Economics and Philosophy, **34** (2018) 31–44 © Cambridge University Press
doi:10.1017/S0266267117000165 First published online 4 September 2017
cambridge.org/eap

A GAME-THEORETIC ANALYSIS OF PASCAL'S WAGER

Ahmer Tarar*

Abstract: Formal analyses of Pascal's Wager have almost all been decision-theoretic, with a human as the sole decision-maker. This paper analyses Pascal's Wager in a game-theoretic setting in which the deity whose existence the human is considering wagering on is also a decision-maker. There is an equilibrium in which the human chooses to wager that the deity exists and Pascal's Wager thus operates, but also one in which the human does not wager. Thus, in a game-theoretic setting, Pascal's Wager is indeterminate: wagering and not wagering are both consistent with equilibrium behaviour.

Keywords: Pascal's Wager, Decision Theory, Game Theory

1. INTRODUCTION

In the second half of the 17th century, Blaise Pascal introduced the rudiments of decision theory by making one of the most provocative and intriguing arguments in the history of philosophy as well as theology: that the infinite reward of heaven implies that a person who assigns any positive probability, no matter how small, to God existing, should rationally choose to 'wager' that God exists and live one's life accordingly (Pascal 1995 [1670]). There is now an immense philosophy literature that criticizes, defends and further develops Pascal's Wager.[1]

* Department of Political Science, Texas A&M University, 4348 TAMU, College Station, TX 77843-4348, USA. Email: ahmertarar@tamu.edu URL: http://people.tamu.edu/~ahmertarar/

[1] Pascal actually proposed the wager argument under the supposition that the probability that God exists is one-half, but the argument holds for any positive probability (Brams 2011: 72). Jordan (2006) presents a comprehensive treatment of many of the philosophical debates surrounding Pascal's Wager, and Jordan (1994) collects many classic articles on

Following Pascal's approach, virtually all of this literature is decision-theoretic.[2] There is a single decision-maker, a human, who (i) has actions (wager or don't wager), (ii) has preferences over outcomes, and (iii) assigns probabilities to possible states of the world that could be true, in this case whether or not a supreme being or deity exists.[3] The human then chooses the action that provides the highest expected utility.

But from a game-theoretic perspective, such analyses are incomplete. In particular, if the deity exists, then he presumably has the option of revealing his existence to the human, in which case the human would not have to wager at all, and could instead act on certainty. Alternatively, he might choose to not reveal his existence to the human, in which case the human would choose whether or not to wager on the basis of the prior probability that she assigns to the deity existing.[4] If the deity's strategy is to reveal his existence, then if the human does not observe the deity revealing his existence, the human's Bayes' rule inference in a Bayesian equilibrium is to assign probability 0 to the deity existing, and therefore she chooses to not wager even though she begins the interaction assigning positive probability to the deity existing, and assigns an infinite utility for heaven (the outcome where she wagers and the deity exists).

This suggests that a game-theoretic analysis of Pascal's Wager is warranted, which I conduct here. For the human, I use the preferences that are standard in decision-theoretic analyses of Pascal's Wager (e.g. Hájek 2003: 28), and for the deity, I use the reasonable preference ordering stipulated in Brams (1982, 2007, 2011).[5] It turns out that there is an equilibrium in which the human chooses to wager that the deity exists and Pascal's Wager thus operates. However, unlike in the decision-theoretic setting, there is also an equilibrium in which the human chooses to not wager. Because both equilibria exist, we can say that in a game-theoretic setting in which both the human and the deity are decision-makers, Pascal's Wager is indeterminate: wagering and not wagering are both consistent with equilibrium behaviour.

the topic. There is also a literature in economics that applies decision theory to a variety of aspects of religious choice (Durkin and Greeley 1991, 1992; Tabarrok 2000; Østerdal 2004; Melkonyan and Pingle 2009, 2010; Pingle and Melkonyan 2012); for critiques, see Montgomery (1992, 1996).

2 Brams (1982, 2007, 2011) is the only exception, which I discuss in more detail below.

3 Because arguments similar to Pascal's Wager have been made in a number of different religious contexts and traditions (Ryan 1994), I will henceforth simply use the term deity.

4 I use male pronouns for the deity, and female pronouns for the human.

5 Brams presents and analyses what he calls the 'Revelation Game', which is a 2×2 simultaneous-move game of complete information with a deity and a human as players. I analyse a sequential-move game of imperfect information. I discuss the relation between the two models and their results in more detail later.

