

---

Game theory and belief in God

Author(s): Paddy Jane McShane

Source: *International Journal for Philosophy of Religion*, Vol. 75, No. 1 (February 2014), pp. 3-12

Published by: Springer

Stable URL: <https://www.jstor.org/stable/24709202>

Accessed: 13-02-2019 21:37 UTC

---

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

Springer is collaborating with JSTOR to digitize, preserve and extend access to *International Journal for Philosophy of Religion*

## Game theory and belief in God

Paddy Jane McShane

Received: 3 August 2012 / Accepted: 30 January 2013 / Published online: 22 February 2013  
© Springer Science+Business Media Dordrecht 2013

**Abstract** In the last few decades game theory has emerged as a powerful tool for examining a broad range of philosophical issues. It is unsurprising, then, that game theory has been taken up as a tool to examine issues in the philosophy of religion. Economist Steven Brams (1982), (1983) and (2007), for example, has given a game theoretic analysis of belief in God, his main argument first published in this journal and then again in both editions of his book, *Superior Beings*. I have two main aims in this paper, one specific and one general. My specific aim is to show that Brams' application of game theory to examine belief in God is, in particular, deeply flawed in two respects. My general aim is to show that *any* game-theoretic model in which a human being and God are players can only succeed at the cost of abandoning the assumption that God is omnibenevolent.

**Keywords** Game theory · Belief in God · Divine perfection · Omnibenevolence

Game theory is, most basically, the study of multi-agent strategic interactions. Taking agents' preferences as inputs, game-theoretic models present outcomes of strategic interactions among agents in terms of agents' preferences. In the last few decades game theory has emerged as a powerful tool for examining a broad range of philosophical issues. David Lewis (1969) has employed game theory to undergird a theory of conventions, generally, and a theory of semantics and linguistic conventions, more specifically ; Kenneth Binmore (1994) and (2008) has used game theory to underwrite a Rawlsian conception of distributive justice ; David Gauthier (1986) has relied on game theory to defend a contractualist account of morality that, he argues, is consistent

---

P. J. McShane (✉)  
Georgetown University, 2145 Goss Circle, Apt. 4, Boulder, CO 80302, USA  
e-mail: [pjm84@georgetown.edu](mailto:pjm84@georgetown.edu)

with his economic conception of rationality. These are merely a few of the many uses of game theory for understanding philosophical issues.

It is unsurprising, then, that game theory has been taken up as a tool to examine issues in the philosophy of religion. Economist Steven Brams (1982), (1983) and (2007), for example, has given a game-theoretic analysis of belief in God, his main argument first published in this journal and then again in both editions of his book, *Superior Beings*. Many philosophers have tackled the issue of whether it is rational to believe in God by relying on decision theory, the study of principles of rational choice that apply to *single* agents who are in situations of uncertainty or risk. A well-know example of this is, of course, “Pascal’s Wager.” Many critics have voiced concerns, by now familiar, with using decision theory to recommend belief in God in the way that Pascal does. However, while philosophers have explored in depth the application of decision theory to the philosophical issues surrounding belief in God, until Brams’ work there has been relatively little discussion about game theory and belief in God. My aim in this paper is to deepen this discussion. I am motivated to do so, most basically, for two reasons: (1) I think that Brams’ use of game theory to examine belief in God is deeply flawed, and (2) given that game theory has proved to be a powerful tool for examining other philosophical issues, it is worthwhile to see if it is suited for the task of examining the rationality of belief in God and to see what limits, if any, there are lurking in the application of game theory to philosophical issues.

In this paper I first lay out Brams’ “Belief Game”: his game-theoretic representation of belief in God with “Man” and God as players. Next, I argue that the “Belief Game” is flawed in (at least) two respects: it relies on a mischaracterization of God’s preferences and it also fails to adequately address the fact that from the perspective of one of the players (“Man”), it is uncertain whether or not the other player (God) even exists. Finally, I present a deeper problem for the “Belief Game,” arguing that *any* game theoretic model in which a human being and God are players can only succeed at the cost of abandoning the assumption that God is omnibenevolent.

