expectations. Wagering for God has heavenly expectation

$$1.p + 0.(1 - p) = p > 0.$$

Wagering against God has zero heavenly expectation, and so it is automatically trumped by wagering for God, whatever positive value p has—even infinitesimal. Thus, the Requirement of Overriding Utility

has been met. And trumped also is any gamble that has wagering for God as an outcome with probability q < 1. For the heavenly expectation of such a gamble is q.p < p, so its earthly expectation is irrelevant—heavenly always trumps earthly. Wagering for God uniquely maximizes your expectation, just as Pascal wants, and the higher the probability of wagering for God, the higher the expectation. The Requirement of Distinguishable Expectations has been met.

It may be tempting to think that this lexicographic representation is equivalent to a surreal representation along the lines of section 4.1. I want to stress that there is *no* such equivalence. As we saw in section 4.1, infinitesimal probabilities could "cancel" with the infinite utility so that wagering for God was not the optimal act. But here, any infinitesimal probability for God's existence still dictates wagering for God, for even an infinitesimal amount of heavenly value trumps any amount of earthly value. This, by the way, is as close as I can come to vindicating Pascal's remark that one should wager for God even "if there were an infinity of chances, of which one only would be for you." In fact, Pascal could stipulate that any positive heavenly expectation exceeds even infinite earthly expectation: a tiny chance at salvation—even infinitesimal—is better than a guarantee of playing the St. Petersburg game. In that case, we could even allow e_1–e_4 to be infinite.

4.3 Salvation Has Finite Utility for an Infinite Period of Time

> *I'm beginning to understand eternity, but infinity is still beyond me.*
> —Cartoon caption in Harris 1989

Pascal thought of salvation as being incomparably better than any earthly pleasure—"an infinity of infinitely happy life," as we saw in section 2.[19] However, one could conceivably attribute infinite utility to even an earthly pleasure, provided that pleasure persisted forever—an infinity of *finitely* happy life, as we might say—and this possibility is not incoherent. Here the economist will be quick to point out that under the assumption that the agent discounts the future at a sufficient rate, such an infinitely protracted good will still yield a finite total utility upon integration over infinite time. Then let us not make this assumption: assume instead that the agent's discount rate is sufficiently small to yield an infinite total utility. For definiteness, suppose this discount rate is zero: the agent puts equal weight on periods of time of equal length, irrespective of how distant they are in the future. It is quite consistent to do this.

We need a device for comparing finite states of well-being that extend for infinite time. How do we recognize the superiority of a superb cognac to a mediocre cup of coffee—both "bottomless," as they say in some cafés—in the case that we can enjoy both forever? After all, the total utility is infinite in both cases.

Vallentyne (1993) offers a principle that gives the right verdict here (even though his concern is really utilitarianism, and aggregating utility over individuals):

> PMU*: An action, a1 *produces more utility* than action, a2, if and only if there is a time t such that for *any* later time t' the *cumulative* amount of utility produced by a1 up to t' is *greater* than that produced by action a2 up to t'. (215)

Pick any time t that you like (even t = 0 will do). It is certainly true that, for any later time t', the cumulative amount of utility produced by drinking the cognac up to t' is greater than that produced by drinking the coffee up to t'. For definiteness, suppose that drinking the cognac produces 2 units of utility per unit time, and that drinking the coffee produces 1 unit of utility per unit time, on some suitable scale. Then the cumulative amount of utility produced by drinking the cognac up to t', namely $2t'$, is greater than that produced by drinking the coffee up to t', namely t'. So drinking the cognac produces more utility than drinking the coffee.

PMU* allows us to make qualitative, but not quantitative, judgments of betterness: it allows us, when its conditions are met, to make verdicts of the form "a1 is better than a2," but it does not tell us how *much* better a1 is than a2. Suppose that one can also enjoy forever a schnapps that is slightly more exquisite than the cognac: it rewards one with 2.1 units of utility per unit time on the same scale. PMU* correctly ranks the three infinite pleasures, but it does not tell us that the liquors are closer to each other in quality than they are to the coffee, let alone how much so.

But there is a way to make such quantitative comparisons. An approach familiar to economists is to consider the *long-run average* utility of each of these: calculate the total utility of each up to time t, for various t; divide this in each case by t; then take the limit as t tends to infinity. The total utility up to time t of the cognac is 2t; dividing this by t yields 2; the limit as t tends to infinity of 2 is 2, namely, the long-run average of the cognac. This exceeds the long-run average utility of the coffee, namely 1, thus accounting for our preference for the cognac.

The long-run average of the schnapps is 2.1, which is better still, but only slightly.