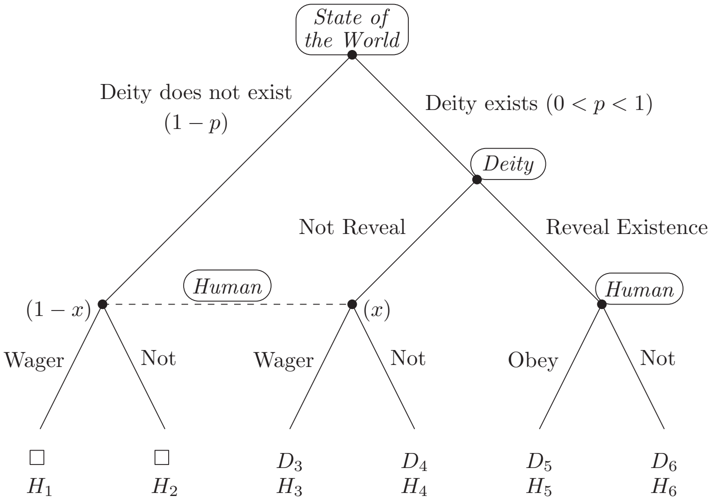

FIGURE 1. The Model

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 presents the main results. To examine the robustness of the main results, I then consider a variety of alternative specifications. Section 4 considers Pascal's first wager payoffs rather than the more standard second wager payoffs (for the human). Section 5 considers an alternative specification of the deity's payoffs. Section 6 considers when all of the human's payoffs are finite. Section 7 concludes.

2. MODEL

The model is shown in Figure 1. It begins with 'chance' or 'nature' probabilistically choosing the state of the world, in this case whether or not the deity exists. The state of the world where the deity exists occurs with probability $0 < p < 1$, and the state where the deity does not exist occurs with probability $1 - p$. The human does not observe this move by chance, but knows the probabilities. Therefore, these probabilities essentially represent the human's prior belief that the deity exists.[6]

[6] That is, portraying the model this way does not imply that p is some objective probability that the deity exists, and the human knows that probability. Instead, probabilistic moves by the fictional player 'chance' are just a device to tractably introduce uncertainty into a model (Harsanyi 1967), and p should be interpreted as this human's subjective prior probability that the deity exists. Assuming that $0 < p < 1$ means that this human assigns positive prior probability to the deity existing, but probability less than 1; Pascal's Wager is directed towards agnostics (Hacking 1972: 188).

https://doi.org/10.1017/S0266267117000165 Published online by Cambridge University Press

If the deity exists, then he decides whether or not to reveal his existence to the human.[7] If he reveals his existence, then the human has to decide whether or not to obey the deity's wishes. If the deity does not reveal his existence, then the human has to decide whether or not to wager that the deity exists (i.e. live her life as though the deity exists).[8] This decision node for the human is part of an information set containing another decision node for the human in which the history that has occurred is that chance chose the state of the world where the deity does not exist. That is, if the human does not observe the deity revealing his existence, the human does not know whether this is because the deity does not exist, or because the deity exists but chose to not clearly reveal his existence.

Decision-theoretic models of Pascal's Wager essentially conduct the entire analysis at the information set, and take the probabilities labelled x (the deity exists) and $1 - x$ (the deity does not exist) as exogenous. The argument is that if the human's utility for heaven (H_3) is infinite and the human's other three payoffs at the information set are finite (e.g. Hájek 2003: 28), then for any $x > 0$ the expected-utility-maximizing choice is to wager. But in a Bayesian equilibrium of a game-theoretic analysis, x is endogenously derived from the deity's strategy using Bayes' rule, and the key result will be that there exists an equilibrium in which $x = 0$ even though $p > 0$, and hence the human chooses to not wager.

Turning to payoffs, there are six possible outcomes of the interaction. Going from left to right, the human's payoffs for these six outcomes are labelled H_1—H_6. Similarly for the deity's payoffs, although these are only defined for outcomes 3—6. For the deity's payoffs, following Brams (1982, 2007: 17—19; 2011: 82), I assume that $D_3 > D_5 > \max\{D_4, D_6\}$. That is, the deity's most-preferred outcome is where he does not reveal his existence and the human chooses to wager, and his second-most-preferred outcome is where he reveals his existence and the human obeys his wishes. Both of these are preferred to the outcomes where the human chooses to not

[7] The interpretation can either be that he reveals his existence to all humans by providing public, universally accessible evidence of his existence, or that he reveals his existence in a personal way to this particular human. For a discussion of this in the context of the 'divine hiddenness' literature in the philosophy of religion, see Schellenberg (1993: 47). Future work might usefully explore a scenario where if the deity exists, he has the opportunity to partially ('noisily') reveal his existence, i.e. send a message that also occurs with positive probability even if the deity does not exist. This would capture the idea that revelation need not be an all-or-nothing matter. My thanks to an anonymous reviewer for pointing this out.

