

Pascal's Wager Is A Possible Bet (But Not A Very Good One): Reply To Harmon Holcomb

III

In "To Bet The Impossible Bet"ⁱ, Harmon Holcomb III argues that "[W]e should not accept [the terms of Pascal's Wager] since they violate the preconditions of it making sense to bet. The wager is useless because of a structural breakdown in the conditions that determine the relations between the act of betting and the payoff matrix options. ... A bet is genuine only if there is a chance of success in reaping the bet's rewards. Any chance vanishes if the wagerer's act of betting something (what is bet) on something (a betting option) puts the same thing in both roles: we cannot be forced to bet by believing or not believing on the options of believing or not believing. Neither an irresistible temptation nor a bribe, the wager is an inconceivable bet. No question of its prudence arises, any more than that of how best to square the circle."ⁱⁱ

I disagree. If there is a structural breakdown in the formulation of Pascal's Wager, it comes with the introduction of infinite utilities; otherwise, the wager is a straightforward application of decision theory in the calculation of the expected utilities of alternate (courses of) actions. Moreover, I disagree with Holcomb's further claim that "If it is a possible bet, the wager argument works; if it is not a possible bet, it doesn't work". Subject to the possible qualification involving infinite utilities, Pascal's Wager is "a possible bet"; but there are numerous reasons why it is not a very good one. In particular: (i) as it stands, the Wager gives infinite utility to every course of action which is open to those who are giving consideration to the wager; (ii) there are many alternative formulations of the payoff matrix in Pascal's Wager, and not all of these formulations give the same result; and (iii) there are many alternative wagers (most involving alternatives to Pascal's God), and consideration of these wagers undermines the apparent cogency of Pascal's argument. Holcomb claims that these criticisms rest on a mistaken logicist methodology; however, I shall argue that this suggestion rests on a misunderstanding of the nature of the distinction between theoretical and practical reasoning. I begin with this last point.

According to Holcomb, there are two broadly different methodological approaches which might be taken in evaluating Pascal's wager. On the one hand, there is a *logicist methodology* which "evaluates the target argument solely in terms of the norms of logic, letting the chips fall where they may on the implications for the rationality of human practices".ⁱⁱⁱ And on the other hand, there is a *pragmatist methodology* which "evaluates the target argument in terms of pragmatic norms, letting the chips fall where they may on the implications for the adequacy of logical norms to reconstruct the sense of human practices".^{iv} Holcomb claims that the pragmatist methodology provides a "ready response" to logicist critiques of Pascal's wager: "By divorcing logic from practice, logicist evaluations ignore the fact that humans are committed to choosing between the alternatives they envisage. Human rationality concerns the relative merits of a belief-choice among given options, not the relation of belief to absolute standards. A decision is a practical decision only if (i) it is a decision to be made by particular people in particular circumstances capable of being used by those people in those circumstances, to be made justifiable or reasonable as products of socially, historically situated people rather than as sentences abstracted from human commitments and communicative acts; and (ii) it is a decision whose character contrasts with what is (a) theoretical; (b) speculative; (c) ideal; (d) unrealistic; and (e) imaginative." In particular, Holcomb claims that my allegedly logicist critique of Rescher's defence of Pascal's Wager^v "begs the question" against the second half of this characterisation of practical decision: "In Oppy's case, the wager is measured against the following: (1) purely theoretical possible assignments of numerical probabilities to God's existence of finitely versus infinitely miniscule probabilities; (2) merely speculative constructions of irrelevant scenarios of alternative gods to the traditional Christian God that are not even taken seriously or given credence by believers or non-believers; (3) epistemically ideal Cartesian expectations that our whole theological belief and

background system can be defended; (4) blatantly unrealistic logical demands that each theological claim can be reconstructed as the conclusion of a sound argument such that its inference contains no gaps and none of its premises lack their own sound arguments; and (5) unceasing imaginative doubts that block any personal involvement in deciding whether to believe, even though believing or not believing in God is a wager we are forced to make.”^{vi}

