Social Evil*

Ted Poston University of South Alabama

Draft. Please do not cite without permission.

Word Count: 8210

Abstract: Social evil is any pain or suffering brought about by game-theoretic interactions of many individuals. This paper introduces and discusses the problem of social evil. I begin by focusing on social evil brought about by game-theoretic interactions of rational, moral individuals. The problem social evil poses for theism is distinct from problems posed by natural and moral evils. Social evil is not a natural evil because it is brought about by the choices of individuals. But social evil is not a form of moral evil because each individual actor does not misuse his freewill. Traditional defenses for natural and moral evil fall short in addressing the problem of social evil. The final section of this paper discusses social evil and virtue. I argue that social evil can arise even where virtue is lacking. Further, I explore the possibility of an Edwardsian defense of social evil that stresses the high demands of true virtue. The conclusion of this paper is that social evil is problematic and provides a new ground for exploring the conceptual resources of theism.

Discussion on the problem of evil assumes that there are two classes of evils: natural evil and moral evil.¹ Richard Swinburne divides "the world's evils in the traditional way into moral evils and natural evils."² Swinburne characterizes moral evils as "those brought about by human intentional choice, or knowingly allowed to occur by humans, together with the evils of their intentional bad actions or negligence."³ Natural evils, according to Swinburne, are "all other evils, such as bad desires that we cannot help, disease, and accidents."⁴ Alvin Plantinga provides a similar division of evils. He writes, "In addition to "natural" evils such as earthquakes, tidal

[·] Thanks to Trent Dougherty, Allan Hillman, Jonathan Jacobs, Clayton Littlejohn, Kevin Meeker, Alvin Plantinga, Alexander Pruss, Richard Swinburne, Peter Vallentyne, William Wainwright, Paul Weirich, and Heath White for excellent comments on earlier drafts.

¹ The distinction between two classes of evil goes back to, at least, Augustine. On natural evil see Augustine's *De Ordine* and on moral evil see *Confessions* and *On Free Will*.

² Swinburne, (2004) p. 236; cf. van Inwagen (2006, p. 9).

³ Ibid., 236.

⁴ Ibid.

waves, and virulent diseases there are evils that result from human stupidity, arrogance, and cruelty." William Rowe's famous cases of Bambi and Sue provide representative cases of each kind of evil. The Bambi case involves the prolonged suffering and death of a fawn. And the Sue case relates the horrible beating, rape, and murder of a five-year old girl. Both these forms of evil exercise the conceptual resources of theism to explain how an omnipotent, omniscient, and morally perfect being might allow these evils. But there is another form of evil that has not received the attention of philosophers working on the problem of evil. This form of evil I call social evil. For the purposes of this paper I assume that evil is any form of pain or suffering and that types of evils are individuated by the processes that bring them about. Thus, moral evil is a form of pain and suffering that is brought about by the direct agency of a person, and natural evil is a form of pain and suffering brought via the operation of laws of nature on matter. Social evil is a form of pain and suffering that results from the game-theoretic interactions of many individuals. Until the last section of the paper I focus on the most problematic form of social evil, social evil that occurs because of the game-theoretic interactions of rational, moral individuals. It is widely acknowledged in the literature on game theory that rational, wellintentioned agents can bring about horrible social outcomes. Russell Hardin (1995) provides an extreme example of this by offering a game-theoretic analysis of violent group conflict. Once we see the game-theoretic machinery in play it is hard to resist the thought that much evil in our world is the unintended result of collective agency among individually rational participants. My goal in this paper is to introduce and explain social evil and the problems it poses for traditional defenses. Furthermore, I will discuss the prospects for a theistic treatment of social evil. I intend

_

⁵ Plantinga, (1974), 8.

⁶ The original "Bambi" case comes from Rowe (1979) and the "Sue" case is in Rowe (1988). William Alston (1991) dubbed these cases "Bambi" and "Sue" and, following Alston, the terminology has stuck.

my discussion to be a starting point for a deeper reflection on the nature of social evil and the theoretical lessons it has for theism.⁷

I. What is social evil?

I begin with a clear case of pain and suffering that results from the game-theoretic interactions of rational, well-intentioned individuals. Suppose you are a resident of suburban Chicago and the Chicago area is facing a serious water shortage. The reservoirs are running dry; the Illinois and Fox rivers are near record lows; even Wisconsin's bountiful lakes and rivers are ominously diminished. Apart from a significant decrease in overall water consumption the Chicago area will run out of an adequate water supply. City planners foresee the possibility of severely restricting residential water supply. However, if most everyone significantly decreases their water consumption—stop watering lawns, washing cars, and letting the tap run any longer than absolutely necessary—the Chicago area will manage until the fall rains come. Obviously, it is in the best interest of all that most everyone follows this advice. But, this represents a considerable cost to each person. If, for example, you decrease your water usage your carefully cultivated garden and fruit trees will wilt and die. This is a hefty burden to pay. However, if no one decreases his or her water usage each will pay an even greater cost. Yet, you realize that if most everyone decreases his or her water consumption then you may continue your normal usage without any ill consequence. Moreover, because the benefit of decreased water usage requires a very large number of participants—well over a million homeowners—your own contribution