[8] As is well-discussed in the literature on Pascal's Wager, Pascal was not as cynical as the term 'wager' implies. He suggested that living one's life *as though* the deity exists, e.g. by interacting with religious people and engaging in their practices, could eventually lead to true belief (e.g. Hacking 1972: 188).

wager, or chooses to not obey. The deity wants the human to follow his wishes, but ideally without revealing his existence (i.e. on the basis of 'faith' or wagering). (Later, I consider an alternative preference ordering for the deity.)[9]

For the human's payoffs, it is natural to assume that $H_5 > H_6$: if the human *knows* that the deity exists, then presumably the human prefers obeying the deity to not obeying. For the human's payoffs at the information set, I adopt the standard assumption in the decision-theoretic literature on Pascal's Wager that $H_3 = \infty$ (presumably $H_5 = \infty$ as well, although all that is needed for the results below is that $H_5 > H_6$), and that H_1, H_2 and H_4 are finite (e.g. Hájek 2003: 28). Following Pascal's second wager (the major focus of the literature) rather than the first wager, I will also assume that $H_2 > H_1$. That is, if the deity does not exist then the human would rather not wager, and hence the human's choice at the information set is a dilemma (Melkonyan and Pingle 2009). (Later, I consider the first wager, under which $H_2 = H_1$: 'if you lose you lose nothing' by wagering (Pascal 1995 [1670]: 123). I also later consider when all utilities are finite.)[10]

3. RESULTS

For sequential-move games of imperfect information, the standard solution concept is perfect Bayesian equilibrium (henceforth PBE). The proofs of all results are in the Appendix.

Proposition 1 *This game has two PBE:*

- *(a) ('Wager Equilibrium') The deity's strategy is {not reveal}. The human's strategy is {wager; obey}. By Bayes' rule, $x = p$.*
- *(b) ('Don't-Wager Equilibrium') The deity's strategy is {reveal}. The human's strategy is {not wager; obey}. By Bayes' rule, $x = 0$.*

In any equilibrium, the human's strategy is to obey if the deity reveals his existence. In the Wager Equilibrium, the human's strategy is to wager. Because of this, if the deity exists, he chooses to not reveal his existence, as this leads to his most-preferred outcome. Given this strategy of the deity, by Bayes' rule $x = p$ (>0) (no updating from the prior occurs when the

[9] Brams assumes that $D_4 > D_6$, but that is not required for the results below, which hold for any relationship between D_4 and D_6. He analyses a 2×2 simultaneous-move game of complete information. For the human's actions, Brams uses the labels Believe and Don't Believe. In a recent critique of Brams, McShane (2014) disagrees with this preference ordering and argues that religious texts suggest that the deity prefers to reveal his existence.

[10] For a detailed discussion of the first versus the second wager, see Toner and Toner (2006). Hacking (1972) calls the first wager the argument from dominance, and the second wager the argument from dominating expectation.

human does not observe the deity revealing his existence), and hence the human's optimal strategy is indeed to wager (given that $H_3 = \infty$). When $p > 0$, i.e. the human begins the interaction assigning positive probability to the deity existing, there exists an equilibrium in which the human chooses to wager, and hence in which Pascal's Wager operates.

However, there also exists an equilibrium in which the human's strategy is to not wager. Given this strategy choice of the human, if the deity exists, he chooses to reveal his existence, as this leads to his second-most-preferred outcome. And given this strategy choice of the deity, by Bayes' rule $x = 0$ (updating occurs if the human does not observe the deity revealing his existence), and hence the human's utility-maximizing choice is indeed to not wager (given that $H_2 > H_1$).

Thus, in a game-theoretic setting, there exists an equilibrium in which the human does not wager, even though the two foundational premises of Pascal's Wager are being maintained: (i) the human begins the interaction assigning positive probability ($p > 0$) to the deity existing, and (ii) the human assigns infinite utility ($H_3 = \infty$) to heaven.

Because both equilibria exist,[11] we can say that in a game-theoretic setting Pascal's Wager is indeterminate: wagering and not wagering are both consistent with equilibrium behaviour.[12]

[11] In games with multiple equilibria, an interesting question that arises is whether one equilibrium Pareto-dominates the others, i.e. is strictly preferred by all decision-makers. The deity strictly prefers the Wager Equilibrium, as it results in his most-preferred outcome (with payoff D_3). If we assume that H_5 is infinite, then *ex ante* the human strictly prefers the Don't-Wager Equilibrium. This is because in both equilibria, with probability p the human gets an infinite payoff. With probability $1 - p$, the human gets H_2 in the Don't-Wager Equilibrium, but only H_1 in the Wager Equilibrium. Hence, neither equilibrium Pareto-dominates the other if $H_5 = \infty$. On the other hand, if H_5 is finite, then even the human strictly prefers the Wager Equilibrium, which thus Pareto-dominates the Don't-Wager Equilibrium.