I disagree with nearly all of this. The root problem lies in Holcomb’s misunderstanding of Pascal’s premise that “reason can settle nothing here”. Holcomb seems to think that Pascal here supposes that belief in God is not “rationally defensible” -- and hence he is led to criticise the pragmatist approach on the grounds that “we have no firm conception of the sort of rationality left after reason departs”^{vii}. But surely what Pascal meant, or at least ought to have meant, is that merely *theoretical* reason can settle nothing here; thus leaving open the question whether *practical* reason can provide a good argument in support of traditional theism. Very roughly, we might characterise the distinction between kinds of reasons which is being adverted to here as follows: Theoretical reason operates solely in the domain of belief; it is guided solely by considerations of truth and evidence, or, more generally, by considerations concerned with the theoretical virtues of beliefs (simplicity, explanatory power, coherence, etc). Practical reason operates in the domain of belief, desire, and action; it is guided by considerations concerning the consequences of courses of action (satisfaction of desires, upholding of values, etc.). Of course, there is much more to be said about these two kinds of reasons, and about their interplay -- but even this rough and ready characterisation shows that it is unnecessary to suppose that we need to seek for some other “sort of rationality left after reason departs”.

The standard modern interpretation of Pascal’s wager -- from which I see no reason to depart, and with which Holcomb provides no reason to disagree -- takes it to be an exercise in decision theory. Some people hold that decision theory provides a complete theory of practical reason; others disagree. We don’t need to take a stand on this issue here: all that matters is that we should accept the relatively uncontroversial claim that decision theory models at least part of the operations of practical reason. Given this much, we can take Pascal’s wager to be an argument

within the sphere of practical reason for the conclusion that the expected utility of belief in God is greater than the expected utility of non-belief. Moreover, we can see that Holcomb's distinction between logicist and pragmatist methodologies is indeed a red herring (though not for the reason that he supposes): there is no proposal to evaluate the argument solely in terms of the norms of theoretical logic ("divorced from practice"); rather, the proposal is to evaluate the argument in terms of the norms of practical reason, i.e. in terms of the norms which actually govern our decisions amongst possible courses of action. Decision theory just is (a part of) the theory of "decisions to be made by particular people in particular circumstances capable of being used by those people in those circumstances, to be made justifiable or reasonable as products of socially, historically situated people rather than as sentences abstracted from human commitments and communicative acts".

Of course, there is at least one important question about the connection between theoretical reason and practical reason which remains to be addressed here. Decision theory is a theory about the rational combination of beliefs and desires in the production of actions -- i.e. the data for decision problems includes both an assessment component (desires, values, evaluation of outcomes) and a theoretical component (beliefs, probabilities of outcomes). Consequently, decision theory does not operate in entire independence from theoretical reason: one of the input factors in any decision problem, viz. the beliefs about probabilities of outcomes, belongs to the province of theoretical reason. *Prima facie*, this fact may seem to be in conflict with Pascal's claim that "reason can establish nothing when it comes to the question of the existence of God". But, of course, there need be no genuine conflict here: if Pascal is claiming that theoretical reason cannot give a decisive answer to the question whether God exists, then he need not be taken to be denying that theoretical reason can make some lesser contribution to the assessment of the claim that God exists. And it does seem, at least *prima facie*, that some such lesser contribution is required.

Suppose that one took Pascal to be claiming that theoretical reason can make NO contribution to the assessment of the claim that God exists. What probability should one then assign to the claim that God exists? Since theoretical reason can make NO contribution, it seems that there is no

probability which it is reasonable to assign -- i.e. one should think that no one value in the interval $[0, 1]$ is better than all the others. But in that case, one doesn't have a value to feed into the decision-theoretic calculation; and, moreover, there are values in the interval which lead to different results when one does the calculations. (Could one support Pascal's conclusion with the observation that almost all of the values in the interval will give the result that one ought to believe in God? Even if one ignores the possibility that there might be infinitesimals in the interval, it seems not. After all, if theoretical reason makes NO contribution, then it doesn't tell you that any value is just as good as the next; rather, it is simply silent on this question as well.)

Of course, there is an assumption underlying this line of argument, viz. that it is theoretical reason which determines (or should determine) the probability which is assigned to the claim that God exists. But perhaps this assumption can be disputed. Perhaps practical reason can play a part in the determination of this probability; or perhaps there can just be a brute matter of fact that a certain probability is assigned, a fact which is untouched by the further workings of theoretical reason. (Moreover, it might be this point which is the real target of Holcomb's distinction between logicist and pragmatist methodologies.)