_

⁷ The concept of social evil has wider applications to ethical theory. Reflection on social evil shows that, in at least some cases, an individual's immoral actions are not the result of a bad will but rather the result of an individual being caught in a tragic game-theoretic scenario. One might think that social evil suggests that there are far fewer moral monsters that we might otherwise think.

does not (and will not) affect whether or not the benefit is realized.⁸ In this case let us assume that you are rational and blameless. You do not suffer from a failure to realize that you are in this kind of circumstance. Moreover, you do not suffer from a moral fault; you do not want to harm anyone and you do not want your action to bring about a worse circumstance. But given the logic of the situation you realize that whatever you do it will not affect the social outcome. Taking into consideration all the relevant factors, the best option for you is not to conserve water. But all other individual residents of Chicago face an identical situation in which the best option for each individual is not to conserve water. Thus, if everyone is rational and blameless each individual will play his best strategy and the collective result will be unintended disaster. The disaster that results is a social evil.

This is a standard form of a multiplayer prisoner's dilemma, also known as "the tragedy of the commons." Cases of this sort are widespread. Achieving adequate healthcare, decent education, effective inoculations, safe freeways, fertile fishing waters, and pristine national parks all require the cooperation of a sufficiently large group of individuals. In these cases the goods achieved and the evils avoided require solving a multiplayer prisoner's dilemma. Because each individual's dominant strategy is to defect, governments aim to avoid the destructive logic of these games by, among other things, imposing significant penalties for defecting. While this is an imminent practical problem, there is a theoretical problem for theism that has not been addressed in the literature to date.

See G. Hardin (1968).

⁸ I assume that this is a simultaneous move game. A simultaneous move game requires that you (and everyone else) act in ignorance of what the other people do.

To guide our discussion I will present a standard two-person prisoner's dilemma case.¹⁰ We can consider a two-person dilemma as the smallest case of a social evil, a case in which an unintended and worse outcome results for each player enacting his best strategy. Consider the following game.

Two-Person Prisoner's Dilemma

		Player II	
		Defect	Cooperate
Player I	Defect	2,2	4,1
	Cooperate	1,4	3,3

In this game Player I has two options: either defect (maintain his water usage) or cooperate (decrease his water usage). Player I achieves his *best* result in which he defects given Player II's cooperation, and he achieves his *worst* result if he cooperates given Player II's defection. The numbers 4-1 represent an *ordinal* ranking of the players' preferences. An outcome 4 is preferred to 3 which is preferred to 2 which is preferred to 1. Ordinal rankings reflect a series of ordered preferences and they do not reflect the *strength* of preference. With an ordinal ranking you cannot infer that (e.g.) Player I's highest ranked option (e.g. a 4) is twice as desirable as his second ranked option (e.g. a 2). Player II's rankings are symmetrical with Player I's. Thus, for Player II his highest ranked option is where he maintains his current water usage and Player I decreases water usage.

What should each Player do? Given Player I's situation he can always improve his outcome if he defects. To see this suppose Player 1 cooperates. Then if Player II defects, Player

¹⁰ I give a multiplayer prisoner's dilemma below in section IV. At this point in our discussion, we will gain no appreciable advantage by using the more complex multiplayer PD.

I would do better by *defecting*. He would move from a payoff of 1 to a payoff of 2. If Player II cooperates, then, just as in the other situation, Player I would do better by *defecting*. He would move from a payoff of 3 to a payoff of 4. In this situation Player I's option of defecting *strictly dominates* cooperating because whatever Player II does, Player I can do better by defecting. Player II, of course, faces the exact similar reasoning and so Player II should defect. Thus, these two rational, well-intentioned players achieve an outcome neither wants, viz., mutual defection. Both players would prefer mutual cooperation to mutual defection but the logic of the game lands rational, well-intentioned players in mutual defection.

The water shortage case illustrates the painful logic of group action. Each participant realizes the structure of the game they are in and each participant is blameless. But *given* the structure of the game it is rational for each player to defect. Note that because of the nature of group action, each player does not directly bring about a bad state of affairs. In large multiplayer prisoner dilemmas the benefit (or detriment) is achieved regardless of what any one individual does. In more realistic games—like the water shortage case—the benefit (or determinate) is insensitive to at least one tenth of a percent of the total number of players. Thus if a multiplayer prisoner's dilemma is realized with 10⁶ players then the social outcome will obtain regardless of what 10³ players do. So, if 1,000 players don't make a difference to this game your decision surely does not make a difference. Consequently, each individual player's choice does not affect the outcome of the game.

I assume that this discussion is sufficient to motivate the thought that social evils are distinctive because they occur as the result of game-theoretic interactions among many individuals and *not* on account of some individual's choice (as in the case of moral evil) or some

natural process (as in the case of natural evil). Let us now examine challenges to the idea that *social evil* is a distinctive kind of evil.

II. The distinctness of social evil

One objection to the claim that social evil is fundamentally different in kind from natural and moral evil is that social evil requires pain and suffering, and pain and suffering are not social conditions. Since pain and suffering are not social conditions the problem of social evil is just the problem of pain and suffering. Traditional defenses address the problem of pain and suffering so we can just apply those to the problem of social evil. There's no special problem here.