[12] Brams (1982, 2007, 2011) presents the only other game-theoretic analyses in which a deity and a human are decision-makers. His 'Revelation Game' (Brams 2011: 82), which is the one most similar to this model, is a 2×2 simultaneous-move game of complete information. The human's two actions are Believe and Don't Believe, and the deity's two actions are Reveal and Don't Reveal. A game of imperfect information, as I analyse here, is more appropriate for modelling an interaction in which the human is uncertain about whether or not the deity exists. The Revelation Game has a unique Nash equilibrium, in which the human chooses Don't Believe (contrary to Pascal's Wager). The game has two 'nonmyopic equilibria' (a non-standard solution concept based on the 'theory of moves'; Brams 1994), in both of which the human chooses Believe (consistent with Pascal's Wager; but Brams emphasizes transitions between belief and nonbelief over generational time, as shifts between the two nonmyopic equilibria occur due to the players alternately possessing 'moving power'). Thus, the validity of Pascal's Wager in the Revelation Game depends on what solution concept is used. I use a standard solution concept, PBE, and find that there exists a PBE in which Pascal's Wager operates, but also one in which it does not.

4. FIRST WAGER PAYOFFS

Although the second wager is the major focus of the literature because the human's choice is actually a dilemma, it turns out that the Don't-Wager Equilibrium exists even when the human has more of an incentive to wager, namely the first wager payoffs $H_2 = H_1$: there is nothing to lose in wagering.

Proposition 2 *When $H_2 = H_1$, the game has three PBE:*

(a) ('Wager Equilibrium') The deity's strategy is {not reveal}. The human's strategy is {wager; obey}. By Bayes' rule, $x = p$.
(b) ('Don't-Wager Equilibrium') The deity's strategy is {reveal}. The human's strategy is {not wager; obey}. By Bayes' rule, $x = 0$.
(c) (Mixed Strategy Equilibria) The deity's strategy is {reveal}. The human's strategy is {wager with any probability $0 < y \leqslant (D_5 - D_4)/(D_3 - D_4) \in (0, 1)$ and not wager with probability $1 - y$; obey}. By Bayes' rule, $x = 0$.

The Don't-Wager Equilibrium still exists, because when $x = 0$ the human can be choosing to not wager, as she is indifferent between wagering and not wagering. Even if there is nothing to be lost by wagering, not wagering is still consistent with equilibrium behaviour. (Because she is indifferent, there also exists a third class of equilibria (c) in which the human mixes. In them, she chooses to wager with low enough probability that the deity chooses to reveal, thus ensuring that $x = 0$ and hence the human is indifferent. In these equilibria the human is choosing to not wager with positive probability, reinforcing the point about not wagering being consistent with equilibrium behaviour.)[13]

5. AN ALTERNATIVE PREFERENCE ORDERING FOR THE DEITY

It turns out that there is a way of resolving this game-theoretic 'problem' for Pascal's Wager. Suppose that the deity's preference ordering is instead $\min\{D_3, D_4\} > \max\{D_5, D_6\}$. That is, the deity's optimal strategy is to not reveal his existence, regardless of what the human's strategy is: not revealing is a strictly dominant strategy.[14]

Proposition 3 *Regardless of whether $H_2 > H_1$ (second wager payoffs) or $H_2 = H_1$ (first wager payoffs), this game has a unique PBE, a 'Wager Equilibrium': the*

[13] If the first wager is instead interpreted to imply that $H_2 < H_1$, then of course the Wager Equilibrium is the unique equilibrium. At one point Pascal (1995 [1670]: 125) goes this far, writing: 'Now what harm will come to you from choosing this course? You will be faithful, honest, humble, grateful, full of good works, a sincere, true friend ... It is true you will not enjoy noxious pleasures, glory and good living, but will you not have others? I tell you that you will gain even in this life ...'

[14] Presumably $D_3 > D_4$ and $D_5 > D_6$, but this is not needed for the results.

deity's strategy is {not reveal}. The human's strategy is {wager; obey}. By Bayes' rule, $x = p$.

With this alternative preference ordering, in any PBE, the deity's strategy is to not reveal, and hence by Bayes' rule $x = p$ (>0). The human therefore chooses to wager (given that $H_3 = \infty$). If the deity's strictly dominant strategy is to not reveal his existence, then a game-theoretic setting poses no problem for Pascal's Wager; there is a unique equilibrium, and in it the human chooses to wager.

6. ALL UTILITIES ARE FINITE

Given that the assumptions that justify expected utility maximization as being the appropriate decision-making principle under risk (von Neumann and Morgenstern 1944; Savage 1954) do not allow for infinite utilities (e.g. McClennen 1994; Østerdal 2004; Melkonyan and Pingle 2009), it is worth considering how this game-theoretic 'problem' for Pascal's Wager fares under the alternative assumption that H_3 and H_5 (the heaven payoffs) are finite. We continue to assume that $H_5 > H_6$ and $H_3 > H_4$.