Suppose that salvation consists of a finite pleasure over infinite time—something that you accord, say, one unit of utility for each unit of time. (We can always rescale the utilities to make this so.) Then of course the long run average utility of salvation is 1. And suppose that in the case that God does not exist, or if you wager against God, you get some earthly rewards, but only for the finite amount of time until your death. Forever after you get zero units of utility for each unit of time. (We can always choose a utility scale so as to get this value, too.) Then whatever happens up till your death makes no contribution to the long-run average: the subsequent infinite period of zero utility overwhelms it.

Here is the decision matrix, now with long-run average utilities rather than total utilities:

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

Note that we have long-run averages in this matrix that have the same values as the "heavenly" components in the previous matrix. And since it was those "heavenly" components that did all the work in determining what one should do, the calculations look similar here: the expected utility of wagering for God is p, which exceeds the expected utility of wagering against God, 0. Furthermore, outright wagering-for exceeds in expectation any gamble at wagering-for, and in general the higher the probability q of wagering-for, the higher the expectation: q.p is an increasing function of q.[20] All this is so even if we allow infinitesimal probabilities. The Requirements of Overriding Utility and of Distinguishable Expectations have both been met.

4.4 Salvation Has Finite Utility

We might say that Pascal held the utilities fixed and "solved for" who should wager for God (that is, everyone but strict atheists who assign zero probability to God's existence and, we might now add, certain near-atheists who assign it infinitesimal probability). But we could instead hold fixed a set of people, and solve for those utilities that would mandate wagering for God for all people in that set. In particular, consider the set *S* of all people who ever lived and who ever will live who assign positive, finite probability to God's existence. *S* is clearly a

huge set of people, including even someone who assigns it probability one-in-a-googolplex, if such a person ever existed/exists. (*S* does not include the strict atheists or near-atheists, but we already knew that the Wager had no hope of convincing them.) Huge though *S* is, it is surely finite, as the human race will surely not persist forever. Now consider the smallest probability that anyone in *S* assigns to God's existence. Call this p_{min}. It represents, as it were, Pascal's hardest sell—the assignment of the most skeptical person in *S* as far as God's existence is concerned.

Now we need not assume that salvation brings infinite utility at all. A finite utility *f* (on a chosen scale) will suffice, provided *f* is sufficiently large.[21] In that case salvation could be a finite, finitely happy life. The decision matrix becomes:

	God exists	God does not exist
Wager for God	f	f_2
Wager against God	f_1	f_3

All the numbers in the matrix are now finite. How large is *f*? We merely need it to be sufficiently large that the expectation of wagering for God exceeds that of wagering against God even for the most skeptical person in *S:*

$$f.p_{min} + f_2.(1 - p_{min}) > f_1.p_{min} + f_3.(1 - p_{min}).$$

Thus, even in this hardest case (so to speak), wagering for God is the act of maximal expectation. There will obviously be a range of candidates for *f* with the required property, and any of them will yield a valid reformulation of the Wager. For definiteness, we could pick one just slightly greater than the "break even" point, the value at which the two expectations are exactly equal. But nothing more really needs to be said about the size of *f*; its exact value does not matter. For that reason, we could even treat it as a variable, ranging over all utility values that are sufficiently large.

Again, this reformulation meets the Requirement of Overriding Utility. Thanks to *f*'s largeness, all the other utilities in the matrix are overridden. And the exact value of your positive probability for God's existence is irrelevant—whatever it is, wagering-for exceeds wagering-against in expectation. For with probability p (guaranteed to be $\geq p_{min}$) for God's existence, your expectation of wagering-for is

$$f.p + f_2.(1 - p)$$

while your expectation of wagering-against is

$$f_1.p + f_3.(1 - p).$$

By hypothesis, $f.p$ is so large that the former value exceeds the latter.

Furthermore, a mixed strategy, with probability q > 0 for wagering-for and probability $(1 - q) > 0$ for wagering-against, has expectation

$$q(f.p + f_2.(1 - p)) + (1 - q)(f_1.p + f_3.(1 - p)),$$

which is less than the expectation of wagering-for, and indeed is a strictly increasing function of q. Thus, the Requirement of Distinguishable Expectations is met. Ironically, it is in this sense that wagering for a God that offers a sufficiently large finite reward is rationally required, while wagering for a God that offers an infinite reward of ∞ is not (as we saw in section 3). Moreover, this finite reformulation parries a major objection (noted in section 2) that decision theorists such as Jeffrey and McClennen had to the original wager: that the very notion of infinite utility is suspect.