### Brams’ belief game

In his article, Brams presents a two-person, non-zero sum game—the “Belief Game”—and argues that the game sheds light on why it is rational for God not to reveal Himself and why our reasons for belief in God are tenuous. A game is zero-sum if it is a feature of the game that, for each potential outcome, the gains of one player are entirely offset by the losses of another, and vice versa (in economic terms, at each possible outcome point, the sum of the players’ utility functions is zero); a non-zero sum game lacks this feature.

Brams’ game involves two players—God and “Man.” God has two possible strategies: to reveal Himself or not to reveal Himself; as the game theory matrix below reveals, Brams argues that these two strategies are equivalent to God establishing or not establishing His existence, respectively. “Man” also has two strategies: to believe or not to believe in God’s existence. To come up with a payoff matrix, Bram assumes the following ordered preferences of God and “Man”:

God (1) wants man to believe in his existence and (2) prefers not to reveal Himself, while “Man” (1) wants belief (or nonbelief) in God’s existence to be confirmed by the evidence (or lack thereof) and (2) prefers to believe in God’s existence. The strategies and outcomes of Brams’ Belief Game, then, are as follows:

	Brams’ belief game	Man	
		Believe in God’s existence	Don’t believe in God’s existence
God	<b>Reveal Himself</b> (i.e., Establish His existence)	3, 4 (Man faithful with evidence: belief in existence confirmed)	1, 1 (Man unfaithful despite evidence: nonbelief in existence disconfirmed)
	<b>Don’t reveal Himself</b> (i.e., Don’t establish His existence)	4, 2 (Man faithful without evidence: belief in existence disconfirmed)	2, 3 (Man unfaithful without evidence: nonbelief in existence confirmed)

Brams (1982, p. 122)

The “Belief Game” rests on a number of key assumptions. First, the numbers in the payoff matrix have only ordinal—and not cardinal—significance. The first number in each box represents the payoff to God and the second represents the payoff to “Man”. 4 is the best outcome for a player, 3 the next best, 2 the next worst, and 1 the worst. Second, God is assumed to be an “active entity that is capable of making choices,” and an entity who has preferences. Third, the “Belief Game” is a game of complete information; as Brams says, “. . . man knows God’s preferences as well as his own.” Finally, the players’ strategies are assumed to be entirely independent of one another and “the players’ choices,” Brams (1982 pp. 122–123) assumes, “are made in ignorance of each other”.

As the payoff matrix above reveals, Brams argues that God most prefers not to reveal Himself while “Man” faithfully believes in His existence, while “Man” most prefers to believe in God’s existence while God reveals Himself. Given the payoffs, regardless of what “Man” does, God gets the best payoff if he does not reveal Himself. Not revealing Himself, then, is God’s dominant strategy. Anticipating that God will not reveal Himself regardless of what “Man” does, “Man” rationally chooses not to believe in God’s existence in order to get a higher payoff (3 rather than 2). The outcome of the strategic interaction between God and “Man,” then, is 2, 3. Notice, however, that there is another outcome—namely 3, 4—at which each player could be better off without thereby making the other player worse off; 3, 4 is the Pareto-optimal outcome. Brams thus calls the Belief Game a paradox “somewhat akin” to a prisoners’ dilemma because the outcome of the interaction between God and “Man” (2,3), acting rationally, end up at a Pareto-inferior equilibrium (2, 3 rather than 3, 4).<sup>1</sup>

<sup>1</sup> Since only one player—God—has a dominant strategy (Don’t reveal Himself), Brams’ Belief Game differs from a standard prisoners’ dilemma in which both players have a dominant strategy (defect). This sort of dilemma, in which one player has a dominant strategy and the players end up at a Pareto-inferior outcome is what is known as a “common-knowledge prisoners’ dilemma. In his entry on the *Stanford Encyclopedia of Philosophy*, “Prisoner’s Dilemma,” Kuhn identifies the conditions of a common knowledge prisoners’ dilemma. He begins by laying out a standard prisoners’ dilemma as follows:

## Two problems with Brams' belief game

There are (at least) two problems with Brams' Belief Game: (1) Brams mischaracterizes God's preferences, a fault which leads him to incorrectly fill in the payoff matrix; (2) Brams fails to adequately address the fact that from the perspective of one of the players ("Man"), the other player (God) may or may not exist. I will address (1) and (2), in turn.