Then at the information set, the human strictly prefers to wager if and only if $(x)(H_3) + (1 - x)(H_1) > (x)(H_4) + (1 - x)(H_2)$, which can be rewritten as $x > (H_2 - H_1)/[(H_2 - H_1) + (H_3 - H_4)]$. That is, the human chooses to wager if and only if the probability that the human assigns to the deity existing when choosing whether or not to wager exceeds a certain threshold (McClennen 1994: 126; Mougin and Sober 1994: 383). Let us call this threshold $x_{critical}$. If the human has the first wager payoffs, i.e. $H_2 = H_1$, then $x_{critical} = 0$. If the human has the second wager payoffs, i.e. $H_2 > H_1$, then $x_{critical} \in (0, 1)$.

6.1. First Wager Payoffs ($H_2 = H_1$)

With the first wager payoffs, the finite-utility results are exactly the same as the infinite-utility results, for both preference orderings for the deity.

Proposition 4 *If the deity has the original contingent-revealer preference ordering $D_3 > D_5 > \max\{D_4, D_6\}$, then the game has the same three PBE as in Proposition 2:*

(a) ('Wager Equilibrium') The deity's strategy is {not reveal}. The human's strategy is {wager; obey}. By Bayes' rule, $x = p$.
(b) ('Don't-Wager Equilibrium') The deity's strategy is {reveal}. The human's strategy is {not wager; obey}. By Bayes' rule, $x = 0$.
(c) (Mixed Strategy Equilibria) The deity's strategy is {reveal}. The human's strategy is {wager with any probability $0 < y \leqslant (D_5 - D_4)/(D_3 - D_4) \in (0, 1)$ and not wager with probability $1 - y$; obey}. By Bayes' rule, $x = 0$.

https://doi.org/10.1017/S0266267117000165 Published online by Cambridge University Press

Proposition 5 *If the deity has the strictly-dominant-non-revealer preference ordering* $\min\{D_3, D_4\} > \max\{D_5, D_6\}$*, then the game has the same unique PBE of Proposition 3, a 'Wager Equilibrium': the deity's strategy is {not reveal}. The human's strategy is {wager; obey}. By Bayes' rule,* $x = p$.

To the point, if the deity has the original contingent-revealer preference ordering, then there exists a Don't-Wager Equilibrium. But if the deity's strictly dominant strategy is to not reveal his existence, then the Wager Equilibrium is the unique equilibrium.

6.2. Second Wager Payoffs ($H_2 > H_1$)

With the second wager payoffs, however, the results change.

Proposition 6 *Suppose that the deity has the original contingent-revealer preference ordering* $D_3 > D_5 > \max\{D_4, D_6\}$.

(i) If $p < x_{critical}$ *then the game has a unique PBE, a 'Don't-Wager Equilibrium': the deity's strategy is {reveal}. The human's strategy is {not wager; obey}. By Bayes' rule,* $x = 0$.

(ii) If $p \geqslant x_{critical}$ *then the game has two PBE in pure strategies:*
  - *(a) ('Wager Equilibrium') The deity's strategy is {not reveal}. The human's strategy is {wager; obey}. By Bayes' rule,* $x = p$.
  - *(b) ('Don't-Wager Equilibrium') The deity's strategy is {reveal}. The human's strategy is {not wager; obey}. By Bayes' rule,* $x = 0$.

(iii) If $p = x_{critical}$ *(a knife-edge condition) then the game also has the following PBE in mixed strategies: the deity's strategy is {not reveal}. The human's strategy is {wager with any probability* $y_{critical} \leqslant y < 1$ *and not wager with probability* $1 - y$*, where* $y_{critical} = (D_5 - D_4)/(D_3 - D_4) \in (0, 1)$*; obey}. By Bayes' rule,* $x = p\ (= x_{critical})$.

(iv) If $p > x_{critical}$ *then the game also has the following PBE in mixed strategies: the deity's strategy is {not reveal with probability* $z = [(1 - p)x_{critical}]/[p(1 - x_{critical})] \in (0, 1)$ *and reveal with probability* $1 - z$*}. The human's strategy is {wager with probability* $y = (D_5 - D_4)/(D_3 - D_4) \in (0, 1)$ *and not wager with probability* $1 - y$*; obey}. By Bayes' rule,* $x = [pz]/[pz + (1 - p)] = x_{critical}$.