This objection is right that social evil requires pain and suffering. But the objection fails because it does not recognize distinct problems with the *route* by which pain and suffering occur. Different defenses are deployed to handle different kinds of evil, and kinds of evil are individuated by the processes that bring them about. It's one thing for pain and suffering to occur by natural processes (why would a perfect being allow those kind of natural processes to occur?), and it's an entirely different thing for pain and suffering to occur by the direct result of human choice (why would a perfect being allow individuals that can directly harm each other?). If the present objection worked against social evil it would work equally as well on the distinction between natural and moral evil; for as the objection goes there would just be one problem here, the problem of pain and suffering. But clearly the problems posed by natural and moral evil are different. And since social evil is neither brought about by natural law nor by direct choice, the problem it poses for theism is different as well.

One might object to the above claim that moral evil and natural evil pose different problems by arguing that since the free will defense shows at least the possibility that natural evil is a result of moral evil then there's no real distinction between moral and natural evil. There's just one problem, the problem of pain and suffering. What shall we say to this objection? We might concede that *if* one doesn't think that natural evil poses any special problem over the problem of moral evil then one would understandably think that further distinctions among types of evil were unnecessary. But we can ask how plausible is it that natural evil doesn't pose any special problem. Natural evils occur because of states of nature and laws which humans lack any appreciable control over. Why would a perfect being bring about a world like that? The present objection notes that the free will defense can appeal to some *past* misuse of free will. But I don't find this move plausible at all. It seems a very harsh punishment for past misuse of free will that it's present *proper* use results in so much pain and suffering. We will return to this point in the discussion of the free will defense in section III below.

Another objection to the distinctness of social evil is that social evil assumes natural evil and so an adequate defense for natural evil will carry over to social evil. The water shortage example is a case of scarce resources. The present objection claims that scarcity of resources is a natural evil and so, fundamentally, social evil is a case of natural evil. But what reason do we have for accepting the claim that scarcity of resources is a natural evil? Natural evil is a form of pain and suffering that results from natural processes. The standard cases of these evils are instances of pain and suffering brought about by hurricanes, earthquakes, and lightning strikes. But why should we think that the *mere* scarcity of a resource (i.e., a limited amount of a resource) brings about pain and suffering? Even in normal circumstances water is a scarce resource. Everyone cannot run every faucet in his home all day long. But we can effectively

manage the situation so that no one is adversely affected. It is true that some cases of scarcity bring about pain and suffering, but it is often the case that the *magnitude* of the pain and suffering depends on the collective response of society. In the water shortage case the Chicago area is facing a severe drought. But the magnitude of the drought's effects depends on the collective response of Chicago area residents. According to the present objection the residents of Chicago are faced with the prospects of *costs* by the drought and that's a form of pain and suffering. That's right as far as it goes, but it brings us back to the first objection that social evil is really natural evil because it requires pain and suffering. But as I argued that objection falls because moral evil requires pain and suffering.

A third objection to the distinctness of social evil is that it is a subtle form of natural evil, specifically human stupidity. If an individual defects fully realizing that everyone else faces the exact same reasoning then the person is stupid. But why should we think that the individual defector is stupid? One reason is that the individual does not realize that his defection brings about a worse state of affairs, viz., 1 less the number of confederates. But this reason rests on a serious misunderstanding of the logic of multiplayer prisoner's dilemmas. The individual defector does not *bring about* a worse state of affairs by bringing it about that there is 1 less number of confederates. This is because the social benefit will be achieved (or not) regardless of what an individual person does. That is, for any n, n±1 confederates do not change the outcome of the game. As remarked above, in very large multiple player games this holds true for upwards of n±10,000. To suppose that an individual's choice affects the outcome of a multiple player game is a gross misunderstanding of the logic of such games.

The principle that for any n, $n\pm 1$ confederates do not change the outcome of the game is true of cases in which the threshold of confederates required to achieve the good and avoid the

evil is vague. In cases of vagueness one cannot apply the principle iteratively without at some point losing knowledge about whether the threshold is met. In the Chicago water shortage case let us suppose that three million confederates will achieve the social good but one hundred thousand will not. One can gradually diminish the numbers from three million participants to one hundred thousand but at some point in this series it becomes vague whether that number of participants will achieve the good. Still, it remains true that small changes, i.e., plus or minus 1, will not affect the outcome to be achieved.

A fourth objection to the special nature of social evils is that it is a form of moral evil. Specifically, an individual defector is morally to blame for defecting. While there *might* be something to this claim (though see sections III & IV below), it will require serious argumentation that goes beyond standard characterizations of moral evil. On the traditional view, moral evil is a form of pain and suffering that directly results from the agency of another person. Typically, this involves the *misuse* of freewill. 11 Clearly, an individual defector does not bring about any (relevant) pain or suffering. The individual defector is not an outcome cause of the effect. The effect will occur regardless of what the individual does.

One response to this line of argument is that it mistakes the logic of rational choice with the logic of moral choice. A prisoner's dilemma represents the interaction of preferences between various agents. If one agent has a dominant strategy one cannot infer from dominance alone that an agent is morally in the clear to enact that strategy. In particular, to be morally just an agent's action must be universalizable and clearly in a multiplayer prisoner's dilemma an individual's strategy to defect is not universalizable. Thus the option of defecting in a multiplayer prisoner's dilemma is immoral.

¹¹ Though see Robert Adams "Involuntary Sin."

