7-1-1994

Discussion: On Pascal's Wager And Infinite Utilites

John Byl

ON PASCAL'S WAGER AND INFINITE UTILITIES

John Byl

In this paper I discuss some objections to Pascals' Wager based on the notion of an infinite utility. It is alleged that infinite utilities result in decisional and mathematical indeterminacies that invalidate Pascal's Wager. Although various resolutions to these objections have been proposed, these in turn have shortcomings. It is argued that the indeterminacies can be readily avoided by treating the infinities as limits. It is suggested that, in situations where only one bet can be placed, the expected utility should be replaced by the most probable average utility. By this standard the Wager is found to fall short if the probability of God's existence is taken to be small.

Pascal's Wager purports to demonstrate that it is rational to try to bring about theistic belief. The form of Pascal's Wager that concerns us goes as follows. If God exists, the pay-off to the believer is infinite; if God does not exist the loss to the believer and the gain to the non-believer are both finite. Hence as long as there is a finite chance, no matter how small, that God exists the expectation of belief exceeds that of unbelief. Thus it is rational to take steps to bring about belief.

Various objections have been raised against the usage of infinity in this argument.[1] These involve primarily the charge that infinite pay-offs give rise to decisional and mathematical indeterminacies. While these objections have recently been responded to by Jeffrey Jordan,[2] his resolution in turn exhibits some deficiencies. I shall propose a simple alternative approach for dealing with the indeterminacy objections.

The Wager differs from most wagers in that everything must be staked on one bet. In such cases it is questionable whether the expected utility is the proper criterion to use. In the final section I shall discuss an alternative, the most probable average utility, which seems to be more rational.

I

The first objection to the usage of infinite utilities goes as follows. Suppose we can choose two courses of action. Action A makes an infinite reward likely with a probability of 0.5; action B with probability 0.001. The utility of an outcome consists of the benefits that would result if that outcome obtains; the expected utility of an act is determined by multiplying the utility of each possible outcome of the act with its associated probability. In this case the expected utilities for acts A and B are given by:

1. $EU(A) = (0.5)(\infty) + (0.5)(0) = \infty$

2. $EU(B) = (0.001)(\infty) + (.999)(0) = \infty$

It seems that $EU(A) = EU(B)$. Since the expectation is infinite for both A and B, why should we choose A? Here we have what Jordan calls a decisional indeterminacy: it would be rational to choose the act leading to the larger expected utility, but here the two are equal and we can thus make no rational choice on the basis of expected utility.

Jordan accepts the equality of the two expected utilities and concurs that the Wager, in its original form, fails. However, he asserts that the Wager can be saved by augmenting it with an additional decision-theoretic principle. He defends the plausibility of the principle that, if a number of acts all have infinite expected utility, one should perform that one that is most likely to bring about the pay-off.

While this principle does not seem unreasonable, and while it does seem to rescue Pascal's Wager in this particular case, it has a rather limited range of application. As we shall see in the next section, Jordan must construct a different method to solve various mathematical indeterminacies. We shall now present a simple alternative approach that resolves both types of indeterminacies.

Consider how the infinite utilities arise. Pascal writes:

> ...it would be foolish of you when you are forced to gamble not to risk your life in order to win three lives in a game. But in fact there is an eternity of life and of happiness at stake.[3]

The emphasis here is on the fact that the future life is eternal. One must choose between an eternity of happiness on the one hand or an eternity of annihilation or wretchedness on the other hand.[4] Eternity starts at death and goes on forever.[5]

Let us therefore take the infinite utility to arise primarily from the notion of an *unending* happy existence, the essential characteristic being that of unbounded time. Suppose that the infinite reward consists of H units of happiness per day multiplied by an endless number of days. For believers H would be some large positive number corresponding to an appropriate state of happiness; for unbelievers H would be zero (for annihilation) or some large negative number corresponding to an appropriate state of wretchedness. We could then consider the infinite utility to be the limit of HT, where T is a large number of days, as T increases to infinity. Taking the future life as consisting of a finite time T, the expected utilities of acts A and B are:

