Phil. 244: Epistemic Normativity
Feb. 28, 2011

Pascal's Wager

I. Pascal's Argument

Today I will defend Pascal's wager against its two most prominent objections. More specifically, I will argue:

> If we accept the decision-theoretic framework within which Pascal couches his argument, and if we're not scared of infinities, then there is a sound argument for a wager-like conclusion.

Hájek's way of formulating Pascal's argument:

P1. Either God exists or God does not exist, and you can either wager for God or wager against God, and the utilities (for you) of the relevant possible outcomes are as follows:

	God exists	God doesn't exist
wager for God	∞	u_2
wager against God	u_1	u_3

where u_1, u_2, and u_3 are finite real numbers. [*premise*]

P2. Rationality requires that the probability, p, that you assign to God existing be positive and not infinitesimal. [*premise*]

P3. Rationality requires you to perform the act that maximizes your expected utility (if there is one). [*premise*]

C1. Wagering for God is the act that maximizes your expected utility. [*follows from P1 and P2*]

C2. Rationality requires you to wager for God. [*follows from C1 and P3*]

Calculation of the relevant expected utilities:

E.U.(wager for God) $= p \bullet \infty + (1-p) \bullet u_2 = \infty$

E.U.(wager against God) $= p \bullet u_1 + (1-p) \bullet u_3 =$ something finite

(I assume, for simplicity, that the relevant states of the world are independent of what the agent does.)

Some minor worries we can set aside:

- *worry #1:* There should be more infinities in the decision matrix; in particular, the utility of wagering against God if He exists should be $-\infty$.

  reply: Whether we take the disutility of damnation to be finite or infinite doesn't affect any of the central philosophical issues raised by Pascal's wager. So to simplify things, let's assume it's finite.

- *worry #2*: It's rationally permissible to assign a probability of 0 to God's existence.

  reply: Would you really be willing to enter a bet in which you gain \$1 if God doesn't exist and lose \$10,000,000 if He does?

- *worry #3*: Since the utilities here are utilities *for the agent* (i.e. how much *he* or *she* values each outcome), P3 at best formulates a principle of *prudential* rationality, not a principle of *all-things-considered* rationality.

  reply: Fine, then replace every occurrence of "rationality requires you to..." with "rationality requires you, insofar as your own well-being is only what's at issue, to...."

II. First Objection: Theoretical vs. Practical Rationality

Suppose we interpret "wager for God" to mean "believe in God." Then Pascal's argument is open to the following objection:

> "To offer this as an argument for the conclusion that rationality requires one to believe in God is to conflate *practical rationality* with *theoretical rationality*. The decision-theoretic principle embodied in P3 is a principle of practical rationality: it concerns what it is rational *to do*. However, belief is not action, so no principle of practical rationality can establish a conclusion about what it is rational *to believe*. Thus C2 does not follow from P1, P2, and P3."

Two not-entirely-satisfactory responses:

1. Reformulate the argument purely as an argument about theoretical rationality, and then insist that there exist *pragmatic reasons for belief*, and the balance of these reasons is what makes belief in God rational.

   problem: As we saw last meeting, the existence of pragmatic reasons for belief is controversial.

2. Instead interpret "wager for God" to mean "try one's best to get oneself to believe in God." What would that involve? Pascal's answer: hang around true believers, mimic their practices, go to mass, take holy water, etc.

   problem: Upon dying, one will only get the eternal bounty of salvation if one successfully manages to inculcate in oneself a genuine belief in God. But one might doubt that associating with believers and aping their practices is a reliable means of fostering such a belief.

A better response:

> We can avoid the objection by switching the issue from being a question about *one's own beliefs* ("Should *I* believe in God, or try to get *myself* to believe in Him?") to being a question about *other people's beliefs* ("Should I try to get *the people whom I care about* to believe in God?").

In particular, let us assume that one has children, and that one cares about their well-being. Then we can argue as follows:

P1′. Either God exists or God does not exist, and you can either raise your children to believe in God or not raise your children to believe in God, and the utilities (for your children) of the relevant possible outcomes are as follows:

	God exists	God doesn't exist
you raise your children to believe in God	∞	u_2
you don't raise your children to believe in God	u_1	u_3

where u_1, u_2, and u_3 are finite real numbers. [*premise*]

P2′. Rationality requires that the probability, p, that you assign to God existing be positive and not infinitesimal. [*premise*]

P3′. Rationality requires you, insofar as you care about your children, to perform the act that maximizes their expected utility (if there is one). [*premise*]

C1′. Raising your children to believe in God is the act that maximizes their expected utility. [*follows from P1′ and P2′*]

C2′. Rationality requires you, insofar as you care about your children, to raise them to believe in God. [*follows from C1′ and P3′*]

What's at issue here is whether to engage in a certain course of action, so there are no longer any worries about practical rationality being conflated with theoretical rationality. Moreover, the conclusion of this argument seems just as disturbing (at least for non-believers!) as the conclusion of Pascal's original argument.

Finally, although it is not absolutely certain that one will successfully inculcate a belief in God in one's children if one raises them as believers, it seems undeniable that raising them to be believers greatly increases the likelihood that they end up believers.

> [Actually, there's a complication here. Since there's a chance your children will be believers even if you raise them to be non-believers, and a chance they'll be non-believers even if you raise them to be believers, the decision matrix in P1′ is not accurate, and we actually need to formulate the argument using either evidential or causal decision theory.]