Put like this, these suggestions don't seem terribly attractive. If theoretical reason says that there is no theoretical reason to choose one probability value over another, then it is simply in conflict with the brute assignment of a value (no matter what it happens to be). And if theoretical reason says that there is no theoretical reason to believe in God, then surely it is just a case of wishful thinking to believe on the basis of something other than theoretical reasons (e.g. on the basis of desires). However, it seems to me that the availability of these responses shows that we still haven't managed to make the correct construal of Pascal's original suggestion. Pascal did not think that theoretical reason says nothing to theists about the question of God's existence; indeed, he thought that several of the traditional theistic proofs were both sound and convincing. Pascal did not think that theists needed to appeal to desires or brute probability assignments if called upon to give reasons for their belief in God. Rather, what Pascal thought was that there is no purely theoretical argument which should persuade agnostics and atheists -- i.e. those who are not

already theists -- to adopt the belief that God exists. What he sought was an *argument* which would persuade agnostics and atheists to change their minds; and he thought that his decision-theoretic argument could fit the bill. Of course, he didn't think that theoretical considerations could play no role in the calculation -- after all, any target of the argument must give some prior probability to the belief that God exists -- but he did think that theoretical considerations alone would clearly be insufficient.

If Pascal's wager is understood this way -- i.e. as an argument intended to persuade atheists and agnostics to mend their ways -- then it is clearly important to ask about the probability judgements and beliefs which atheists and agnostics should, or are likely to, make and have. I suggest that reasonable atheists and agnostics will subscribe to at least one of the following claims: (i) there is a no more than infinitesimal probability that God exists (such that the product of this infinitesimal with the value to be accrued by belief if God exists is not maximal in the calculation of expected utility); or (ii) there are other formulations of the pay-off matrix which deserve (equal) considerations and which lead to different results (e.g. that God rewards everyone, regardless of what they believe); or (iii) there are other wagers, involving alternative deities, which are no less deserving of attention, but which give contradictory advice about what to believe (e.g. there is the God who rewards all and only those whose favourite number is 17).

It is these suggestions which Holcomb labels "purely theoretical", "merely speculative", "epistemically ideal", "blatantly unrealistic", "unceasing imaginative doubts", etc. But in my view these labels are completely misplaced. What is at issue is what atheists and agnostics believe, and the reasonableness of their so-doing. It seems to me eminently likely that atheists and agnostics think that there are numerous alternatives to the traditional Christian God which should be taken no less seriously than it, or, at least, not a whole lot less seriously than it. But, of course, this is not to say that they think that these alternatives should be taken seriously, or given credence, by either believers or non-believers. Moreover, it isn't mere "unceasing imaginative doubts" or "purely theoretical speculations" which are the source of atheism and agnosticism: rather, that atheism or agnosticism emerges with a view of the world which finds many straightforward reasons for

doubting that the Christian God exists (e.g. in the lack of good arguments or positive evidence, in the problem of evil, in what Hume calls “the natural history of religion”, in the adequacy of a purely naturalistic conception of the world, and so on). And, moreover, there are no “Cartesian expectations that our whole theological belief and background belief systems can be defended” or “unrealistic logical demands that each theological claim can be reconstructed as the conclusion of a sound argument” in any of this: the point isn’t to establish that theists must give up what they believe; rather, it is to show why non-theists won’t be persuaded to share those beliefs.

In sum: there are various reasons why Holcomb’s distinction between logicist and pragmatist methodologies is indeed a red herring. First, it misconstrues the distinction between theoretical and practical reason. And second, it supposes that standard atheistic and agnostic responses to Pascal’s Wager are intended to undermine the justification for the beliefs of theists when, in fact, those responses are simply meant to show why atheists and agnostics do not find the argument persuasive. There is some irony in the fact that many atheists and agnostics would count themselves heirs of the American pragmatists. There is no conflict between the use of the objections which I want to press against Pascal’s wager and the endorsement of a broadly coherentist, non-foundationalist epistemology.

II

I turn next to the question of the merits of Pascal’s wager, given the concession that it is “a possible bet”. Holcomb claims: “Logicist and pragmatist methods share a common counterfactual: If we were to grant the truth of the theologies and background beliefs taken as the basis for making the wager, e.g. as a positive, significant probability that God exists and the elimination of other Gods as relevant alternatives, our result would be that the wager is a good one.”^{viii} “... There may indeed be an insight behind the wager, namely, that a finite being should sacrifice everything of finite value in order to attain the infinite value of oneness with God if doing so is possible. ... If

it is a possible bet, the wager argument works; if it is not a possible bet, it doesn't."^{ix}

In order to evaluate these claims, we shall need to have the wager argument before us. Very briefly, we may encapsulate it as follows: Let the value of your life, given that you don't believe in God and that God does not exist, be zero. (This is an arbitrary calibration, intended to make the calculations go more easily). Let the cost of belief in God be $-B$. And let your subjective probability that God exists be p . You are to decide whether or not to believe in God; your decision is guided by the following table, which enables you to calculate the expected values of belief and non-belief under the conditions that God exists and that he does not exist:

	God exists (Prob= p)	God does not exist (Prob= $1-p$)
Believe in God	infinity	$-B$
Don't believe in God	little or nothing (perhaps even negative)	0

Clearly, the expected value of belief is: $p \cdot \text{infinity} - B = \text{infinity}$; and the expected value of non-belief is no more than $p \cdot (\text{little or nothing}) + 0 = \text{little or nothing}$. So, by the principle which enjoins one to maximise one's expected utility, one should believe in God.