3. $EU(A) = (0.5)(HT) + (0.5)(0) = 0.5\ HT$

4. $EU(B) = (0.001)(HT) + (0.999)(0) = 0.001\ HT$

The relative utility is EU(A)/EU(B) = (0.5 HT)/(0.001 HT) = 500. Since T always cancels out, this result will hold even as T increases to infinity:

5. $\lim_{T \to \infty} [EU(A)/EU(B)] = \lim_{T \to \infty} [(0.5HT)/(0.001HT)] = 500$

In short, we still merely choose that act that will maximize the relative expected utility. Hence the indeterminacy can readily be removed by taking the limit of a finite quantity as it increases to infinity. No additional decision-theoretic principle need be applied.

Note that this is similar to the question of who is richer: person A who accumulates x dollars per day for an eternity, or person B who accumulates 2x dollars a day for the same eternity? It is clear that at any particular time B will be twice as rich as A. Thus B will always be richer than A, even though both ultimately become infinitely rich. In the case that both options are equally likely, the limit approach to the expectations will prefer B. Jordan's principle, however, breaks down here since both options have an infinite utility and have equal probability, leading to a decision indeterminacy.

If it be insisted that the infinite reward is actually a double infinity of an eternity of infinite happiness, this could be handled in a similar fashion by taking also the limit of a boundless H as it approaches infinity.

One might object that the infinity thus approached is only a *potential* infinity, not an *actual completed* infinity. To this we respond that, first of all, it is not clear that Pascal's Wager requires an actual, completed infinity. All that is necessary is that this timespan is greater than any finite number we can think of. Certainly, from a human point of view, the future eternity is only an indefinite, potential infinite that will never be completed. Moreover, it is evident that if, as in the above example, a quantity has the same value at any future time then this value will be valid for all eternity, regardless as to whether we consider this eternity to be an actual completed infinity or merely a potential infinity.

The limit approach has the advantage of being compatible with standard Bayesian principles, except for the fact that many standard formulations of Bayesian theory reject unbounded utilities. Since unbounded utilities are realistic (at least for those who believe, along with Pascal, in a future eternity), any incompatibility of such utilities with Bayesian principles must be considered a deficiency in Bayesian theory.

II

A second indeterminacy arises as follows. Suppose there is a very small probability that there exists a devil who punishes theistic believers with infinite hell. Jordan considers the case where a certain act (belief in a theistic god) has a (0.45-x) probability of bringing about an infinite outcome (heaven), a 0.55 probability of no afterlife and a very remote probability of x resulting in an infinite disutility (hell). The EU of this act would be:

6. $EU = (0.45\text{-}x)(\infty) + (0.55)r + (x)(-\infty) = \infty + -\infty = ?$

where r is a finite utility and x is a very small number. In this case we seem to have a mathematical indeterminacy. How can we decide what value the EU should have?

Jordan accepts the view that the mathematical indeterminacy is real. His resolution is to remove the infinite disutility ($-\infty$) from the calculation. He justifies this on the ground that this possibility is so remote that it warrants nothing but neglect.

One might object that any non-zero probability, however small, multiplied by infinity still yields an infinite utility. To this Jordan responds that *every* act carries with it possible outcomes that involve infinite utilities, since every act might be punished with an infinite disutility by some bizarre god. Hence every act carries with it the above sort of indeterminacy. Thus, Jordan argues, just as we properly neglect very remote decisions in mundane decisions, we are justified in doing so in Pascalian decisions also.

He concedes, however, that if x is as large as 0.01 then the indeterminacy is not thus removable. Jordan then goes on to argue that the latter case is, from Pascal's point of view, not realistic: Pascal directed his Wager against people for whom the only real outcomes were either theism or naturalism. They would not have attached much value to x. Thus, according to Jordan, we are justified in ignoring the indeterminacy.