Recall that Brams' assumes that God has the following ordered preferences: God (1) wants "Man" to believe in his existence and He (2) prefers not to reveal Himself. Brams, on the assumption that the Old Testament is true, calls upon the God of the Old Testament to support (2), his assumption that God prefers not to reveal Himself. However, there is reason to doubt that the Old Testament really does support Bram's understanding of God's preferences.

First, God does reveal Himself in the Old Testament. For example, He appears to Moses:

And the angel of the Lord appeared unto him in a flame of fire out of the midst of a bush: and he looked, and, behold, the bush burned with fire, and the bush was not consumed. And Moses said, I will now turn aside, and see this great sight, why the bush is not burnt. And when the Lord saw that he turned aside to see, God called unto him out of the midst of the bush, and said, Moses, Moses. And he said, Here am I (Exod 3:1–4 RSV).

The most reasonable interpretation of this Biblical passage from the Old Testament, if we take the Old Testament to be true, is that God did in fact reveal Himself to Moses in the burning bush. Puzzling enough, Brams (1982 p. 127) himself tentatively acknowledges that God revealed Himself to Moses; he writes, "He [God]

Footnote 1 continued

*Standard prisoners' dilemma*

	Cooperate	Defect
Cooperate	Rr, Rc	Sr, Tc
Defect	Tr, Sc	Pr, Pc

Here R represents the reward for both players if they both cooperate, S represents the "sucker" payoff for the player who, alone, cooperates, T represents the temptation payoff for the player who, alone, defects, and P represents the punishment that each player receives if both players defect. In a standard prisoners' dilemma, the following inequalities hold true, and both players have a dominant strategy of defection:

- (a)  $Tr > Rr$  and  $Pr > Sr$
- (b)  $Tc > Rc$  and  $Pc > Sc$
- (c)  $Rr > Pr$  and  $Rc > Pc$

However, Kuhn points out that "the force of the [prisoner's] dilemma can be felt under weaker conditions"—the conditions of the common knowledge prisoners' dilemma. In a common knowledge prisoners' dilemma, either (a) or (b)—but not both—fails to hold. We can see that in the Belief Game (b) is not satisfied. If God reveals Himself, "Man's" payoff if he doesn't believe in God's existence (1) is not greater than "Man's" payoff if he does believe in God's existence (4), so  $Tc$  is not greater than  $Rc$ . This entails that "Man" does not have a dominant strategy; however, since both players are rational and since both players know the ordering of the other player's preferences, the players end up at a Pareto-inferior equilibrium, as in the standard prisoners' dilemma. The Belief Game, then, is a common knowledge prisoners' dilemma.

never revealed Himself in any physical form, except possibly to Moses before he died, though he continually demonstrated His powers in other ways, notably by punishing those he considered transgressors”. However, Brams is quiet on the issue of why it is that he thinks that God only *possibly* revealed Himself to Moses. Until Brams gives us further evidence to undermine the most straightforward reading of the above quoted passage from Exodus, we are left to doubt Brams’ assumption that God prefers not to reveal Himself.