Several points of clarification should be added to this sketch. *First*, it is unclear when one is to believe in God if one is to get the reward. Perhaps all that is required is belief at the moment of death; or perhaps there should be belief over a longer period. If all that is required is belief at the moment of death, then, ignoring practical difficulties, it seems that one should be able to have this without cost: just wait until one is dying, and then form the belief. Of course, this suggestion is blatantly unrealistic -- how can one be sure that one will know that one is dying? how can one be sure that the belief which one will form will be sufficiently sincere, etc. -- but these worries can all be subsumed under our *second* point of clarification, viz. that it is an unrealistic assumption to

suppose that one can just choose to believe in God. The actual wager argument tells you that you want to be a believer; but it doesn't tell you how to achieve this end, and so it doesn't actually tell you what to do. However, it does tell you something: it tells you that you should enter into courses of action which have some chance of leading to the result that you end up believing. Pascal suggests that you should hang out with believers, go to mass, etc. Perhaps there are other strategies which will work for other people. *Finally*, it is sometimes alleged that if one came to believe in God on the basis of the wager argument (*via* some complex chain of intermediate events), God would not reward the belief because of the venal motives upon which it is based. This objection also seems misconceived. For, on the assumption that God would not reward a belief based on venal motives, the idea will be that the venal motives are merely supposed to start one down a path which leads eventually to a state which is genuinely deserving of divine reward; i.e. the venal motives will be supposed to vanish from the scene at some point as one travels down the path.

Even given these points of clarification, there are various objections which can be brought against the argument. I shall begin with two objections which "grant the truth of the theologies and background beliefs taken as the basis for making the wager" -- i.e. objections which (at least for the most part) do not contest the claim that there is a finite, non-zero probability that God exists, nor the claim that there are no other relevant alternative Gods which need to be taken into account in a fuller calculation.

First, as Anthony Duff^x points out, the Wager argument in itself provides no guidance at all to present action. No matter what you do, there is some chance that it will be the beginning of a chain of actions which leads eventually to your believing in God and obtaining the infinite reward. Consequently, no matter what you do, your action has infinite expected utility -- no matter what you do, it has a non-zero chance of leading you to heaven. Perhaps it might be replied that one ought to engage on a course of action which gives the greatest chance of obtaining the reward -- but what would be the justification for this reply? It is no part of standard decision theory: standard decision theory merely tells one to maximise expected utility -- and, as we have seen, it

seems by Pascalian lights that every action has the same expected utility. Perhaps it might be said that the case involving infinite utility is special: there is a tie-breaking rule which applies in just this case. Maybe; but one could reply that one does better to repudiate the extension of decision theory to cases involving Pascalian infinite utilities. There are independent reasons -- drawn from consideration of cases like the two envelope paradox and the St. Petersburg game -- for excluding distributions of values which lack a finite mean from decision theoretic calculations. In particular, these reasons suggest that one ought to exclude individual infinite values from decision theoretic calculations. Rather than extend decision theory with *ad hoc* rules aimed at accomodating Pascalian reasoning, an opponent of the wager might well think that the case is one to which decision-theoretic reasoning is known to be inapplicable, on pain of paradox. (I shall consider possible replies to this argument later).

Second, one might doubt the values which appear in the table. Surely there is some chance that God will reward everyone, regardless of whether they believe or not. If so, then the expected value of disbelief will also be infinite -- and so the Pascalian argument will give one no reason to start down a path calculated to lead to belief. Furthermore, picking up on the suggestion made in the last paragraph, one might think that all the values in the table must be finite -- there can be no such thing as *infinite* utility, strictly construed. Of course, this suggestion doesn't entail that the wager argument won't work -- indeed, it might well be construed as a way of developing an argument which has far more chance of succeeding^{xi} -- but it does entail that the success of the argument will depend upon the precise magnitudes which are assigned to the probabilities and values. (One might conjecture that reasonable atheists and agnostics are *bound* to assign these magnitudes in such a way that the calculation of expected utility favours disbelief.) Alternatively, pursuing the opposite tack, one might think that if there can be infinite values, then there can also be infinitesimal probabilities -- and one might further think that the proper location for the calculation of expected utility is some non-standard kind of number theory (e.g. Robinson's non-Archimedean number field, Nelson's internal set theory, Conway's extended number system). Again, this suggestion won't entail that the wager argument doesn't work -- it, too, might be construed as a way of developing an argument which has more chance of succeeding^{xii} -- but it