Frankly, I don't find this resolution very convincing. First of all, the removal of the infinity disutility strikes me as directly contrary to the spirit of Pascal's Wager. Pascal argued that as long as there was *any* finite chance that God exits — no matter how small — then infinite reward will favour the EU for belief. In the words of Pascal:

> For in this game you can win eternal life, eternal happiness. You have one chance of winning against a finite number of chances of losing, and what you are staking is almost nothing. Surely that settles it. Wherever there is infinity, and where there is not an infinity of chances of losing against the chance of winning, why hesitate? Surely you must stake everything then.[6]

Pascal does not allow for some finite threshold probability that allows us to ignore remote but non-zero possibilities. Such a manoeuvre would permit naturalists to defuse the Wager by assigning sufficiently low probability to God's existence and then removing the infinite utility. Also, Jordan leaves Pascal's Wager as indeterminate, thereby invalidating it, for those who estimate probability x to be small but still non-negligible (e.g., in the order of 0.01).

These difficulties arise because Jordan again concedes too much strength to the objection. An easier resolution is to use the same approach as above. Again, take T as the number of days in the life here-after and take the limit as T approaches infinity. This yields:

7. $\lim_{T\to\infty}$ [(0.45-x)HT +0.55 r +(x)(-HT)] = $\lim_{T\to\infty}$ [(0.45-2x)HT +0.55 r] = $+\infty$

for x < 0.225. Thus, as long as the probability of God's existence outweighs that of the disutility, the EU is positive infinity.

This is similar to the question of what our assets would be if we acquire 2x dollars a day for an eternity but spend x dollars per day. It is clear that, however far we look into the future, our income will always exceed our expenses and our net assets will increase to $+\infty$, rather than ending up as an indeterminate [$\infty + -\infty$].

In summary, the infinite utilities in Pascal's wager need not involve any decisional or mathematical indeterminacies. Nor do they require any augmentation of the standard version of the Wager with certain decision-theoretic principles. Of course, this conclusion does not demonstrate the validity of Pascal's Wager; it only asserts that the usage of infinite utilities need not, in itself, be problematic.

III

While I have thus far defended Pascal's Wager, I am nevertheless convinced that it is flawed. What concerns me is the fact that the Wager, unlike most gambling games, asks us to stake everything on one bet. The obvious significance of an infinite utility for an outcome is that, even though that outcome may be very improbable, it still generates an infinite expected utility. Yet I wonder whether in such cases the EU is indeed the pertinent quantity that should determine our rational choice.

Consider a simple example. Suppose we have an opportunity to stake our entire fortune (say $100,000) on a 0.001 chance of winning a billion dollars. The expectation of betting is

8. EU(bet) = (.999)(0) + (0.001)(1,000,000,000)= 1,000,000

If we refrain from betting, thereby keeping our $100,000, the expectation is EU(no bet) = (1)(100,000) = 100,000. Thus, by Pascal's logic we would be foolish not to bet. Yet if we do bet we stand a very good chance (99.9%) of losing everything. Out of every thousand persons who take on the bet all but one ends up bankrupt. Is it then really rational to risk the family fortune on such a long-shot, even if the possible reward may be very great?

The difficulty is that the EU will yield us the average expected winnings only *in the long run*, if we bet many times. If we can place only a few bets than it seems more rational to consider the *most probable average utility* (hereafter referred to as MU): the most probable distribution of outcomes, multiplied by their respective utilities and then averaged. If we bet n times then the most probable distribution of outcomes with probabilities (p_1, p_2,...) will be (np_1, np_2,...), rounded off to the nearest integer. Thus, in the above

case, if we can place 10000 bets the most probable distribution for prob-
abilities (.999, 0.001) is (9990, 10). Then the MU is given by:

   9.   MU(bet) = [(9990)(0) + (10)(1,000,000,000)]/10000 = 1,000,000

In this case it would certainly be rational to bet. Note that here the MU equals
the EU. This will always be the case if n, the number of bets, is sufficiently
large. For small n, however, this will not hold true, particularly not for n =
1. In that case the probability distribution (0.999, 0.001) leads to a most
probable outcome of (1, 0). This yields an MU:

   10. MU(bet) = [(1)(0) + (0)(1,000,000,000)]/1 = 0
   This compares with a MU(no bet) of 100,000.
   I maintain that in such situations as these, where everything is risked on a
small number of bets, it is MU, rather than EU, that should be used when
comparing the values of different acts. Let me stress that I am concerned here
only with *high-risk* situations: there may well be low-risk cases where one
might be rationally justified in a single-case bet on a low-probability out-
come. Of course, when the relative probabilities are of the same order (e.g.,
a 49% chance of winning) the matter becomes less clear.
   To illustrate this further, consider the St. Petersburg Paradox. This "para-
dox" is often raised in discussions of Pascal's Wager; Jordan also brings it
up. In the St. Petersburg game one gets a pay-off of 2^m if the first "heads"
appears on the mth toss. The probability of the first "heads" appearing on the
nth toss is 2^{-m}. The EU is therefore:

   11.   EU = (1/2)(2) + (1/4)(4) + (1/8)(8)...= 1 + 1 + 1+... = ∞

The paradox is that, even though the EU is infinite, in practice no one would
pay much to play the game. How can it be resolved? Jordan suggests that we
should simply ignore the higher m's since their probabilities are very small.
This seems to me to be not only unwarranted, but also contrary to the spirit
of the Wager, where even small probabilities can lead to a significant EU. I
suggest that, here, too, it is more instructive to consider the MU, rather than
the EU. The probability distribution is (1/2, 1/4, 1/8, ...). Thus for n = 1 we
have a most probable outcome distribution of (1/2, 1/4, 1/8,...) which rounds
off to (1, 0, 0,...), yielding:

   12.   MU(1) = [(1)(2)]/1 = 2

For n = 3 we get a most probable outcome distribution of (2, 1, 0,...) yielding

   13.   MU(3) = [(2)(2) + (1)(4)]/3 = 8/3

For n = 7 the most probable outcome distribution is (4, 2, 1, 0,...), with

   14.   MU(7) = [(4)(2) + (2)(4) + (1)(8)]/7 = 24/7

In general, it turns out that

    15. $MU(n) = [(n + 1)(\log_2 (n+1))]/n$

Thus the value of the game as determined by the MU depends on how often you play, increasing to infinity only as n approaches infinity. This resolves the paradox: we can expect an infinite average payoff only if we play infinitely often; for a finite number of plays we can expect only a finite average payoff. Once again, the MU seems more pertinent than the EU as a criterion for rational choice.

What about Pascal's Wager? On the basis of expected utility, Pascal argues that it is rational to risk one life to gain an infinity of lives, as long as the chance of God's existence is non-zero. Yet, if this chance is small, say 0.01, the most probable distribution of outcomes is (.99, 0.01), which rounds off to (1, 0) and we have:

    16. $MU(belief) = (1)(0) + (0)(\infty) = 0$

    17. $MU(nonbelief) = (1)(1) + (0)(0) = 1$ life

Thus, using the MU criterion, the nonbeliever may be pardoned for hesitating to risk his entire life for a one percent chance that God exists. From this point of view the Wager works only if the probability of God's existence can be shown to be considerably higher. To the extent that this requires additional lines of argumentation, it significantly weakens the Wager.

Trinity Western University

NOTES

1. See, for example, Richard Jeffrey, *The Logic of Decision* (Chicago: University of Chicago Press, 1983), pp. 153-54; Antony Duff, "Pascal's Wager and Infinite Utilities" *Analysis,* vol. 46 (1986), pp. 107-9.

2. Jeffrey Jordan, "Pascal's Wager and the Problem of Infinite Utilities" *Faith and Philosophy,* vol. 10 (1993), pp. 49-59.

3. Blaise Pascal, *The Mind on Fire: An Anthology of the Writings of Blaise Pascal*, ed. James M. Houston (Portland: Multnomah Press, 1989), p. 132.

4. *Ibid.*

5. Pascal, *op. cit.,* p. 126.

6. Pascal, *op. cit.,* p. 151.