A second reason to doubt that Brams’ assumption that God prefers not to reveal Himself comes to light when we recognize that Brams is working with far too narrow an understanding of what it means for God to reveal Himself. Consider, again, Brams’ claim that “[God] never revealed Himself in any physical form, except possibly to Moses before he died, though he continually demonstrated His powers in other ways, notably by punishing those he considered transgressors.” Here, Brams suggests that he believes that what it is for God to reveal Himself is for God to reveal Himself in some physical form. Further, it seems that Brams takes it to be necessary for this physical form to be or resemble some bodily form. Even if we grant that for God to reveal Himself He must reveal Himself in some physical form, why must we think that this physical form must be or resemble a bodily form? To narrow our understanding of what it is for God to reveal Himself in such a way is arbitrary and unduly anthropomorphic. While God may choose to take on or resemble a bodily form, He is not essentially embodied. So, why not think that by exercising His powers in other ways—e.g., by “punishing those he considered transgressors”, etc.—God reveals Himself in the physical world and provides evidence of His existence?

As it turns out, the Old Testament is filled with stories of God revealing Himself in ways other than by physical embodiment resembling a human form. There are several passages in the very first book of the Old Testament, Genesis, in which God reveals Himself by way of dreams. In the Old Testament God also reveals Himself in other physical ways. For example, in Exodus it is written, “And the Lord went before them by day in a pillar of a cloud, to lead them the way; and by night in a pillar of fire, to give them light; to go by day and night” (Exod 13:21 RSV). Additionally, the Old Testament reminds us that God reveals Himself in the natural world: “The heavens declare the glory of God; and the firmament sheweth his handywork” (Ps 19:1 RSV). So, we can see that evidence of God revealing Himself abounds in the Old Testament.

But there’s a further point to be made: not only does Brams overlook numerous pieces of textual evidence that show that God reveals Himself in the Old Testament, Brams also overlooks that fact that according to Abrahamic religions the Old Testament is itself, in its entirety, revelation.

Finally, for Christians who accept the New Testament and recognize Jesus as the Son of God, the assumption that God prefers not to reveal Himself is even more dubious. Even allowing Brams’ unduly narrow understanding of what it is for God to reveal Himself, it is clear that—if we accept the New Testament as true—God reveals Himself by becoming man.

So, we have identified four reasons for doubting Brams’ assumption that God prefers not to reveal Himself: (1) God revealed Himself in physical form to Moses in the Old Testament; (2) there is lots of textual evidence that God revealed Himself in other ways in the Old Testament; (3) the Old Testament itself is revelatory; and, finally, (4) if we

believe in the New Testament, then we accept that God revealed Himself as Jesus. It seems, then, that if we accept the Old Testament as true, it is most reasonable to conclude that often times God in fact prefers to reveal Himself; if we accept the New Testament as true, our case for believing that God prefers to reveal Himself is even stronger.

If we correct Brams' Belief Game on the assumption that God, in fact, prefers to reveal Himself, we get the following game theory matrix:

The Belief Game on the assumption that God prefers to reveal Himself

		Man	
		Believe in God's existence	Don't believe in God's existence
God	Reveal Himself (Establish His existence)	4, 4 (Man faithful with evidence: belief in existence confirmed)	2, 1 (Man unfaithful despite evidence: nonbelief in existence disconfirmed)
	Don't reveal Himself (Don't establish His existence)	3, 2 (Man faithful without evidence: belief in existence disconfirmed)	1, 3 (Man unfaithful without evidence: nonbelief in existence confirmed)

We see that once we correct for Brams' mistaken assumption that God prefers not to reveal Himself, Brams' purported paradox goes away. For, once again, God has a dominant strategy. But this time it is to reveal Himself. And while "Man" does not have a dominant strategy (he would prefer to believe in God if He reveals Himself, and not to believe in God if He doesn't reveal Himself), if "Man" is rational he will know that God has a dominant strategy and he will thus believe in God. We will end up, then, in equilibrium in the top-left box of the matrix—(4,4)—the Pareto-optimal outcome. Bram's "paradox" is no longer.

Now on to another problem with Bram's Belief Game. As noted above, Brams assumes that the Belief Game is a game of complete information; Brams (1982 p. 123) writes, "man knows God's preferences as well as his own". However, Brams also claims that "Man" is uncertain as to whether or not God exists:

...if the human player in the Belief Game is agnostic, which is presumable the kind of player—as opposed to an avowed theist or atheist—who would take this game seriously, he will play it while still uncertain about God's existence. In other words, he will not know God's strategy choice, or even whether He exists (1982 p. 124).