does entail that the success of the argument will depend upon the precise magnitudes which are assigned to the probabilities and values. (Once more, one might conjecture that reasonable atheists and agnostics are *bound* to assign these magnitudes in such a way that the calculation of expected utility favours disbelief.) Perhaps it might be objected that one could not reasonably assign magnitudes to the probabilities and values in such a way that the calculation of expected utility favours disbelief-- or, more moderately, that, as a matter of fact, reasonable people do not assign magnitudes to the probabilities and values in such a way that the calculation of expected utility favours disbelief. However, as we shall see, atheists and agnostics can -- and do -- marshal independent considerations in favour of those assignments of magnitudes which undermine the calculations of expected utility.

Apart from the objections canvassed thus far, there are very strong objections to Pascal's wager which deny the assumption that there are no other relevant alternative Gods (or wagers) which need to be taken into account. If one considers alternative Gods who require alternative beliefs of those upon whom the infinite reward is to be bestowed, and if one awards a similar kind of credence to these alternative Gods, then the wager argument will not go through. Here is a simple example to illustrate the point. Suppose that one also considers the Perverse God who infinitely rewards all and only those who fail to believe in any God. Suppose further that the cost of belief in the Perverse God is C. The pay-off table might then look like this:

		God alone exists (Prob=p)	Perverse God alone exists (Prob=q)	Both exist (Prob=r)	Neither exist (Prob=1-p-q-r)
Believe in God alone	infinite	-B	-B	-B	
Believe in Perverse God alone		-C	-C	-C	-C

Believe in both God and

Perverse God

$\infty - (B+C)$ $\infty - (B+C)$

Don't believe in any God

0

∞ ∞

0

If one does the calculation now, the only alternative which does not yield infinite expected utility -- on the assumption that p , q , and r are all finite and non-zero -- is belief in the Perverse God alone. But the bet is forced: either one believes in God, or the Perverse God, or both, or neither. So what is one to do? The principle which enjoins us to maximise expected utility does not say -- and any other principle which we might choose to invoke will fail in some relevantly similar wager, or so I claim. (In this case, not believing in God might be preferred on the grounds that it costs least, considerations of expected utility apart. But, in other cases, this need not be so.) Since it seems to me to be clear that atheists and agnostics do -- or at any rate would -- think that they have no more reason to believe in God than they have to believe in any of a (possibly even infinite) range of alternative deities, it is clear that this kind of objection triumphs, regardless of the result of the previous two objections. Of course, as I said before, this is not to say that theists have no more reason to believe in God than they have to believe in any of a (possibly even infinite) range of alternative deities; on the contrary, I think that theists have good reason to think that belief in God is much to be preferred to belief in any of the range of allegedly possible alternative deities. But the question before us concerned the dialectical efficacy of Pascal's wager in converting atheists and agnostics, not the reasonableness or coherence of theistic beliefs. Given the intended use of Pascal's wager, the "many gods" and "many wagers" objections are evidently decisive.

In sum: It certainly is not the case that the wager argument is a good argument, given that the wager is "a possible bet". Moreover, it is not even the case that the wager argument is a good argument given both that it is "a possible bet" and that "alternative" gods and wagers are ignored. (It is surely a serious question -- at least for atheists and agnostics -- whether God would reward everyone regardless of their beliefs. How could an all-wise, all-powerful, and all-good God do

otherwise, given the scantiness of the evidence for his existence?) Of course, there may be a coherence requirement on the beliefs of (some) theists which requires them to hold that a relevant wager argument is sound -- if one believes that God will reward all and only those who believe, then one will think that the expected utility calculation is sound -- but this point is simply irrelevant to the question whether the wager argument is any good. Given that the wager argument is meant to persuade atheists and agnostics, it is surely obvious that it will almost never meet with success.