5. A Problem for the Reformulations—And a Dilemma for *Any* Reformulation

Four reformulations of Pascal's Wager are now before you. These proposals, I submit, yield valid arguments for wagering for God, where Pascal's argument was invalid. Yet still they may not meet Pascal's needs.

The problem in each case, not altogether surprisingly, concerns the utility of salvation. For while all the proposals meet the Requirement of Overriding Utility, they still do not seem adequately to capture Pascal's reasoning. Recall that according to him, "[u]nity joined to infinity adds nothing to it …. the addition of a unit can make no change in its nature" or symbolically, $\infty + 1 = \infty$. Likewise, $\infty + 2 = \infty$, $\infty + 3 = \infty$, and indeed $\infty + x = \infty$ for all positive x. Let us call this property of ∞ *reflexivity under addition*. When the utility of salvation is reflexive under addition, one cannot increase it by adding something to it. We can see why Pascal would regard the utility of salvation to be reflexive under addition: as I noted in section 2, he thought of salvation as the best possible thing. But if that utility is a surreal infinite number such as ω, or a long run average of 1, or a finite number f, then adding 1 (or 2, or 3, or indeed any positive x) to it *does* increase it; these quantities are not reflexive under addition. Salvation, then, is no longer the best possible thing after all.

It is an apparent virtue of the two-dimensional representation, I think, that it is less obviously susceptible to this objection. Of course, the objection would have to be restated so that the addition here is well-defined: it makes no sense to add a scalar to a vector. The natural generalization, to the two-dimensional case, of reflexivity under addition is that one cannot increase the utility of salvation by adding another vector to it. To be sure, one *can* increase that utility by increasing its earthly component: for example, adding $(1, 0)$ to $(e_1, 1)$ yields $(e_1 + 1, 1)$, which is a greater utility than $(e_1, 1)$. Strictly speaking, then, reflexivity under addition fails again. But notice that no addition can raise the heavenly component above its maximal value of 1, and this is the "trumping" component in the lexicographic ordering. The increase in utility here, then, is in this sense comparatively negligible, and thus the failure of reflexivity under addition is comparatively negligible. We might even maintain that the earthly component is determined solely by one's rewards in one's earthly life, and that the heavenly component concerns something that happens (or not) thereafter. What is at issue with salvation, we might insist, is the heavenly component, which is maximal. In this important sense, we might say that salvation is as good as it gets as far as the two-dimensional representation is concerned.

The objection would have to go more along these lines: God settled for just *two* dimensions of value, when he could have created three, or four, or … And we could then proceed much as before: any gain in the third dimension trumps any losses in the first two, and so on. Suppose, for example, that we have three dimensions of value. The two-dimensional strip that previously contained all expectations of the form (x, y) is now identified with the set of triples of the form $(x, y, 0)$. The utility of salvation corresponds to a vector of the form $(e_1, 1, 0)$. Understood this way, the natural generalization of reflexivity under addition—so that the notion applies at all—fails, this time in an important way. Adding a finite vector to the utility of salvation can make a big difference, indeed an overwhelming difference—for example, $(e_1, 1, 0) + (0, 0, 1) = (e_1, 1, 1)$, which is *much* better than the utility of salvation itself. So the utility of salvation is far from being reflexive under addition after all: adding a finite amount of higher-dimensional value makes a huge difference. Again, salvation is not the best possible thing after all—it's not even close.[22]

The problem for each of the reformulations stems from our giving, as I said, infinitude a finite-looking gloss. Alas, I see no way of squaring this with Pascal's view of infinity. "The finite is annihilated in the pres-

ence of the infinite, and becomes a pure nothing," he writes. But in each of the reformulations of the utility of salvation, the finite is *something*. To be sure, the utility of salvation was carefully chosen to swamp all other terms in the expectation calculations, and when it comes to merely ordering expectations, swamping is as good as annihilation. Still, the chosen utility for salvation, in turn, is not merely bettered, but *swamped* to the same degree by *another* conceivable utility: for instance, ω^2 stands to ω as ω stands to 1; and so on. And that other utility, in turn, is swamped by still another (ω^3, say), and so on *ad infinitum*—an "infinitum" of the form that Pascal *would* recognize! Far from being the best possible thing, salvation isn't even close; in fact, in the eyes of Pascal it becomes a pure nothing. It is hardly surprising, then, that the notion of infinity that he envisages is reflexive under addition. At least that way infinitude stays infinite-looking.

And yet as we saw in section 3, a kindred property of ∞ is the undoing of the Wager: $\infty . x = \infty$ for all positive, finite x. Let us call this property of ∞ *reflexivity under multiplication*. Of course, it was just this property that we exploited in showing that all the mixed strategies, with their various weights for wagering for and against God, have the same expected utility as outright wagering for God.