With a finite utility for heaven and the deity having the original contingent-revealer preference ordering, the game-theoretic 'problem' for Pascal's Wager is even stronger than in the infinite-utility case. As in the infinite-utility case, a Don't-Wager Equilibrium always exists. But in the infinite-utility case, a Wager Equilibrium also always exists, whereas in the finite-utility case a Wager Equilibrium only exists when $p \geqslant x_{critical}$, i.e. the human's prior belief that the deity exists is at least as large as her belief-threshold for wagering. If the human begins the interaction assigning a sufficiently low probability to the deity existing (i.e. $p <$

https://doi.org/10.1017/S0266267117000165 Published online by Cambridge University Press

$x_{critical}$), then the Don't Wager Equilibrium is the *unique* one.[15] (When $p \geqslant x_{critical}$, then mixed-strategy equilibria also exist. In them, the human chooses to not wager with positive probability, reinforcing the 'problem.')

Proposition 7 *Suppose that the deity has the strictly-dominant-non-revealer preference ordering* $\min\{D_3, D_4\} > \max\{D_5, D_6\}$.

(i) *If* $p < x_{critical}$ *then the game has a unique PBE, a 'Don't-Wager Equilibrium': the deity's strategy is {not reveal}. The human's strategy is {not wager; obey}. By Bayes' rule,* $x = p$.
(ii) *If* $p > x_{critical}$ *then the game has a unique PBE, a 'Wager Equilibrium': the deity's strategy is {not reveal}. The human's strategy is {wager; obey}. By Bayes' rule,* $x = p$.
(iii) *If* $p = x_{critical}$ *(a knife-edge condition) then the above two pure-strategy PBE exist, as well as the following PBE in mixed strategies: the deity's strategy is {not reveal}. The human's strategy is {wager with any probability* $y \in (0, 1)$ *and not wager with probability* $1 - y$*; obey}. By Bayes' rule,* $x = p$ $(= x_{critical})$.

With a finite utility for heaven and the deity having the strictly-dominant-non-revealer preference ordering, the 'problem' is again stronger than in the infinite-utility case. In the infinite-utility case, this preference ordering for the deity ensures that the Wager Equilibrium is the unique one, and hence completely resolves the 'problem'. But in the finite-utility case, this requires $p > x_{critical}$; if $p < x_{critical}$, then the Don't-Wager Equilibrium not only exists, it is in fact the unique one. With a finite utility for heaven, even the deity having a strictly dominant strategy of not revealing his existence does not ensure that the human chooses to wager; instead, the human's strategy choice depends on the value of the prior p relative to her belief-threshold $x_{critical}$. (When the knife-edge condition $p = x_{critical}$ holds, then mixed-strategy equilibria also exist.)

7. CONCLUSION

Almost all existing formal analyses of Pascal's Wager are decision-theoretic, with a human as the sole decision-maker. However, if the deity whose existence the human is considering wagering on exists, then presumably the deity has the option of choosing whether or not to clearly reveal his existence to the human. This suggests that a game-theoretic analysis of Pascal's Wager is warranted, which is what I conduct here.

[15] The intuition is that when $p < x_{critical}$, then even if the deity's strategy is to not reveal his existence, by Bayes' rule $x = p$ ($<x_{critical}$), and hence the human would choose to not wager.

Using the standard infinite-utility-for-heaven Pascal's Wager preferences for the human, and Brams' (1982, 2007, 2011) plausible contingent-revealer preference ordering for the deity, it turns out that a Wager Equilibrium exists in which the human chooses to wager and Pascal's Wager thus operates, but a Don't-Wager Equilibrium also exists (even if the human has the first wager preferences in which there is nothing to lose by wagering). Thus, in a game-theoretic setting Pascal's Wager is indeterminate: wagering and not wagering are both consistent with equilibrium behaviour.

If the deity instead has a strictly dominant strategy of not revealing his existence, then wagering is the unique equilibrium outcome and hence Pascal's Wager is 'restored'. But with a finite utility for heaven, the problem is more severe, and assuming that the deity's strictly dominant strategy is to not reveal his existence does not necessarily resolve it. In summary, Pascal's Wager faces some novel problems in a game-theoretic setting that do not arise in the decision-theoretic setting.

ACKNOWLEDGEMENTS

My thanks to the participants in a Program in Ethics and Political Science workshop at Texas A&M University in May 2016, as well as to two anonymous reviewers.

References

Brams, S. J. 1982. Belief in God: a game-theoretic paradox. *International Journal for Philosophy of Religion* 13: 121–129.

Brams, S. J. 1994. *Theory of Moves*. Cambridge: Cambridge University Press.

Brams, S. J. 2007. *Superior Beings: If They Exist, How Would We Know?* (2nd edn). New York, NY: Springer.

Brams, S. J. 2011. *Game Theory and the Humanities*. Cambridge, MA: MIT Press.

Durkin, J. and A. Greeley. 1991. A model of religious choice under uncertainty: on responding rationally to the nonrational. *Rationality and Society* 3: 178–196.