III

Finally, we come to Holcomb's arguments in favour of the claim that "Pascal's wager is not a possible bet", i.e. in favour of the claim that "there is a structural breakdown in the conditions that determine the relations between the act of betting and the payoff matrix options". Holcomb provides two arguments, one which formalises the considerations which he thinks support his conclusion, and one which is intended to parallel Pascal's wager to its discredit, I shall consider these arguments in turn.

First, then, Holcomb provides a formal argument for the conclusion that there is a structural breakdown in Pascal's wager, the core of which is as follows (I use my own numbering):

- (1) According to Pascal's wager, believing and betting are closely related. To bet is to believe or not, in this context. (Justification: this is what is intended by the initial presumption emphasised throughout: the bet is forced.)
- (2) According to Pascal's wager, one bets, not on God's existence, but on belief that God exists or else on belief that God does not exist (Justification: the formulation of Pascal's payoff matrix)
- (3) According to Pascal's wager, believing in God is betting one way, and disbelieving in

God is betting the other way. (Justification: the formulation of Pascal's payoff matrix)

(4) (Hence) Pascal's wager implies that what one bets on is identical to what one bets.

(Justification: logical inference from (1)-(3).)

In this argument, Premise (1) would seem to be redundant, given Premise (3). Consequently, the important questions concern the validity of the inference from (2) and (3) to (4), and the truth of (2) and (3). I shall only address what I take to be the most obvious problem, namely, the evident falsity of the conjunction of (2) and (3).

Referring to the payoff table for Pascal's wager, we see that the relevant possible states of the world are that God exists and that God does not exist, and the possible actions between which we are to decide are belief in God and non-belief in God. Of course, as we have already seen, there are problems hidden in the description of the possible actions: we can't just decide to believe in God, even given that this is what the principle of maximising expected utility enjoins us to do. However, we can come to believe that we ought to believe in God, and we can enter into whatever course of behaviour we deem best suited to bring it about that we do end up believing. Given these complications, what would be the best way to describe the decision problem which we confront in the language of "betting"?

In order to help us answer this question, let us introduce a more standard betting scenario for comparison. Suppose that there is a two horse race, and that we are asked to predict which horse will win. Suppose further that we shall win a huge reward if we are correct in our prediction, and shall incur no cost if we are wrong. In this case, we may think of our act of prediction as the making of a bet: in predicting that Horse 1 will win, I am making a bet that Horse 1 will win. Moreover, in this case, there is a clear sense in which what I bet on is Horse 1; in this sense, obviously, I do NOT bet on my predicting that Horse 1 will win. But with what do I bet? Well, in one sense, with nothing (I put up no stakes) or at least not very much (it doesn't take much effort to make a prediction). But, in another sense, I bet with my prediction (my act of making the prediction constitutes my act of making a bet). We can construct a payoff table for this scenario as

follows:

	Horse 1 wins	Horse 2 wins
Predict Horse 1 wins	BIG	0
Predict Horse 2 wins	0	BIG

Note that it would make no difference to the case if one were to bet by forming a belief rather than by making a prediction. One *could* study the form guide, and form the belief that Horse 1 will win. Of course, one might not be able to form either the belief that Horse 1 will win, or the belief that Horse 2 will win. And in this case -- as in the case of Pascal's wager -- one might only be able to conclude that it would be to one's advantage to form one belief or the other. However, if one were to form the belief that Horse 1 will win, then one would be betting that Horse 1 will win.

Return now to the case of Pascal's wager. By analogy with the previous case, it seems that we should say something like the following: The act of believing is the act of making a bet: in believing that God exists, I am making a bet that God exists. Thus, there is a clear sense in which what I bet on is God's existence (and NOT my believing that God exists). But with what do I bet? Well, in one sense, I bet with whatever it costs me to form and/or have the belief that God exists. (Pascal says that one bets with one's life; roughly, what this means is that one bets with the quality of one's life, given the assumption that, other things being equal, one's life will have less value if one believes (and/or strives to believe) than if one does not.) But, in another sense, I bet with my belief (my act of forming and/or having the belief that God exists constitutes my act of betting on God). Of course, as we noted before, it may be that the most that I can do now is to form the belief that I ought to believe in God -- and, in that case, there is the problem of how to describe the behaviour which one undertakes as one endeavours to bring it about that one believes in God. Does this behaviour constitute betting on God? Well, yes and no. On the one hand, you

can't get the reward unless you believe -- so you might hold that unsuccessful attempts to come to believe in God won't count as bets on God. But, on the other hand, given that the behaviour in question is motivated by the calculation of expected utility, you could equally well count the behaviour as a bet on God, albeit a bet which is not guaranteed to pay off even if God exists.