Also, Brams (1982 p. 125) writes, "In a sense, a thoughtful agnostic plays the Belief Game all his life, never certain about God's strategy choice—or even that He exists—and wavering between his own". Brams is inconsistent, then, in claiming that the Belief Game is one of complete information and also that one of the players, "Man," is deeply uncertain as to whether the other player, God, even exists. "Man" can *know* God's preferences only if he knows that God exists. So, if the Belief Game is a game of complete information, then God exists, according to the Belief Game. Brams is trying to represent the rationality of belief in God from the perspective of a

player, “Man,” who is uncertain as to whether or not God exists, but the way he has set up the Belief Game makes it the case that it is necessary to assume that “Man” already knows that God exists.

### A deeper problem: God and game theory

Up to this point I have only pointed out problems that are specific to Brams’ Belief Game. However, in addition to these, I think there is a deeper problem with *any* attempt to use game theory to model interactions between God—as He is conceived of in the Abrahamic tradition—and humans. The deeper problem is, most basically, that game theoretic models involving God and human beings as agents undermine God’s omnibenevolence. The goal of this section is to explain and motivate this problem.

Recall from the introduction: game theory is, most basically, the study of multi-agent strategic interactions. Game theorists make simplifying assumptions about the preferences of agents in game theoretic models, and among these assumptions is an assumption of non-tuism.<sup>2</sup> Most basically, non-tuism requires that agents not take an interest in the preferences of those with whom they interact. An important point of clarification: saying that an agent has non-tuistic preferences is not tantamount to saying that an agent is selfish.<sup>3</sup> The assumption of non-tuism is consistent with agents having preferences that refer to other agents with whom they interact, including aspects constitutive of the well-being of other agents. What the assumption of non-tuism precludes is taking an interest in another agent’s preferences in the context of the strategic interaction, while at the bargaining table, so to speak. In game-theoretic models agents come to the bargaining table with ordered preferences and these preferences may make reference to others; but while at the bargaining table, predicted outcomes are determined based on the assumption that agents will act solely so as to maximize their ordered preferences. So non-tuism precludes one agent from acting, in the context of the strategic interaction, for the reason that his or her action will maximize another agent’s utility. Phillip Wicksteed (1910) and (1933), the early twentieth century economist widely credited for coining the term “non-tuism,” explains the concept as follows: “as far as the relation is really economic, the significance to us of what we are doing is measured not by its importance to the man for whom it is done, but by the degree to which it furthers our own ends”.<sup>4</sup> Gauthier (1986 p. 86), in *Morals By Agreement*, explains non-tuism as taking “no interest in one another’s

<sup>2</sup> This is actually a bit of an oversimplification. Christopher Morris (1991 pp. 76–95) distinguishes the concepts of mutual unconcern, egoism, asociality, and private consumerism from non-tuism. Each of these other assumptions is stronger than the assumption of non-tuism (for example, mutual unconcern—or Rawls’ mutual disinterest—requires that agents not take an interest in the preferences of others). I have cashed my argument out in terms of non-tuism to show that even on the weakest assumption, God’s omnibenevolence is undermined if we attempt to cast Him as a player in a game theoretic model.

<sup>3</sup> Another point of clarification: I take it that game theorists, in assuming that agents have non-tuistic preferences, do not take themselves to be describing the real preferences of real agents. In other words, the assumption of non-tuism is a *simplifying* assumption that need not, and does not, in fact track the behavior of real agents facing real situations.

<sup>4</sup> While Wicksteed spelled out the concept of “non-tuism” in earlier works, he first used the term “non-tuism” in *The Common Sense of Political Economy and Selected Papers and Reviews on Economic Theory*.

interests,” (p. 100) thereby ensuring that “no person gains or loses simply from the utility of others”. Brams (1982 p. 125) himself, in laying out his game-theoretic model, makes an assumption of non-tuism. He assumes that the players’ strategies are entirely independent of one another and, furthermore, that the players’ choices “are made in ignorance of each other”.