This response appeals to those with Kantian moral sensibilities. But the problem with this reply is that universalizability is not a necessary condition for moral permissibility. Suppose you are considering what kind of life you should live. Should you be a medical doctor by which you may save many lives or should you develop your talent as an artist through which you may enrich the lives of many? Each decision is made against the backdrop of a diverse population with different aims and goals. A life as a medical doctor assumes that many people are not doctors and a life as an artist assumes that not everyone is an artist. One cannot simply universalize the maxim *I will become a doctor to save the lives of many* because if everyone became a doctor to save the lives of many that specialization would collapse.

This objection to the requirement of universalization shows that if the requirement is to be plausible at all it must allow that diversity is represented within the universalization test. To test the moral appropriateness of one's choice to become a doctor one must universalize the maxim assuming that other people choose different lives to lead. But once we allow for diversity in the universalization test, the original objection that defecting is immoral falls apart. For one can represent in the universalization test that one's action will not influence what the group will do. Even more strongly, one can represent that one knows this to be the case. And once this is represented then one's maxim passes the universalization test.

Some people persist in thinking that there is an attenuated sense of 'bring about' on which the individual brings about a worse state of affairs. But, to repeat a point made above, this is a failure to realize the logic of collective action. A related, but confused, objection is that an individual's act of defecting represents a cost itself, and so the individual is not faultless for defecting. This objection is confused because it changes the logic of the game. A prisoner's dilemma is a mathematical object that represents moves and preferences. If one thinks that

defection itself is a cost then that should be reflected in the system of preferences. Either the new game will be a prisoner's dilemma or not. If it is then the problem of social evil will arise and if not then the problem may not arise.¹²

III. Standard Defenses

In this section I examine standard defenses to determine how they might handle social evil. I argue that social evil is difficult to incorporate within standard theistic responses to evil. I consider four defenses: the value of natural laws, a soul-making defense, the free will defense, and skeptical theism.

A. The value of natural laws

Richard Swinburne argues that the problem of natural evil is lessened by the value of natural laws.¹³ A universe with recognizable regularities allows persons to successfully predict the consequences of their actions. If I want to help you by offering you nourishment, it is valuable for me to know that if I offer you bread it won't kill you. Similarly, if I want to study the possibility of cold fusion, it will help to know that there are reliable ways to bring about freezing temperatures. Because a world with laws is beneficial in this respect, it may be a consequence that some natural evils occur. For instance, a world in which substances have stable properties may imply that in some cases the stable properties of substances harm individuals (e.g., a tree falls and breaks Joe's leg). Or, a world in which environments have

12

¹² I leave open the question of whether prisoner dilemmas are a necessary condition for social evil

¹³ See Swinburne (2004), 245ff.

stable properties may imply that at boundary zones (e.g., a cold front meets a warm front) violent storms occur.

Regardless of what one makes of Swinburne's claim, it is clear that this will not help with social evils. Social evils arise because of the collective agency of rational, well-intentioned individuals. This requires that individuals can effectively reason about their options, which in turn requires observable regularities in the world, but that's true of moral evil as well. The crucial difference between social evil and natural evil is that social evil occurs because of human agency. Thus, if the appeal to natural laws cannot cope with moral evil it cannot tackle social evil.

B. Soul-Making

A different defense appeals to the value of certain kinds of character traits—patience, fortitude, courage, and compassion—and then argues that these valuable traits require evils. Compassion requires the possibility of coldness; courage requires the possibility of cowardice. The soul-making defense stresses that a world that contains the great goods of character must also contain the great evils of character as well. John Hick provides a nice summary of this defense:

The value-judgement that is implicitly being invoked here is that one who has attained to goodness by meeting and eventually mastering temptation, and thus by rightly making responsible choices in concrete situations, is good in a richer and more valuable sense than would be one created ab initio in a state either of innocence or of virtue. In the former case, which is that of the actual moral achievements of mankind, the individual's goodness has within it the strength of temptations overcome, a stability based upon an accumulation of right choices, and a positive and responsible character that comes from the investment of costly personal effort.¹⁴

_

¹⁴ Hick (1977), 255-6.

Does the soul-making defense offer promise for handling social evil? An initial hurdle for the soul-making defense is that social evils are not cases in which an individual *brings about* a worse outcome. There is no causal connection between an individual's choice and the outcome that is realized by the group. This verdict is especially clear if we apply the distinction between outcome causation and aspect causation. An outcome cause is a difference maker to whether or not the effect occurs. An aspect cause is a difference maker to the effect occurring as *it does*. In cases of social evil each individual's action is not a difference-maker to the outcome being achieved or not. And it is only in the most attenuated sense that an individual's action is an aspect cause of the effect. So, there is very little room for an individual in such a case to achieve valuable character traits. Perhaps, though, a defender of this move will stress the value of solidarity or the value of self-inflicted loss even when those choices do not have any larger social consequences. But it's doubtful whether this move could be sustained because there is no causal connection between what an agent does and what social outcome is achieved. If an individual's choices have no effects there is little of moral value in them.

There is another problem with the soul-making defense apropos social evil. The soul-making defense is deeply individualistic. It focuses on the value of an individual's own character traits. Thus the soul-making defense doesn't have the conceptual resources to explore the value of traits of societies, traits that do not reduce to traits of individuals. Exploring the value of traits of societies seems to be a ripe area for reflection on social evil. Perhaps further investigation on the value of certain types of society will uncover reasons God has to permit social evils. But unfortunately our current discussion of evil is largely informed by individualistic assumptions to the extent that we lack the language and perspective to discuss societal goods and evils.