Durkin, J. and A. Greeley. 1992. Response to Montgomery. *Rationality and Society* 4: 121–122.

Hacking, I. 1972. The logic of Pascal's wager. *American Philosophical Quarterly* 9: 186–192.

Hájek, A. 2003. Waging war on Pascal's wager. *Philosophical Review* 112: 27–56.

Harsanyi, J. 1967. Games of incomplete information played by 'Bayesian' players I. *Management Science* 14: 159–182.

Jordan, J., ed. 1994. *Gambling on God: Essays on Pascal's Wager*. Lanham, MD: Rowman and Littlefield.

Jordan, J. 2006. *Pacal's Wager: Pragmatic Arguments and Belief in God*. Oxford: Clarendon Press.

McClennen, E. F. 1994. Pascal's wager and finite decision theory. In *Gambling on God: Essays on Pascal's Wager*, ed. J. Jordan, 115–138. Lanham, MD: Rowman and Littlefield.

McShane, P. J. 2014. Game theory and belief in God. *International Journal for Philosophy of Religion* 75: 3–12.

Melkonyan, T. and M. Pingle. 2009. Religion and faith: a decision theory perspective. *Interdisciplinary Journal of Research on Religion* 5: 1–19.
Melkonyan, T. and M. Pingle. 2010. Ambiguity, pessimism, and rational religious choice. *Theory and Decision* 69: 417–438.
Montgomery, J. D. 1992. Pascal's wager and the limits of rational choice: a comment on Durkin and Greeley. *Rationality and Society* 4: 117–122.
Montgomery, J. D. 1996. Contemplations on the economic approach to religious behaviour. *American Economic Review* 86: 443–448.
Mougin, G. and E. Sober. 1994. Betting against Pascal's wager. *Noûs* 28: 382–395.
Østerdal, L. P. 2004. Pascal's and Tabarrok's wagers. *Theory and Decision* 57: 1–4.
Pascal, B. 1995 [1670]. *Pensées*, ed. A. J. Krailsheimer. New York, NY: Penguin Books.
Pingle, M. and T. Melkonyan. 2012. To believe or not believe ... or not decide: a decision-theoretic model of agnosticism. *Rationality and Society* 24: 408–441.
Ryan, J. K. 1994. The wager in Pascal and others. In *Gambling on God: Essays on Pascal's Wager*, ed. J. Jordan, 11–19. Lanham, MD: Rowman and Littlefield.
Savage, L. J. 1954. *The Foundations of Statistics*. New York, NY: Wiley.
Schellenberg, J. L. 1993. *Divine Hiddenness and Human Reason*. Ithaca, NY: Cornell University Press.
Tabarrok, A. 2000. Believe in Pascal's wager? Have I got a deal for you! *Theory and Decision* 48: 123–128.
Toner, P. and C. Toner. 2006. Pascal's first wager reconsidered: a virtue theoretic view. *International Philosophical Quarterly* 46: 75–90.
von Neumann, J. and O. Morgenstern. 1944. *Theory of Games and Economic Behaviour*. New York, NY: John Wiley and Sons.

APPENDIX

This appendix presents the proofs of all of the propositions. Since we assume that $H_5 > H_6$ throughout, in every PBE the human chooses to obey.

PROPOSITION 1: Suppose there is a PBE in which $x > 0$. Then the human's strictly optimal choice is to wager. Given this, the deity's strictly optimal choice is to not reveal. Then $x = p$ (>0) by Bayes' rule. This gives PBE (a).

Suppose there is a PBE in which $x = 0$. Then the human's strictly optimal choice is to not wager. Given this, the deity's strictly optimal choice is to reveal. Then $x = 0$ by Bayes' rule. This gives PBE (b).

We have analysed every possible value of x, and hence there are no other PBE. Q.E.D.

PROPOSITION 2: For PBE (a), the existence argument is exactly the same as in Proposition 1.

Suppose there is a PBE in which $x = 0$. Then the human is indifferent between wagering and not wagering. CASE (I): Suppose the human is choosing to not wager. Given this, the deity's strictly optimal choice is to reveal. Then $x = 0$ by Bayes' rule. This gives PBE (b). CASE (II): Suppose the human is choosing to wager. Given this, the deity's strictly optimal choice is to not reveal. But then $x = p$ (>0) by Bayes' rule, inconsistent with our supposition that $x = 0$, and hence there is no PBE here. CASE (III): Suppose the human is choosing to wager with probability $y \in (0, 1)$ and not wager with probability $1 - y$. Maintaining our supposition that $x = 0$ requires that the deity be choosing to reveal. The deity at least weakly prefers to reveal if and only if $(y)(D_3) + (1 - y)(D_4) \leqslant D_5$, which can be rewritten as $y \leqslant (D_5 - D_4)/(D_3 - D_4) \in (0, 1)$. This gives PBE (c).

https://doi.org/10.1017/S0266267117000165 Published online by Cambridge University Press

We have analysed every possible value of x, and hence there are no other PBE. Q.E.D.