Finally, we can apply these observations to Holcomb's argument. *Via* the conjunction of (2) and (3), Holcomb claims that, according to Pascal's wager, one bets on belief that God exists (and not on God's existence) by believing that God exists. But, as we have just seen, *first*, it isn't clear that actual believing is necessary for betting on God, though it is indeed sufficient; and, second, if actual believing is necessary for betting on God, then, in believing that God exists, one bets on God's existence, and not on "belief that God exists". As the parallel with the case of betting on horses shows, there is no reason at all to think that there is a structural breakdown in the conditions that determine relations between the act of betting and the payoff matrix options: Pascal's wager is a straightforward case of betting on outcomes under conditions of uncertainty.

Of course, even if there had been a problem in the application of the language of "betting" to the case of Pascal's wager, it would not (necessarily) have followed that there is something wrong with Pascal's argument. As I have already stressed, "Pascal's wager" is an exercise in decision theory. Perhaps any piece of decision-theoretic reasoning can be recast in the language of "betting"; but, if not, it is hard to see why the failure of the recasting should be taken to discredit the decision-theoretic reasoning. (In many cases of decision-theoretic reasoning -- e.g. in cases in which rewards are contingent simply on the making of predictions -- there is no obvious candidate for the role of "stake". To some extent, this is true in the case of Pascal's wager -- where rewards are contingent simply on the formation of beliefs -- though, as I explained above, one could think that the "stake" is (some aspect of) the quality of one's life.)^{xiii}

I turn now to the *second* of the arguments which Holcomb gives in favour of the claim that there is a structural breakdown in Pascal's wager. Holcomb suggests that the wager argument can be paralleled to its obvious discredit:

Suppose our situation is one in which we don't know if it will rain in the near future or not, and are to decide whether or not we shall carry an umbrella. ... To make the assignment of expected returns parallel to the wager argument, let us grant that .. the option of carrying the umbrella when it rains contains a reward of such a higher magnitude than all other possible options that it is infinite and the rest are finite by comparison. Compare the implications of the two matrices, as we reflect on the difference between them that the option in the Rain Matrix is an action and the option in the God Matrix is a belief. ... On the Rain Matrix, carrying an umbrella maximises expected utility. So .. as a rational person acquainted with the Rain Matrix Wager, one ought to believe that one maximises expected utility by carrying an umbrella and one ought to act on that belief by carrying an umbrella. But this does not mean that I should believe that it rains (that is a vertical category, not a decision option), just that I carry an umbrella (whether it rains or not; carrying an umbrella is a decision option). If we construct a prudential argument with exactly the same logic for the decision about whether to believe God exists, we get the following. On the God Matrix, believing God exists maximises expected utility. So ... as a rational person acquainted with the God Matrix Wager, one ought to believe that God exists and one ought to act on that belief by coming to believe that God exists. By parity of reasoning to the Rain Matrix Wager, we should say that this doesn't mean that I should believe that God exists (that is a vertical category, not a decision option), just that I should believe that God exists (whether God exists or not; that is a decision option). But this result is self-contradictory.^{xiv}

There is no genuine contradiction here. In the case of the Rain Matrix, the correct observation is that the fact that carrying an umbrella maximises expected utility *does not entail, via the consideration* that rain is that possible future state of the world which returns the greatest reward for the choice which maximises expected utility, either that it rains or that one should believe that it rains . But, of course, this observation is silent on the question whether the fact that carrying an umbrella maximises expected utility entails either that it rains or that one should believe that it rains *via some other* relevant considerations -- e.g. (one's belief in) the efficacy of one's belief

that it rains in bringing it about that it rains. (If carrying an umbrella while it rains matters so much to one, and if one (thinks that one) can bring about rain merely by believing that it is raining, then presumably one will believe that it is raining. In this case -- given that one believes that one should always act so as to maximise expected utility -- the fact that carrying an umbrella maximises one's expected utility may well entail that (one believes that) it rains.) In the case of the God Matrix, the correct observation is that the fact that believing in God maximises expected utility *does not entail, via the consideration* that God's existence is that possible future state of the world which returns the greatest reward for the choice which maximises expected utility, either that God exists or that one should believe that God exists. But, again, this observation is silent on the question whether the fact that believing in God maximises expected utility entails that God exists or that one should believe that God exists, *via some other* relevant considerations -- e.g. (one's belief) that one's beliefs about supernatural entities always tracks the truth. Moreover -- granted the assumption that one should act always so as to maximise expected utility -- this observation also overlooks the point that the fact that believing in God maximises expected utility *directly* entails that one should believe that God exists (i.e. it entails this without taking the erroneous detour *via* the observation that God's existence is that possible future state of the world which returns the greatest reward for the choice which maximises expected utility). Holcomb's "this doesn't mean that I should believe that God exists, (but) I should believe that God exists" is no contradiction -- there are two ways in which one might seek to reach the conclusion that God exists from an examination of the pay-off matrix, but one of these ways involves an egregious error (and hence "doesn't mean that I should believe that God exists"), while the other does not (and hence "does mean that I should believe that God exists").