Reflexivity under multiplication, however, ought to be desirable to Pascal when the multipliers are greater than 1: we have $\infty . 2 = \infty$, $\infty . 3 = \infty$, and so on, and this is all to the good. In fact, it is all to the *ultimate* good, since once again the utility of salvation is not bettered. So really Pascal should want to be selective about which reflexivities hold of the utility of salvation: reflexivity under multiplication *by positive, finite probabilities* is a bad thing, since it opens the door to all the mixed strategies; reflexivity under multiplication *by numbers greater than 1* is a good thing, since it underscores the maximality of the utility of salvation (much as reflexivity under addition does).

Thus Pascal, and any advocate of an argument in the spirit of his Wager, faces a dilemma. If the utility of salvation is *both* reflexive under addition and under multiplication by positive, finite probabilities, as ∞ 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). But how do we give infinitude a finite-looking gloss with respect to multiplication by positive, finite probabilities, but an infinite-looking gloss with respect to addition? Said another way (given the correspondence between multiplication and repeated addition), how do we give infinitude a finite-looking gloss with respect to multiplication by positive, finite probabilities, but an infinite-looking gloss with respect to multiplication by numbers greater than 1? How do we represent the utility of salvation in a way that is sufficiently nuanced to make the distinctions that Pascal wants, but not the distinctions that he doesn't want?

It might seem that the dilemma would be resolved if there were some maximum utility level that a human could achieve, some "saturation point" beyond which additional rewards made no perceptible difference. This is plausible for most actual people and monetary wealth, for example: there is some degree of affluence so great that accruing further dollars does not improve one's situation. It is possible that for all people there is such a saturation point even when it comes to salvation. Perhaps it is of the essence of a human to be finite, and perhaps a finite being cannot reap an infinite reward; perhaps an infinite reward can only be finitely appreciated by a human. Suppose, for example, that this saturation point is f units of utility on some chosen scale. Then there could be no complaint against God for his making salvation worth f—any additional reward would go unappreciated. Unity joined to f adds something to it, mathematically speaking, but it adds nothing that makes a difference to us. Likewise for doubling f, or tripling it, and so on. In that case, the two problems would apparently be solved with a single stroke: we could distinguish in expectation the various gambles at wagering for God ($f/2$ is smaller than f, and so on), while salvation would be the maximal good that we could realize. We could generalize this point to the other reformulations. Simply suppose that our capacity to enjoy a reward, while not necessarily finite, is nevertheless represented by a quantity that is not reflexive under addition or multiplication: the saturation point could be ω, or an infinite long-run average of 1, or the vector-valued quantity (e_1, 1).

But Pascal would not obviously be out of the firing line yet. For now the complaint could shift to the level of the saturation point, the objections merely relocated. Why would God create us with a saturation point at all? And having done so, it would be swamped by another

choice of saturation point, which would be swamped by another one, … *ad infinitum.*[23] It matters little whether it is the utility of salvation or the saturation point that is given a finite-looking gloss: either way, we fall far, far short of Pascal's promise of our "gaining everything" if we win the Wager.

6. Conclusion

It seems, then, that Pascal's Wager and all the reformulations of it that I have considered face a serious problem. Moreover, I believe that it is a problem that runs deep, not one that will go away with some clever tinkering. For I see no prospects for characterizing a notion of the utility of salvation that is reflexive under addition *without* being reflexive under multiplication by positive, finite probabilities, or reflexive under multiplication by numbers greater than 1 *without* being reflexive under multiplication by positive, finite probabilities. Yet it seems that nothing less will salvage Pascal's reasoning. So we are left with a dilemma. If the utility of salvation *is* reflexive under both addition and multiplication by positive, finite probabilities (as in Pascal's original argument), wagering for God will be just one of many equally rational courses of action, and our choice among them will be arbitrary. If the utility is *not* reflexive under either addition or multiplication by positive, finite probabilities (as in my reformulations of the argument), salvation will be so far from being the best thing possible as to be unsuitable for Pascal's theology. *I* wager that any future version of the argument will succumb to this dilemma.

California Institute of Technology

References

Allais, Maurice. 1953. Le Comportment de l'Homme Rationnel Devant la Risque: Critique des Postulats et Axiomes de l'École Américaine. *Econometrica* 21:503–46.

Arnauld, A. 1964. *The Art of Thinking (Port Royal Logic).* Translated by J. Dickoff and P. James. Indianapolis: Bobbs-Merrill.

Appiah, Anthony. 1985. *Assertion and Conditionals.* Cambridge: Cambridge University Press.