C. The Free Will Defense

The free will defense focuses on the immense value of free will and its irresponsible use in generating horrendous evils. A free will defense can be used to generate a story for why God might permit moral evil. A crucial part of the free will defense is that persons are responsible for actions that they freely bring about. But as we've seen above social evil lacks the feature that an individual brings about the evil. Social evils result from the collective agency of individuals, not from any particular individual's choice. Another way to see this is that social evils do not result from any specific individual *misusing* her free will. If an individual in a multiplayer prisoner's dilemma enacts her best option she is not thereby stupid or wicked. As explained above, individuals in prisoner dilemmas need not be either irrational or mean. In the water shortage case I stipulated that each individual was rational and blameless. They each know the situation and they don't want to harm anyone by their actions, but given the situation each person's best play is defection. In this case defection is not morally unjust. Thus, the value of an individual's free will doesn't account for social evils because there's no sense to be made of a culpable misuse of free will in this case.

Alvin Plantinga, pursuing a line by Augustine, considers the possibility that natural evil is the result of the free action of non-human spirits. ¹⁵ It's possible that, as Plantinga says, "Satan rebelled against God and has since been wreaking wherever havoc he can. The result is natural evil." Given this possibility, the free will defense shows that there is no logical inconsistency between theism and natural evil. In the past some free agent performed a culpable act and one of the enduring results of that act is natural evil. A similar move may be considered in connection with social evil. It is possible that the occurrence of the destructive logic of a multiplayer

¹⁵ Plantinga (1974), 192.

¹⁶ Ibid

prisoner's dilemma is due to the past misuse of free will. Following Plantinga's Augustinian lead, it's possible that "Satan rebelled against God and has since been wreaking wherever havoc he can. The result is social evil."

How plausible is this response? A crucial difference between natural and social evil is that natural evil is not the product of human agency whether individual or collective, but social evil is the result of collective human agency. Thus it's hard to see how the destructive logic of collective action is the product of some non-human agency and very easy to see how that logic is the collective result of the human choice. Consequently the Augustinian/Plantingian appeal to non-human agency does not provide a viable free will defense for social evil.

A different possibility for the free will defense is that prior to human sin God prevented the occurrence of multiplayer prisoner's dilemmas. ¹⁷ Perhaps, prior to the misuse of free will, there were no scarce resources, or individual's incentives and options never realized the destructive logic of certain games. But upon the misuse of free will God removed these protections. So, even though some social evil arises from no present wrongdoing, the occurrence of multiplayer prisoner's dilemma is the tragic result of previous sin. This strikes me as the best option for a free-will defense. But it stands in need of much further elaboration and defense. One issue is that even if it's true that God prevented the occurrence of multiplayer prisoner's dilemmas prior to human sin, we need some additional account of what might justify God in permitting this destructive logic to occur now. As I've argued horrible social outcomes can result from the collective choices of rational, well-intentioned individuals, individuals, we may say, that are using their free will responsibly. It seems a harsh punishment for previous human

¹⁷ Thanks to Alex Pruss for this suggestion.

sin that the proper use of free will would now result in evil. So the possibility that multiplayer prisoner's dilemmas arise because of past human sin does not exonerate the free will defense.

D. Skeptical Theism

Skeptical theism is the position that for all we know God has reasons for permitting evil that exceed our capacity for discovering. God is a perfect being, complete in every perfection all-knowing, all-powerful, and supremely morally perfect. We are far from achieving any such exalted status and it stands to reason that the limits of God's epistemic and moral capacities are far beyond our imagining. Think of a novice chess player observing a match between Kasporov and Spassky. At a crucial point in the match Kasporov moves his bishop to E7 capturing Spassky's E7 pawn. Why, thinks the novice, did Kasporov do this when it will result in the immediate capture of the bishop and also deploy Spassky's Queen? The novice simply can't see a reason. Should the novice conclude that there is no reason? Of course not. It would be silly to suppose that the novice is in a position to discern every reason Kasporov might have for making this move. Similarly, according to the skeptical theist, our inability to think of a reason God might have for allowing horrendous evils doesn't justify us in believing there are no evils. Applied to social evil, if we find that the traditional defenses fail we shouldn't conclude that God doesn't exist. Rather we should think that for all we know God might have a reason for permitting social evils.

There is certainly something right with the skeptical theist's position that we should be leery of our ability to limn the limits of a perfect being's reasons. Yet how far should we take this humility? On one interpretation of skeptical theism we should be so humble as to think that the existence of evil is evidentially irrelevant to the existence of God. That is, we should think

that the probability that there's a God given evil is the same as the probability that there's a God. Where 'Pr(p)' represents the probability that p is true, 'G' represents that *there is a God*, and *e* is some evil, the present claim is that, for any e, Pr(G|e)=Pr(G). This claim is implausible. Where 'e*' is the proposition that *10 trillion sentient creatures suffer lives of endless torment*, it is overwhelmingly plausible that Pr(G|e*)<Pr(G). This is <u>not</u> to say that one may learn something else, q, on which Pr(G|e*&q)=Pr(G); rather it is the jejune observation that a morally perfect being does not desire the endless torment of sentient creatures. Consequently, the humility that skeptical theism recommends must not yield the result that evil is completely irrelevant to the existence of God. Furthermore, since skeptical theism offers us no prospective on why God might permit social evil and, as we have seen, the traditional defenses fail, we may rightly think that the existence of social evil lowers the probability that there's a God.