But what motivates Brams’ and others’ assumption of non-tuism? Well, on the one hand if you have a preference that the other player fares poorly while at the bargaining table, then you’ll end up at the no-agreement point; strategic interaction won’t ever get off the ground. On the other hand, if you have a preference that the other player fares well while at the bargaining table, then game theory seems otiose. What masquerades as a game theory model is really just a decision theory model. We can find even more motivation for an assumption of non-tuism if we think about both players having tuistic preferences. For example, think about what would happen if I cared about you satisfying your preferences and you cared about me satisfying mine. I want you to get what you want in a particular interaction, but what you want is for me to get what I want in that interaction. There is a kind of incoherence here. The incoherence is probably even more dramatic if I want you to satisfy your preferences, but your preferences are just that I fail to satisfy mine.<sup>5</sup> The idea of a game theoretic payoff matrix makes no sense under these conditions.

Now on to the deeper problem. A serious problem arises, I argue, when we attempt to assume that God has non-tuistic preferences. Why? Because such an assumption is inconsistent with God’s divine perfection and, more specifically, God’s omnibenevolence. According to the Abrahamic tradition, God’s omnibenevolence is in part manifested in the personal, loving relationships He fosters with humans. The assumption of non-tuism is fundamentally at odds with this conception of God. I assume that God, if He exists, cares (and must care) about our well-being in virtue of his omnibenevolence. I expect this assumption to strike most of us as relatively uncontroversial.<sup>6</sup> But it also seems to me that on any plausible theory of well-being, our well-being is closely tied up with our preferences.

Let’s consider two prevalent and promising views of well-being: (1) preferentist accounts and (2) objective list accounts. Preferentism is the view that our well-being is entirely constituted by the satisfaction of our preferences. Economists often assume preferentism. It is rather obvious, I hope, that if preferentism is true, i.e. if it’s true that the satisfaction of our preferences constitutes our well-being, it follows that God must care about our preferences if He cares about our well-being.

However, I recognize that most people who believe in God will not find preferentist accounts of well-being particularly attractive. Indeed, all theists that I know of defend some version of an objective list account of well-being. To those who endorse this sort of account, I’d cash my argument out something like this: God ordains goods. These goods are under-specified. What I mean by this is that, in some sense, these goods are multiply realizable. So, it’s permissible to choose from a variety of paths in pursuing these under-specified goods on an objective account of well-being. Often times these

<sup>5</sup> Thanks to Steve Kuhn for thinking of this possibility.

<sup>6</sup> I’m setting aside, here, debates about the plausibility of the very notion of “well-being,” especially criticisms that the notion of “well-being” is otiose (see, for example, Moore (1903) and Scanlon (1998)).

various paths for pursuing goods constitutive of our well-being are morally on par, so to speak. In these cases, what distinguishes one path from another is not some moral feature, but rather just simply our preferences. God, in virtue of His benevolence, cares about our unique, chosen paths; a loving, personal God who cares about us pursuing goods that are constitutive of our well-being cares about our preferences for how we'd like to realize goods.