(Saint) Augustine. 1966. *The Catholic and Manichean Ways of Life (De Moribus Ecclesiae Catholicae et De Moribus Manichaeorum).* Translated by Donald A. Gallagher and Idella J. Gallagher. Washington, D.C.: The Catholic University of America Press.

Broome, John. 1995. The Two-Envelope Paradox. *Analysis* 55 (1): 6–11.

Brown, Geoffrey. 1984. A Defense of Pascal's Wager. *Religious Studies* 20: 465–79.

Cain, James. 1995. Infinite Utility. *Australasian Journal of Philosophy* 73:401–4.

Carnap, Rudolf. 1963. Carnap's Intellectual Autobiography. In *The Philosophy of Rudolf Carnap*, The Library of Living Philosophers, vol. 11, edited by Paul Arthur Schilpp. La Salle, Ill.: Open Court.

Chalmers, David J. 2002. The St. Petersburg Two-Envelope Paradox. *Analysis* 62:155–57.

Chernoff, Herman, and Lincoln E. Moses. 1959. *Elementary Decision Theory*. New York: Wiley.

Clark, M., and N. Shackel. 2000. The Two-Envelope Paradox. *Mind* 109:415–42.

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

de Finetti, Bruno. 1972. *Theory of Probability*. New York: Wiley.

Diderot, Denis. 1875–77. *Pensées Philosophiques*, 59, *Oeuvres*. Edited by J. Assézat, vol. 1.

Duff, Antony. 1986. Pascal's Wager and Infinite Utilities. *Analysis* 46:107–9.

Edwards, W., H. Lindman, and L. J. Savage. 1963. Bayesian Statistical Inference for Psychological Research. *Psychological Review* 70:193–242.

Ellsberg, D. 1961. Risk, Ambiguity and the Savage Axioms. *Quarterly Journal of Economics* 25:643–69.

Fishkind, D., J. D. Hamkins, and B. Montero. 2002. New Inconsistencies in Infinite Utilitarianism: Is Every World Good, Bad or Neutral? *Australasian Journal of Philosophy* 80:178–90.

Hacking, Ian. 1994. The Logic of Pascal's Wager. In Jordan 1994.

Harris, Sidney. 1989. *Einstein Simplified: Cartoons on Science and Scientists*. New Brunswick: Rutgers University Press.

Hausner, M. 1954. Multidimensional Utilities. In *Decision Processes*, edited by R. M. Thrall, C. H. Coombs, and R.L. Davis (New York: Wiley).

James, William. 1956. The Will to Believe. In *The Will to Believe and Other Essays in Popular Philosophy*. Mineola, N.Y.: Dover Publications.

Jeffrey, Richard C. 1983. *The Logic of Decision*. 2d ed. Chicago: University of Chicago Press.

Jeffreys, Harold. 1961. *Theory of Probability*. 3d ed. Oxford: Oxford University Press, Clarendon Press.

Jordan, Jeff, ed. 1994. *Gambling on God: Essays on Pascal's Wager*. Lanham, Md.: Rowman & Littlefield.

Jordan, Jeff. 1998. Pascal's Wager Revisited. *Religious Studies* 34:419–31.

Kemeny, J. 1955. Fair Bets and Inductive Probabilities. *Journal of Symbolic Logic* 20:263–73.

Lewis, David. 1980. A Subjectivist's Guide to Objective Chance. In *Studies in Inductive Logic and Probability*, vol. 2, edited by Richard C. Jeffrey. Berkeley: University of California Press.

———. 1986. Probabilities of Conditionals II. *Philosophical Review* 95:581–89.

Lindstrøm, Tom. 1988. Invitation to Non-Standard Analysis. In *Nonstandard Analysis and its Applications*, edited by Nigel Cutland. London: London Mathematical Society, Student Texts 10.

Mackie, J. L. 1982. *The Miracle of Theism*. Oxford: Oxford University Press.

Martin, Michael. 1990. *Atheism: A Philosophical Justification*. Temple University Press.

McCall, Storrs, and D. M. Armstrong. 1989. God's Lottery. *Analysis* 49:223–24.

McClennen, Edward. 1994. Finite Decision Theory. In Jordan 1994.

Mougin, Gregory, and Elliot Sober. 1994. Betting Against Pascal's Wager. *Noûs* 28:382–95.

Nelson, Edward. 1987. *Radically Elementary Probability Theory*. Annals of Mathematics Studies. Princeton: Princeton University Press.

Norton, J. D. 1998. When the Sum of Our Expectations Fails Us: The Exchange Paradox. *Pacific Philosophical Quarterly* 79:34–58.