IV. Social Evil and Virtue

To this point I've argued for the claims that social evil is distinct from natural evil and moral evil and also that standard theistic defenses do not handle social evil. In this section I have two goals. First, I weaken the assumptions I made regarding social evil. I have argued above that social evil arises by the collective action of rational, well-intentioned individuals in multiplayer prisoner's dilemmas. This is but one form of social evil. Below I substantiate the claim that social evil can arise from the collective action of rational but *immoral* agents in multiplayer prisoner's dilemmas. This can occur when game-theoretic scenarios make it very difficult to avoid doing the wrong thing. In these kinds of cases an individual can be strongly tempted to perform an act that, apart from the wider game-theoretic scenario, would be a small

fault. But because of the destructive game these individuals find themselves in, small sins add up to horrendous evils.

The second goal of this section is to explore the possibility of an Edwardsian response to social evil. Jonathan Edwards claims that true virtue consists in love for being in general. A truly virtuous person does not love only a limited system—including himself—but loves every being and seeks the good of all. A truly virtuous person will perform the act that is best for all even if that act requires shunning his or her own private good. Since God has a reason to bring about truly virtuous people, one might reasonably think that the value of true virtue can provide God with a reason for permitting the game-theoretic machinery that produces social evil. The ensuing discussion will show that the Edwardsian defense solves a two-person prisoner's dilemma but it does not solve a multiplayer prisoner's dilemma.

A. Immoral man and very immoral society¹⁸

Our initial examination of social evil investigated instances of pain and suffering brought about by the collective action of rational and well-intentioned agents. The key feature that distinguishes social evil from natural and moral evil is that the pain and suffering that occurs does not arise from the direct choice of any individual nor from the result of natural processes. It is the cumulative effect of very many choices within the multiplayer prisoner's dilemma that produces pain and suffering. Where each individual enacts their best strategy, the structure of the game the actors are in produces horrendous outcomes.

This description of social evil does not assume that the individuals are blameless. Social evil requires only that game-theoretic machinery produce an amount of pain and suffering that is

¹⁸ I'm told by a reliable source that Reinhold Niebuhr later thought that this would be a more apt title for his famous book *Moral Man and Immoral Society*.

disproportionate to the individual choices in the game. If a game is realized in which each individual is more self-interested than not and yet each chooses their best strategy they can produce a horrendous outcome that is not mirrored in their own actions.

Examples of this kind of social evil are all too common. Russell Hardin in his book *One* for All provides a game-theoretic account of several well-known conflicts: Yugoslavia, Northern Ireland, Somalia, Rwanda, and the nationalist movement in Quebec. 19 Each conflict shares a basic structure. The roots of the conflict lay in norms of group identifications, norms that govern group inclusion and exclusion. These norms are important because there are significant benefits to be realized by being a member of the group. In 17th century England, for instance, identification in the aristocracy carried with it promise of better jobs, education, and social mobility. But identification requires norms of exclusion, norms that distinguish one group from the other. In 17th century England one powerful norm of exclusion was the duel; an aristocrat was required on pain of loss of honor to risk injury and death at the smallest of offenses. Once the norms of identification are in place, each individual's incentive structure favors identifying with a group. This leads to competition between groups for access to benefits. Often this competition leads to conflict and even violence. But once violence is realized, this provides the tipping point of group conflict in which escalating reprisals are rewarded. In many cases, this leads to open war. Thus there is this basic structure: norms of group identification and consequent benefits from identification; competition; conflict; tipping point; system in which reprisals are rewarded; and escalating violence.

_

¹⁹ See Hardin (1995), 155-177.

The Croatian War of Independence from 1991-1995 provides a case in point of the dreadful game-theoretic scenarios that produce horrendous suffering. Croatia was one of six republics in Yugoslavia, and in 1990, faced with the prospects of Serbian dominance in Yugoslav politics, Croatian leaders decided to move for independence. Unfortunately, Croatian Serbs, in addition to being a substantial minority in Croatia, constituted a sizable portion of the Croatian military and police force. When the Croatians moved for independence they faced the prospects of an internal military revolt from the Croatian Serbs. These Croatian Serbs found themselves in a game in which they could do nothing to ensure loyalty to the new Croatian government. In lieu of this situation Croatian leaders preempted the possibility of a fifth column by dismissing Croatian Serbs from military and police positions. This move, though favored by the game the Croatians found themselves in, created the tipping point for internal Serbian dissent and led to revolt. The ensuing series of moves and countermoves resulted in the devastating Croatian civil war.²¹

This story highlights the awful logic of some games. Further it doesn't assume that each actor is morally in the clear. Arguably, events around the tipping point are replete with bad intentions. But the bad intentions and wrongful actions do not themselves explain the descent into open war and needless carnage that followed. Thus, we can clearly see the effects of gametheoretic scenarios in producing great pain and suffering.

In the discussion of the first several sections I focused on pain and suffering that resulted from the collective action of rational, moral agents. Now we have seen that social evil can result anytime the game-theoretic scenario *magnifies* the consequences of individual action.

_

²⁰ This paragraph is a condensed summary of Hardin's analysis (see pp. 156-163).