III. Second Objection: The Many-Gods Objection

Probably the most widely endorsed objection to Pascal's argument is the so-called Many-Gods Objection:

> The decision matrix in P1 assumes there are only two alternatives that are relevant to one's decision: either (the Christian) God exists, or (the Christian) God does not exist. However, surely this does not exhaust the relevant theological options.
>
> For example, God might end up being *self-loathing* and bestow eternal bliss on all and only those who *don't* believe in Him. Then if we assume that the only way you might acquire infinite utility is if either the self-loathing or Christian God bestows it upon you, the true decision matrix will be as follows:
>
> | | Christian God exists | self-loathing God exists | neither God exists |
> |---|---|---|---|
> | wager for God | ∞ | u_2 | u_3 |
> | wager against God | u_1 | ∞ | u_4 |
>
> If p_1 is the probability one assigns to the Christian God existing, and p_2 is the probability one assigns to a self-loathing God existing, the relevant expected utilities are now:
>
> $$\text{E.U.(wager for God)} = p_1 \bullet \infty + p_2 \bullet u_2 + (1 - p_1 - p_2) \bullet u_3 = \infty$$
>
> $$\text{E.U.(wager against God)} = p_1 \bullet u_1 + p_2 \bullet \infty + (1 - p_1 - p_2) \bullet u_4 = \infty$$
>
> Either way one's expected utility is infinite, so decision-theoretic considerations have nothing to say about whether one should or should not wager for God: rationality does not settle the matter.

However, this is too quick. Consider the following example (which I owe to Peter Bach-y-Rita):

> *Die vs. Coin*: I give you two options: either (1) you can flip a fair coin, and if it comes up heads you will receive an eternity of bliss (otherwise you receive nothing); or (2) you can roll a fair 6-sided die, and if a "1" rolls you will receive an eternity of bliss (otherwise you receive nothing).

Using reasoning similar to that used in the above version of the Many-Gods Objection, the expected utilities for you of each option would seem to be:

$$\text{E.U.(flip coin)} = 0.5 \bullet \infty + 0.5 \bullet 0 = \infty$$

$$\text{E.U.(roll die)} = 0.1666... \bullet \infty + 0.8333... \bullet 0 = \infty$$

So, by parity of reasoning, rationality doesn't settle the matter, as both options have the same expected utility. But surely this is the incorrect result! It would be the height of irrationality to choose the die over the coin.

> [Note that the same example shows that something must be wrong with the Duff/Hájek "mixed strategies" objection to Pascal's argument.]

Even if standard decision theory is not equipped to handle infinite utilities, I take it to be a datum that, whatever the correct infinitary analogue of decision theory is, it *must* yield the result that, in Die vs. Coin, choosing to flip the coin is more rational than choosing to roll the die.

Similarly, in the following case:

Dueling Lotteries: I give you three options: either (1) you can enter, for free, a lottery in which you have a p_1 chance of winning an eternity of bliss (and otherwise receive nothing); or (2) you can enter, for free, a lottery in which you have a p_2 chance of winning an eternity of bliss (and otherwise receive nothing); or (3) you can enter neither lottery,

I take it to be a datum that, if $p_1 > p_2$, the rational course of action to take is option (1). But it is only a short step from here to concluding that if (i) the probability of the Christian God existing is p_1, (ii) the probability of a self-loathing God existing is p_2, and (iii) $p_1 > p_2$, then one should wager for God.

Two complications:

a. Maybe infinite expected utilities do not always trump finite expected utilities. Consider:

> *Risky Lottery*: I give you the option of entering, for free, a lottery in which you have a 1 in a googolplex ($10^{(10^{100})}$) chance of winning eternal bliss, but if you lose you must suffer 10 years of excruciatingly painful torture.

Some might resist the conclusion that rationality requires you to enter this lottery.

> [Advocates of standard decision theory, however, must accept that conclusion. For they are committed to the rationality of entering a variant of this lottery in which what one wins is a finite but insanely large (say: a googolplex raised to the googolplex) number of days of eternal bliss. And surely they are committed to the claim that it is more rational to enter the infinite version of Risky Lottery than the finite version of Risky Lottery.]

b. Not all infinities are created equal. Consider:

> *Dueling Lotteries (Revised)*: I give you three options: either (1) you can enter, for free, a lottery in which you have a p chance of winning an infinite life of daily *infinite* happiness (and otherwise receive nothing); or (2) you can enter, for free, a lottery in which you have a $p + \varepsilon$ chance of winning an infinite life of daily *finite* happiness (and otherwise receive nothing); or (3) you can enter neither lottery.

When ε is small, most of us are inclined to say that the rational course of action is option (1).

Thus if the pay-off for, say, believing in Allah when He exists is a different order of infinity than the pay-off for believing in the Christian God when He exists, the correct infinitary analogue of decision theory might yield the result that one should wager for Allah, even if one thinks it more likely that the Christian God exists.

IV. An Unsettling Conclusion

When one combines the replies to the first and second objections to Pascal's Wager, we reach the following conclusion. Given the general decision theoretic framework assumed by Pascal, and ignoring for ease of exposition complication (b) above, we have an apparently sound argument for the following conclusion:

> Rationality requires you, insofar as you care about your children, to raise them to do whatever is required for salvation by the salvation-bestowing God [or combination of salvation-bestowing Gods] you think has the greatest chance of existing.

So far as I see it, this leaves us three options:

1. Reject the decision-theoretic framework within which Pascal couches his argument.
2. Go strict finitist and reject the very notion of infinity. (A big bullet to bite, given the ubiquity of appeals to infinity in mathematics, set theory, physics, and even decision theory itself.)
3. Accept the conclusion.

(One way in which the conclusion might not end up being so bad: perhaps the salvation-bestowing God most likely to exist only requires that one be a good/virtuous person in order to achieve salvation. But, of course, to say this is to engage in some highly speculative theology.)