Oppy, Graham. 1990. On Rescher on Pascal's Wager. *International Journal for Philosophy of Religion* 30:159–68.

Palacios, M. Asin. 1920. *Los Precedentes Musulmanes del 'Pari' de Pascal*. Santander.

Pascal, Blaise. 1948. *Pascal's Pensées*. Translated by W. F. Trotter. London: Dent; New York: Dutton.

Pratt, John W., Howard Raiffa, and Robert Schlaifer. 1995. *Introduction to Statistical Decision Theory*. Cambridge: MIT Press.

Rawls, John. 1971. *A Theory of Justice*. Cambridge: Harvard University Press.

Rescher, Nicholas. 1985. *Pascal's Wager*. Notre Dame: University of Notre Dame Press.

Robinson, Abraham. 1966. *Non-Standard Analysis*. Amsterdam: North Holland.

Ryan, John. 1994. The Wager in Pascal and Others. In Jordan 1994.

Savage, L. J. 1954. *Foundations of Statistics*. New York: Wiley.

Schlesinger, George. 1994. A Central Theistic Argument. In Jordan 1994.

Shimony, Abner. 1970. Scientific Inference. In *The Nature and Function of Scientific Theories*, edited by Robert Colodny. Pittsburgh: University of Pittsburgh Press.

Skalia, H. J. 1975. *Non-Archimedean Utility Theory*. Dordrecht: D. Reidel.

Skyrms, Brian. 1980. *Causal Necessity*. New Haven: Yale University Press.

Sobel, Howard. 1994. Two Envelopes. *Theory and Decision* 45:69–96.

———. 1996. Pascalian Wagers. *Synthese* 108: 11-61.

Sorensen, Roy. 1994. Infinite Decision Theory. In Jordan 1994.

Stalnaker, Robert. 1970. Probabilities and Conditionals. *Philosophy of Science* 37:64–80.

Thrall, Robert M. 1954. Applications of Multidimensional Utility Theory. In *Decision Processes*, edited by R.M. Thrall, C. H. Coombs, and R. L. Davis. New York: Wiley.

Vallentyne, Peter. 1993. Utilitarianism and Infinite Utility. *Australasian Journal of Philosophy* 71:212–17.

Vallentyne, Peter, and Shelly Kagan. 1997. Infinite Value and Finitely Additive Value Theory. *Journal of Philosophy* 94:5–27.

Weirich, Paul. 1984. The St. Petersburg Gamble and Risk. *Theory and Decision* 17:193–202.

Notes

[1] Pascal's presentation of the Wager is somewhat obscure, even quirky in various ways, frustrating a definitive reconstruction. I do want to insist, however, that I am adopting a perfectly standard interpretation of §233 in the *Pensées*, if cast in the anachronistic terminology of modern Bayesian decision theory (and that casting too is standard).

[2] You may say that Pascal's conclusion is really "Rationality requires you to *believe* in God." But perhaps one cannot simply believe in God at will; and rationality cannot require the impossible. Pascal is well aware of this objection: "[I] am so made that I can-

not believe" (1948, §233), his imaginary interlocutor poignantly replies. But Pascal maintains that one can *resolve to believe* in God—one can *cultivate* such belief. I will adopt the phrase "wagering for God" as shorthand for "resolving to believe in God."

[3] If you wager against God at t, you may still wager *for* God at a later time t′, and I presume that Pascal would regard you as getting infinite utility if your belief in God, or your resolve to believe, then persists until your death. One may fairly object that you might permanently lose such belief or such resolve, and thus the putative infinite utility, so really the decision matrix should be more complicated than Pascal's. This amounts to an objection to premise 2, and as such joins a long list of objections to his premises to which I will turn shortly. Without wishing to dismiss these objections, my primary concern lies elsewhere.

[4] Cf. Augustine's discussion of "the supreme good of man" in his 1966, especially book 1.

[5] Earlier in the *Pensées* (§229) he also writes: "nothing would be too dear to me for eternity."

[6] Martin (1990) reads Pascal himself this way. However, Pascal says: "The justice of God must be vast like His compassion. Now justice to the outcast is less vast … than mercy towards the elect" (65). I take Pascal to be suggesting here that f_1 is not $-\infty$, as does Sobel (1996).

[7] These points carry through even if you think that what you do is not independent of whether God exists. Maybe God helps people wager for him, so that P(God exists|you wager for God) > P(God exists|you wager against God). Still, the expected utility calculations are as before, provided the first conditional probability is positive: infinite for wagering-for, finite for wagering-against. In the rest of the paper, I will make the simplifying assumption of independence, but at no point will this be essential.