²¹ Hardin also notes that the Croatian move from independence made group cooperation in Bosnia "virtually impossible" (p. 159).

Sometimes pain and suffering result from people even though no one has done anything amiss. This is the first kind of social evil. But in other cases social evil results from people when they have done something wrong but the social situation they are in magnifies the wrong actions. This is the second form of social evil.

B. An Edwardsian Defense of Social Evil

Jonathan Edwards in his book *The Nature of True Virtue* argues that true virtue is love for being in general. The truly virtuous person does not seek to benefit only a limited group of participants but rather seeks the good of every being. Edwards' discussion on true virtue includes a penetrating exposition on the consequences of self-love. Edwards offers what is, in effect, a game-theoretic explanation of apparent moral behavior. Edwards observes that a person's self-interest will motivate them to act in apparently non-self-interested ways. Consider giving to the needy. Edwards argues that this act may often arise from a limited benevolence. One can be concerned only with the needy in one's own town, or one can give to the needy to identify with the "moral" crowd and then reap the benefits of being a member in good standing in that crowd. Edwards underscores that any such act does not arise from true virtue. True virtue consists in love for being in general.

Given Edwards emphasis on the high demands of true virtue together with the claim that a perfect being desires to bring about truly virtuous individuals, there may be the makings here for an Edwardsian defense of social evil. A case can be made for an Edwardsian defense, but, if my analysis is correct, an Edwardsian defense answers only social evil arising from a two-person prisoner's dilemma; it does not extend to a general solution to the multiplayer prisoner's

dilemma. Even though the Edwardsian defense is unsuccessful, it holds important lessons for reflection on social evil.

I begin by examining an Edwardsian solution of a two-person prisoner's dilemma. I will assume that the truly virtuous person will favor the option with highest social utility. In a two-person prisoner's dilemma I will take social utility to be the sum of the individual preferences. In the matrix below I put the social utility in brackets. Consider the following game.

Two-person PD with social utility

		Player II	
		Defect	Cooperate
Player I	Defect	2,2 (4)	4,1 (5)
	Cooperate	1,4 (5)	3,3 (6)

By inspection we see that the option with the highest social utility is the one in which everyone cooperates. This option achieves 6 units of social utility. What does this revised game look like for the truly virtuous person? Let us assume that Player I is truly virtuous. Player I faces the following choice: defect or cooperate. If Player I defects then he realizes either the social utility of 4 (in which Player II defects) or the social utility of 5 (in which Player II cooperates). In either case, Player I would realize more social utility by cooperating. Thus the action of cooperating dominates. The truly virtuous person solves the two-person prisoner's dilemma.

The stress Jonathan Edwards places on true virtue requires of a person in a two-person prisoner's dilemma to choose an option that leaves them open to exploitation. It is interesting that Jesus makes similar remarks in the Sermon on the Mount. In Matthew 5:38-42 Jesus addresses the common tit-for-tat strategy. The tit-for-tat strategy enjoins a player facing an

iterative prisoner's dilemma situation to begin by cooperating and to depart from cooperating only in the face of defection (i.e., retaliate). The tit-for-tat strategy is a highly successful strategy for a person to play in an iterative prisoner's dilemma game. Jesus calls his followers not to play tit-for-tat but rather to sacrifice their own benefit for the benefit of others. In Matthew 5:38-42 Jesus says,

You have heard that it was said, 'Eye for eye, and tooth for tooth.' But I tell you, do not resist an evil person. If anyone slaps you on the right cheek, turn to them the other cheek also. And if anyone wants to sue you and take your shirt, hand over your coat as well. If anyone forces you to go one mile, go with them two miles. Give to the one who asks you, and do not turn away from the one who wants to borrow from you.

Later, in Matthew 5:43-36, Jesus address the obvious objection that this will lead to exploitation,

You have heard that it was said, 'Love your neighbor and hate your enemy.' But I tell you: Love your enemies and pray for those who persecute you, that you may be sons of your Father in heaven. He causes his sun to rise on the evil and the good, and sends rain on the righteous and the unrighteous. If you love those who love you, what reward will you get?

From these passages we see that Jesus forbids his followers from playing tit-for-tat and instead enjoins them to place their faith in God's justice and the rewards of the kingdom of God. It is wrong to think that this changes the nature of the game Jesus's followers find themselves in. It's not as if, in virtue of placing their livelihood in God's hands, they no longer are in a prisoner's dilemma situation. Rather in virtue of their love for God and love for being in general they perform the act that yields highest expected social utility. The truly virtuous person forgoes his or her own good for the good of the other. Thus, this Edwardsian position, reflected in Jesus's Sermon on the Mount, solves the two-prisoner dilemma.

Does the Edwardsian position solve the general multiplayer prisoner's dilemma? No. To see this I will first describe the general form of a multiplayer prisoner's dilemma and then

describe an instance of this for the truly virtuous. This shows that multiplayer prisoner's dilemmas can arise for the truly virtuous. Consequently, true virtue doesn't lead to the dissolution of multiplayer prisoner's dilemmas.