Here's an example to illustrate this line of thinking: Mark Murphy (2001), in *Natural Law and Practical Rationality*, gives the following catalog of goods that are perfective of our nature and constitutive of our well-being: life, knowledge, aesthetic experience, excellence in work and play, excellence in agency, inner peace, friendship and community, religion and happiness. So, imagine a recent college graduate, Jane, who double majored in economics and philosophy. Jane, as it turns out, is equally skilled in economics and philosophy, and enjoys both disciplines. Ultimately, Jane decides that she most prefers to become a professional philosopher, and does so. If Murphy is right that excellence in work is constitutive of our well-being, and if I'm right that God must be concerned with our well-being, then God must be concerned with Jane achieving excellence in work. The way Jane prefers to achieve excellence in work, at least in part, is by pursuing a career in philosophy. My contention is that God must care about Jane's excellence in her chosen career, namely philosophy, because she prefers it. That is because part of what it is to be in a loving, personal relationship is to take the preferences of others to be reason-giving. For example, if my partner were to tell me that he prefers cauliflower to broccoli for dinner, it gives me a reason to pass by the broccoli and pick up some cauliflower while at the grocery store; if he were to tell me that he loves knitting and would like to host a meeting of the knitting club of which he is a member, it gives me a reason to open up our home (even though I, personally, find knitting to be rather tedious); I could go on. My contention is that just as my partner's preferences matter to me (and should matter to me) in virtue of the fact that we're in a loving, personal relationship, so too our preferences matter to God in virtue of the facts that He is omnibenevolent and that we're in a loving, personal relationship with Him. Of course there are limits to the reason-giving force of one's preferences. If I were to prefer engaging in self-mutilating behaviors, I doubt that this preference would be reason-giving for my partner; he would not, for example, have reason to stock up on scissors. Similarly, our preferences are not always reason-giving for God. Only preferences that fall in the penumbra of goods constitutive of our well-being seem to count. But this does not undermine my main point because, if you recall, an assumption of non-tuism requires that agents in a game theoretic model take *no interest whatsoever* in one another's interests. I conclude, then, that to assume that God has non-tuistic preferences, i.e., that He does not care about our preferences, even on an objective list account of well-being, is to abandon the assumption that God is omnibenevolent.

## Conclusion

Brams, in his "Belief Game," mischaracterizes God's preferences and he fails to adequately address the fact that from the perspective of one of the players ("Man"), it

is uncertain whether or not the other player (God) even exists. But I hope to have shown at this point that these criticisms are in some sense moot, since there is a deeper problem that undermines the very attempt to treat God as an agent in a game-theoretic model with humans. I have argued that since God, in virtue of His benevolence, cares about our well-being, He must care about our preferences. If that is true, I have argued that an assumption of non-tuism cannot apply to God, or else the assumption that God is omnibenevolent will be undermined.

**Acknowledgments** For helpful comments and discussions, I would like to thank Nick Casalbone, Steve Kuhn, Anne Langhorne, Mark Murphy, Travis Rieder, and the audience of the 2011 South Carolina Society for Philosophy/ North Carolina Philosophical Society annual meeting.

## References

- Binmore, K. (1994). *Playing fair: Game theory and the social contract*. Cambridge, MA: MIT Press.
- Binmore, K. (2008). *Natural justice*. Oxford: Oxford University Press.
- Brams, S. (1982). Belief in God: A game-theoretic paradox. *International Journal of Philosophy of Religion*, 13(3), 121–129.
- Brams, S. (1983). *Superior beings: If they exist, how would we know?*. New York: Springer.
- Brams, S. (2007). *Superior beings, IF they exist, how would we know?* (2nd ed.). New York: Springer.
- Kuhn, S. (2009). Prisoner's dilemma. In E. N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, Online Edn. <http://plato.stanford.edu/entries/prisoner-dilemma/>. Accessed 24 July 2012
- Lewis, D. (1969). *Convention: A philosophical study*. Cambridge: Harvard University Press.
- Gauthier, D. (1986). *Morals by agreement*. Oxford: Oxford University Press.
- Moore, G. E. (1903). *Principia ethica*. Cambridge: Cambridge University Press.
- Morris, C. (1991). Moral standing and rational choice contractarianism. In P. Vallentyne (Ed.), *Contractarianism and rational choice: essays on David Gauthier's morals by agreement* (pp. 76–95). Cambridge: Cambridge University Press.
- Murphy, M. (2001). *Natural law and practical rationality*. Cambridge: Cambridge University Press.
- Scanlon, T. (1998). *What do we owe to each other?*. Harvard: Belknap Press.
- Wicksteed, P. (1910). *The common sense of political economy: Including a study of the human basis of economic law*. London: Macmillan.
- Wicksteed, P. (1933). *The common sense of political economy and selected papers and reviews on economic theory*. New York: Routledge and Kegan.