Moreover, I doubt that using some version of causal decision theory instead would really change matters. We would just replace the assumption that p is positive with the same assumption about whatever probability replaces it—the probability of a counterfactual, an imaged probability, or what have you.

[8] Jeffrey (1983, 153) makes a similar point in his discussion of infinite utility:

> If the agent takes act 1 to bestow probability .99 on the prospect of heaven, and he takes act 2 to bestow probability .01 on that prospect, and if he takes all other consequences of the acts to have finite desirabilities, it seems clear that the agent would and should strongly prefer act 1 over act 2. On the other hand, in the Bayesian account of the matter, the agent is taken to rank the two acts together at the top of his preference scale, since each of them has infinite expected desirability; for we have
>
> $$.99 \times \infty = .01 \times \infty = \infty$$
>
> ...

We will shortly see just how much Jeffrey's point can be generalized.

[9] I follow here the terminology of such authors as Chernoff and Moses (1959) and Pratt, Raiffa, and Schlaifer (1995).

[10] I claim to have shown the argument to be invalid. You may think that I have really disputed premise 2, augmenting the decision matrix to include various actions that Pascal did not consider. In defense of my claim, let me note that it is perfectly standard in decision theory to take a set of pure strategies as given, with their various corresponding pay-offs tabulated in a decision matrix; all the mixed strategies then come for free, as it were, their expected utilities thereby determined. Obviously, it would be impossible to list all of the mixed strategies as extra rows in the matrix, for there are uncountably many of them.

Of course, we have recast Pascal's somewhat obscure text in modern guise—see footnote 1—and there may be some indeterminacy regarding how best to diagnose the

flaw in his original argument. It may not matter much at the end of the day whether we judge the argument to be invalid or to have a false premise—either way, the argument is unsound—although I have given reasons for preferring my diagnosis. However we classify it, the flaw deserves attention, as do its consequences.

[11] Duff (1986) has made the essentials of this point before me. I intend my ensuing discussion to develop his contribution further. I agree with (and in some cases reiterate) almost everything in his version of the argument, with this small caveat: we should make clear that all the probabilities and expectations at issue are *those of a particular agent*. Duff says, for example:

> No course of action can make it absolutely certain that I will not come to believe in God: therefore, every course of action has an infinite expected value—the infinite value of believing in God multiplied by the probability that God exists, and by the probability that I will come to believe in Him. (108)

We should remember that notions such as "certain", "the probability," and "expected value" are subjective, and that these observations will not hold for all agents. That is why I do some extra work in the next section in an attempt to arrive at similar observations at least for all rational agents.

Also, as will become clear in section 4, I disagree with the moral that Duff draws at the end: "This might look like a *reductio ad absurdum* of the Wager. It might more particularly suggest that there is something wrong with trying to capture infinity with probability" (109). Indeed, I will devote much space to "trying to capture infinity" decision-theoretically.

[12] I assume here that belief revision in the face of such evidence goes by conditionalization, Jeffrey conditionalization, or some other rule that keeps probability assignments of 0 and 1 fixed.

[13] And what if the coin toss lands the "wrong" way, dictating that you wager against God? Running Pascal's argument one more time, doesn't that mean that you *then* do worse than you would have if the coin had landed the "right" way, and you are really back where you started: you should act as if the coin had landed the "right" way, and wager for God outright? Not so. First, the best strategy need not result in the best consequences: taking out fire insurance for your house is a good strategy even if your house does not burn down (in which case you would have done better not taking out the insurance). Even Pascal could be prepared to admit that the strategy he advocates may not have the best consequences (namely, if God does not exist). But second, and more importantly, if the coin lands the "wrong" way, your expectation does not change. By Pascal's lights, you still enjoy infinite expectation whatever you do next.

[14] Moreover, many other philosophical discussions involve the notion of infinity, and of infinitesimal probability, and I suspect (though I cannot argue here) that surreal numbers may clarify or illuminate them also. I am thinking, for example, of the recent spate of articles on the two envelope paradox (for instance, Sobel 1994, Broome 1995, Norton 1998, Clark and Shackel 2000, Chalmers 2002); the interpretation of probability known as "hypothetical frequentism"; Lewis's 1980 discussion of the Principal Principle; Skyrms 1980 on causal necessity; Savage 1954 and de Finetti 1972 on countable additivity; McCall and Armstrong 1989 on "God's Lottery"; recent discussions of infinite utilities in utilitarianism, as found in Vallentyne 1993, Cain 1995, Vallentyne and Kagan 1997, Fishkind et al. 2002; and so on.