The multiplayer prisoner dilemma I describe below assumes that there is a vague threshold for cooperation required to achieve the social good and avoid the social evil. As I argued above in the water shortage cases it is imminently reasonable that the threshold for providing adequate water supply is vague. Any change of one liter will not make a difference to the outcome achieved. Standard representations of multiplayer prisoner's dilemmas assume that the threshold for cooperation is exact. To provide a suitable matrix for the game I focus on three states: the state in which that doesn't meet the threshold, the state that is at or exceeds the threshold, and the penumbral state. In the penumbral state it is unclear whether the social good will be achieved or not. I will assume that in the penumbral state there is enough cooperation to achieve at least some (but not full) benefit. When the threshold of cooperation is met then the full benefit will be achieved.

In a multiplayer prisoner's dilemma each player i faces the same matrix. The table below represents the matrix for a random player i.

Multiplayer Prisoner's Dilemma with a vague threshold

		Below the threshold of	The penumbral	At or above the threshold
		cooperation	state	of cooperation
Player	Defect	No cost & no benefit	No cost & some	No cost & full benefit:
i		Ordinal rank: 2	benefit: 4	Ordinal rank: 6
	Cooperate	Cost & no benefit	Cost & some	Cost & full benefit:
		Ordinal rank: 1	benefit: 3	Ordinal rank: 5

What should player i do? As before with the two-person prisoner's dilemma we can see that defection is the dominant strategy. If player i cooperates, then, whatever the other players do, i would do better by defecting. Inspect the above table: if player i cooperates then if the threshold isn't met, i does better by defecting; if the penumbral state is realized i does better by defecting; and if the threshold is exceeded, i does better by defecting. Perhaps, though, if all players were truly virtuous we could avoid this disastrous game. In the following I argue against this by providing a multiplayer prisoner's dilemma for the saints.

Suppose ten saints each manage their own orphanage. Christmas time is around the corner and it is time to solicit gifts. Each orphanage draws upon a common pool of resources which in normal times is not adequate to meet every need, but it is just enough for most. Each orphanage sends out letters to the community asking for donations. Now each saint faces the following decision. Should I *cooperate* by sending out only letters and relying on the donations generated by those letters or should I defect by sending out letters and then making personal phone calls? The strategy to defect should not be thought of as a departure from true virtue. In the case at hand defection amounts to making that extra call to see to it that Johnny gets the red fire truck he wants. It's implausible to think that the saint that defects from the group is acting out of anything but love for her orphanage. But if the saints each defect what they bring about is an "arms race" to give to charity. In addition to raising the required effort each year on the part of the saints, it is likely to lead to less overall charity. Besieged by letters, emails, phone calls, and visits people are likely to grow weary of the increasing intrusions. Thus, the collective effect of departure from the standard is to realize an overall worse situation. However, whatever each individual saint does will not affect the overall call. If Theresa, at the last minute, makes that

extra phone call she will bring about a better situation—Johnny gets the fire truck and all the giving is as it would otherwise be.

The orphanage case illustrates the dreadful logic of multiplayer prisoner's dilemma. Even if each player is truly virtuous, it will not avoid the painful consequences of enacting each player's best option. Thus the Edwardsian defense fails.

Conclusion

Social evil, pain and suffering resulting from the collective agency of rational individuals in multiplayer prisoner's dilemmas, is a pervasive feature of our world. Yet traditional defenses do not address it. Traditional defenses have focused exclusively on pain and suffering that results from either natural processes or from moral agency. Theists should view this problem as an opportunity to further mine the conceptual resources of theism. Additionally, social evil provides strong motivation for theists to be concerned about the structures of society. Theists should aim for societal structures that minimize this dreadful logic. Further, social evil provides theists with even more reason to long for the kingdom of God on earth. When God's presence is manifest to everyone it is plausible that multiplayer prisoner's dilemma will not arise. Everyone will have the *summum bonum* and problems from enacting the dominant strategy will not exist. Perhaps, for all we know, social evil is permitted to induce hope for the kingdom of God and motivate political action. But in lieu of the dreadful evils brought about by social evil this is at best a skeptical solution, a solution for which the best to be said is that we do not know that it is false.

Bibliography

- Adams, R. 1985. "Involuntary Sins." *The Philosophical Review* 94:1, 3-31.
- Alston, W. 1991. "The Inductive Argument from Evil and the Human Cognitive Condition." *Philosophical Perspectives* 5, 29-67.
- Hardin, G. 1968. "The Tragedy of the Commons." Science 162, 1243-1248.
- Hardin, R. 1995. *One for All: The Logic of Group Conflict*. Princeton: Princeton University Press.
- Hick, J. 1977. Evil and the God of Love. New York: HarperCollins.
- Kuhn, S. 2009. "Prisoner's Dilemma", *The Stanford Encyclopedia of Philosophy (Spring 2009 Edition)*, Edward N. Zalta (ed.), URL http://plato.stanford.edu/archives/spr2009/entries/prisoner-dilemma/>.
- Niebuhr, R. 1932. Moral Man and Immoral Society. New York: Charles Scribner's Sons.
- Plantinga, A. 1974. God, Freedom, and Evil. Grand Rapids, MI: Wm. B. Eerdmans.
- Rowe, W. 1979. "The Problem of Evil and Some Varieties of Atheism." *American Philosophical Quarterly* 16, 335-41.
- _____. 1988. "Evil and Theodicy." *Philosophical Topics*, 119-32.
- Swinburne, R. 2004. The Existence of God, 2nd edition. New York: Oxford.
- van Inwagen, P. 2006. The Problem of Evil. New York: Oxford.