PROPOSITION 3: Because not revealing is the deity's strictly dominant strategy, in any PBE the deity chooses to not reveal. Therefore $x = p$ (>0) by Bayes' rule. Therefore the human's strictly optimal choice is to wager. This gives the PBE, and also establishes its uniqueness. Q.E.D.

PROPOSITIONS 4 AND 5: The key thing here is to notice that when $H_2 = H_1$, then $x_{critical} = 0$: the human strictly prefers to wager if and only if $x > 0$, and is indifferent if $x = 0$. Therefore, the exact same proofs as for Propositions 2 and 3 carry through. Q.E.D.

PROPOSITION 6: When $H_2 > H_1$, then $x_{critical} \in (0, 1)$. Suppose there is a PBE in which the human's strategy is to not wager. Given this, the deity's strictly optimal choice is to reveal. Then $x = 0$ ($< x_{critical}$) by Bayes' rule, and hence the human's strategy of not wagering is indeed optimal (in fact, strictly). This gives the existence of PBE (i) and (ii)(b), and establishes that the Don't-Wager Equilibrium exists regardless of the value of p relative to $x_{critical}$.

Suppose there is a PBE in which the human's strategy is to wager. Given this, the deity's strictly optimal choice is to not reveal. Then $x = p$ by Bayes' rule, and hence the human's strategy of wagering is (at least weakly) optimal if and only if $p \geqslant x_{critical}$. This gives PBE (ii)(a), and establishes that the Wager Equilibrium exists if and only if $p \geqslant x_{critical}$.

Finally, suppose there is a PBE in which the human's strategy is: wager with probability $y \in (0, 1)$ and not wager with probability $1 - y$. This requires $x = x_{critical}$, so that the human is indifferent. The deity strictly prefers to reveal if and only if $(y)(D_3) + (1 - y)(D_4) < D_5$, which can be rewritten as $y < (D_5 - D_4)/(D_3 - D_4) \equiv y_{critical} \in (0, 1)$. CASE (I): First suppose that $0 < y \leqslant y_{critical}$ and the deity is choosing to reveal. Then $x = 0$ by Bayes' rule, which is inconsistent with our requirement that $x = x_{critical}$ (>0), and hence there is no PBE here. CASE (II): Now suppose that $y_{critical} \leqslant y < 1$ and the deity is choosing to not reveal. Then $x = p$ by Bayes' rule, and hence this is a PBE if and only if the knife-edge condition $p = x_{critical}$ holds. This gives PBE (iii). CASE (III): Finally, suppose that $y = y_{critical}$, and the deity is choosing to not reveal with probability $z \in (0, 1)$, and reveal with probability $1 - z$. This is a PBE as long as z is such that $x = x_{critical}$. By Bayes' rule, $x = [pz]/[pz + (1 - p)]$, and setting this equal to $x_{critical}$ and solving for z gives $z = [(1 - p)x_{critical}]/[p(1 - x_{critical})]$. Note that $z > 0$. Also, $z < 1$ can be simplified to $p > x_{critical}$, and hence this PBE exists if and only if $p > x_{critical}$. This gives PBE (iv), and also establishes the uniqueness claim in the statement of PBE (i). Q.E.D.

PROPOSITION 7: Because not revealing is the deity's strictly dominant strategy, in any PBE the deity chooses to not reveal. Therefore $x = p$ by Bayes' rule. CASE (I): If $p < x_{critical}$, then $x = p < x_{critical}$, and hence the human's strictly optimal choice is to not wager. This gives PBE (i), and also establishes its uniqueness. CASE (II): If $p > x_{critical}$, then $x = p > x_{critical}$, and hence the human's strictly optimal choice is to wager. This gives PBE (ii), and also establishes its uniqueness. CASE (III): If the knife-edge condition $p = x_{critical}$ holds, then $x = p = x_{critical}$, and hence the human is indifferent between wagering and not wagering. Hence PBE (i) and (ii) both exist, as do mixed-strategy PBE in which the human chooses to wager with any probability $y \in (0, 1)$, and to not wager with probability $1 - y$. This gives PBE (iii). Q.E.D.

BIOGRAPHICAL INFORMATION

Ahmer Tarar is an Associate Professor in the Department of Political Science at Texas A&M University. Most of his research is on international conflict and the causes of war. More broadly, his research applies game theory to a variety of topics in the social sciences and the humanities.

https://doi.org/10.1017/S0266267117000165 Published online by Cambridge University Press