[15] Note that $\omega + 1$ is greater than ω (by 1). Thus, we part company with Pascal when he says, "Unity joined to infinity adds nothing to it." More on this below.

[16] See Conway (1976), chap. 1. One reason I have for preferring Conway's system to those of, say, Cantor or Zermelo-Fraenkel, is the following. Starting with the familiar infinite number ω, both Cantor and Zermelo-Fraenkel generate larger infinite numbers ($\omega + 1$, 2ω, ω^2, and so on). However, they do not generate *smaller* ones. The beauty of

Conway's system is that he generates with equal ease smaller numbers such as $\omega - 1$, $\omega/2$, $\sqrt{\omega}$, and so on—and it is just this feature that we will appeal to in order to evade my main objection to Pascal's original wager.

[17] Strictly speaking, 0 is an infinitesimal (the only real one, in fact), so a probability that is non-infinitesimal is automatically positive. But typically when I speak of "infinitesimals," I mean the positive ones. Context should make my usage clear.

[18] The concept of lexicographic utility originates with Hausner (1954) and Thrall (1954), who established the existence of a lexicographic utility function. This sort of ranking of alternatives will be familiar to readers of Rawls (1971), who calls it a "lexical order."

[19] Maybe something could be incomparably better than any earthly pleasure, and yet still lead to only a finite total of utility: suppose it behaves like a delta function, providing an instantaneous pulse of infinite utility. Pascal, of course, did not think of salvation that way.

[20] If p is finite, the zero pay-offs could be replaced with infinitesimal payoffs without harming the argument.

[21] Mougin and Sober (1994) consider finite versions of the Wager, leaving it open whether Pascal himself offered such a version: "Pascal's theology allows him to describe the *payoffs* that accrue to the theist and to the atheist. Heaven is of great (perhaps infinite) value" (382). I do want to insist, however, that Pascal's talk of "an infinity of an infinitely happy life" rules out our attributing *to him* a finite version. Sobel (1996) also canvasses a finite wager, but it is tailored to a particular person ("'I'") with a particular probability/utility profile, and it is complicated by a concern with not flouting rationality: "Let there be in my view parity between ⅓ chance of eternal bliss, and a certainty of not detracting from my rationality to the extent that willful belief in God would do, so that 'I' value eternal bliss twice as much as 'I' disvalue that detraction" (41). My finite wager is thus more general. And Jordan (1998) advocates that we consider finite versions of the Wager. He then surveys their possible audience, a list that includes: "(5) *the convinced atheist*: one who believes that $0.5 > Pr(G) > 0$, or would so believe if s/he were to think about it" (428), where "*G*" stands for "God exists." Jordan does not give us any quantitative information about the utility of salvation, so I am quite at a loss to explain why he thinks that "the *bottom third* of those described by (5) would be, no doubt, beyond the persuasive scope of a finite wager since their probability assessments of theism are significantly less than one-half" (428, my italics). In any case, assuming (as I think plausible) that no human genuinely assigns infinitesimal probability to God's existence, *all* of those described by (5) would be within the persuasive scope of the finite wager(s) that I suggest.

[22] There is a related problem that can be stated in terms of the Principle of Sufficient Reason—a principle associated with Leibniz rather than Pascal, to be sure, but it still provides a way of sharpening the point. We can understand (even if we don't agree) why Leibniz thought that this is the best of all possible worlds: God would not have a sufficient reason for creating any other world. If he created instead, say, the seventeenth best world (if there is such a thing), he would apparently be acting arbitrarily—why *that,* we might ask, rather than the sixteenth best? Likewise, it is natural to think of salvation as being the best of all possible rewards, but according to each of the reformulations it is not. God, then, is portrayed as acting arbitrarily in making it have the utility that it has and not some other. Why, we might ask, should the utility of salvation be ω rather than $\omega+1$, or $\omega+2$, or … (on a given scale)? Why f rather than $f+1$, or $f+2$, or …? Why should salvation have a long-run average of 1, rather than 2, or 3, or …? Or 2 dimensions of value rather than 3, or 4, or …? To be sure, there are worse things than acting arbitrarily—just ask Buridan's ass—and maybe even God can so act, consistently with his nature. But notice that such questions did not even arise when we understood salvation

as having a utility of "infinity," where "[u]nity joined to infinity adds nothing to it." Pascal would simply note that there is nothing arbitrary about God bestowing a reward of ∞ rather than $\infty + 1$, because they are one and the same.

[23] Furthermore, there would apparently be no sufficient reason why the saturation point should have one value rather than another. The choice of f rather than $f + 1$ or $f + 2$ or … seems arbitrary; likewise for the other reformulations.