In sum: Neither of Holcomb's arguments against Pascal's wager is successful. The decision problem which Pascal envisioned can be consistently described in the language of "betting"; and there is no "violation of the rules of practical reason" involved in the calculations which Pascal makes. Moreover -- waiving considerations about the acceptability of infinite utilities -- there is no reason at all to think that Pascal's wager is not "a possible bet". However, as I argued above, there is every reason to think that it provides no motive at all for atheists and agnostics to make

the bet on God. In particular, it seems incredible that one might think that atheists and agnostics think that there is a straightforward choice between two lone alternatives: the traditional Christian God and nothing. Given the reasons for non-belief to which atheists and agnostics will typically advert -- the lack of good evidence, the weakness of positive arguments, the scope for wishful thinking, etc. -- it is natural for atheists and agnostics to claim that there is every bit as much reason -- or, at the very least, hardly any less reason -- to believe in alterantive deities (or to hold a non-standard conception of the traditional Christian God) as there is to believe in the traditional Christian God. Moreover, this claim is not an *ad hoc* response generated by the argument of Pascal's wager; rather, it is more or less a constitutive feature of traditional atheistic and agnostic worldviews. Consequently, it is plain in advance that the prospects for Pascal's wager are utterly dim: to think otherwise is simply to misunderstand what it is that atheists and agnostics typically believe.^{xv}

ENDNOTES

ⁱInternational Journal for Philosophy of Religion, 35 (1995), pp.65-79. All subsequent numbered references are to this work.

ⁱⁱpp.65-66

ⁱⁱⁱp.66

^{iv}p.66

^v“On Rescher On Pascal's Wager”, International Journal for Philosophy of Religion, 30 (1990), pp.159-168.

^{vi}p.67

^{vii}p.68

^{viii}p.68

^{ix}p.77

^x“Pascal's Wager And Infinite Utilities” Analysis, (1986), pp.107-109.

^{xi}See Alan Hajek's excellent "Waging War On Pascal's Wager", unpublished, for further discussion of this theme.

^{xii}Again, see Hajek's "Waging War On Pascal's Wager" for further discussion.

^{xiii}There is much more which I could say about Holcomb's arguments in favour of the claim that Pascal's wager conflates "what one bets" with "what one bets on". In particular, Holcomb's use of the expression "what one bets" itself involves a conflation between (i) what one bets *with* (i.e. the stakes or the costs); and (ii) what one bets *by* (i.e. what one does in order to make the bet).

Moreover, his claim that one bets *on* belief in God, and not on God's existence, simply ignores the distinction between (i) the ideal case, in which one can simply choose to believe; and (ii) the realistic case, in which the kind of belief in question is not simply and straightforwardly subject to the will. Finally, his distinction between "betting on the scenario 'I believe the bus will come'" and "betting on the scenario 'the bus will come'" pays insufficient attention to the "disquotational" properties of the belief-operator in this context. Armed with these distinctions, one should find it easy to construct replies to the arguments given in sections 3 and 4 of his paper (discussion of which I have foregone for lack of space, time and world.)

^{xiv}pp.73-74

^{xv}My "Weak Agnosticism Defended" International Journal for Philosophy of Religion **36** (1994), pp.147-167, provides a discussion and defence of a kind of agnosticism from a broadly "pragmatist" perspective. This paper marks a departure from the kind of position which I was previously disposed to defend; in particular, it marks a departure from the position which I assumed in "On Rescher On Pascal's Wager". I agree with Holcomb that I there exhibited some tendency towards a "logicist" criticism of the reasonableness of theistic belief. However, while this tendency now seems to me to be mistaken, this fact has no bearing on the cogency of my criticism of Pascal's wager *considered as* an argument which is intended to make atheists and agnostics change their minds. It is one question what reasonable atheists and agnostics typically

believe; it is a quite different question whether reasonable theists should believe this too.