3

The Evidential Weight of Social Evil

Joseph Corabi

INTRODUCTION

In his recent prize-winning essay "Social Evil," Ted Poston examines a variety of evil that typically goes undiscussed in treatments of the problem of evil.[1] In addition to the familiar moral and natural evils, Poston introduces the distinct category of social evils. "Social evil," he says:

> is an instance of pain or suffering that results from the game-theoretic interactions of many individuals. When a social evil occurs, responsibility for the outcome lies with no particular person and no impersonal force of nature; rather it lies with a group of people, each of whom may be morally in the clear.[2]

Although Poston's primary aim is to introduce and describe the phenomenon, he does ultimately conclude that social evil is "problematic" and that it provides an "opportunity to further mine the conceptual resources of theism."[3] In this paper, I take up the challenge that Poston raises. I argue that, while *apparent* social evils are disturbingly common, *real* social evils are surprisingly rare—at least when one adopts an ethical framework that is at home within Christianity (and other traditional theistic religions as well). Moreover, the social evil candidates that remain after scrutiny add little evidential weight to the kinds of non-social evil cases already familiar from mainstream discussions of the problem of evil. Hence, although social evils and apparent social evils raise significant practical challenges for Christians and others deeply invested in promoting a flourishing world, their addition to the mix does not provide significant new evidence against the existence of the Christian God. This is because it does not contribute a large quantity

[1] Ted Poston, "Social Evil," *Oxford Studies in Philosophy of Religion* 5 (2014), 166–85.

[2] Poston (2014), 168.

[3] Poston (2014), 185.

of new suffering to the world and moreover that suffering is not of a problematically novel quality. Thus, my efforts are primarily aimed not at criticizing the general atheological argument from social evil, but at a version of the argument that targets the Christian God and the God of other religions that share ethical commitments with Christianity of the sort discussed below. My hunch, though, is that this will cover a wide range of theistic positions.

My strategy in the paper will be, first, to describe social evil and give some paradigmatic examples of apparent social evils, as well as present a few qualifications and set the issues in context. Second, I will go to work discussing the ways that examples of apparent social evils can be "siphoned off"—i.e. shown not to satisfy the conditions required to count as social evil for one reason or another. This process will involve clarifying the nature of both social evil and Christian moral commitments. Next, I will examine the kinds of social evil candidates that remain after the siphoning process is completed, exploring and tying together what evidential lessons about theism can be learned. Objections will be considered, and the concluding section will then sum up the overall findings.

SOCIAL EVIL: PURE AND IMPURE

To get a clear understanding of social evil, we first need a detailed account of what a social evil is. In his discussion, Poston actually describes two distinct types of social evil, though one receives the bulk of his attention and is clearly his main target. This first variety is pain and suffering that results from "the game-theoretic interactions of rational moral individuals;"[4] this I will call "pure" social evil:

A pure social evil $=_{df}$ A scenario where agents $a_1 \ldots a_j$ all make free decisions and are morally blameless, yet pain or suffering results for agents $b_1 \ldots b_k$. This pain or suffering is brought about at least largely because of the free decisions made by $a_1 \ldots a_j$, and would not result if only a small proportion of those specific free decisions were made.[5] (There may or may not be overlap between the $a_1 \ldots a_j$ and the $b_1 \ldots b_k$, and there must be at least an a_1 and an a_2.)

[4] Poston (2014), 166.

[5] When I refer to "those specific free decisions" here, I mean the specific actions, not the existence of the choices in the abstract. In other words, to offer an example, if a_{10} faces a choice between R and S and chooses S, I mean to refer to a_{10}'s choice of S, not the fact that a_{10} faced this free choice.

I readily admit that it is possible to quibble over this definition. For our purposes, though, it will be good enough—none of my conclusions will trade on any objectionable niceties of formulation, and the definition states at least a necessary condition that any pure social evil will meet. One issue we should be mindful of, however, is that the definition does not specify whether the required blamelessness is restricted to the decision itself (considered in complete isolation) or whether it includes various background decisions in the past that have placed the individual in their present epistemic position. I will not attempt to stipulate a resolution to this issue, as there is no clear precedent to fall back on from previous discussions. In the end, whether we require diachronic blamelessness or merely synchronic blamelessness may matter somewhat for the classification of social evil, but it will not matter to an ultimate appraisal of the evidential importance of the cases. I will return to this issue later in the paper.

There are many classic candidates for status as pure social evils. Poston himself gives a fictional example that he takes to be paradigmatic: residents of Southern California aiming to conserve water during a drought. In his scenario, the Los Angeles area is in danger of running out of an adequate water supply, and so area residents must come together and agree to restrict consumption to prevent more severe consequences. But restricting water use involves a significant cost for most individuals, and many of them are involved in beneficial projects where water use is required. (Some of them take care of beautiful public gardens, for instance, while perhaps others run public swimming pools for kids.) These individuals reason that, if they were to defect and continue to use water as normal, their beneficial projects could continue as before. There would be no effect on the overall plight—after all, they are each just one of millions of agents involved. Each then individually decides to defect, knowing that this will bring about the best overall outcome.[6] But then so many individuals wind up defecting that, together, they produce an outcome which is far worse than would have been the case had they all just cooperated in the first place. (Let us suppose, as Poston does, that what each individual does is completely opaque to all other individuals, and there are enough individuals with good reason to defect that together they create a serious shortage.) The powerful insight is that, blameless and well intentioned though the individuals involved are, their choices collectively lead to a disastrous outcome.

[6] It is important to note that each individual is correct that their decision to defect will bring about the best overall outcome, given how other individuals are deciding. The problem is just that, when large numbers of individuals think this way, the overall outcome is far worse than it would otherwise be.

This fictional case is meant as an illustrative example, but there are many real life scenarios that are widely thought to involve social evils. Among these are various kinds of pollution, anthropogenic climate change, traffic congestion, abuse of antibiotics, and overfishing in the world's oceans, along with many lower-stakes everyday situations.[7]

Poston's second and less central variety of social evil is what I will call "impure social evil." Impure social evils are scenarios where at least some of the individuals are blameworthy, but where the game-theoretic machinery produces "an amount of pain and suffering that is disproportionate to the individual choices in the game."[8] Poston gives as examples here the violent conflicts in Northern Ireland and the former Yugoslavia, as well as suffering associated with the practice of dueling in early modern England. (In these cases, individuals in difficult positions acted in mildly blameworthy ways, but those mildly blameworthy actions often had devastating consequences.) It turns out to be very difficult to offer a precise definition of impure social evil—more difficult than to offer the corresponding definition of pure social evil. The reason is that it is extraordinarily tough to pin down what it means for an amount of pain and suffering to be "disproportionate" to individual choices, even if we think we have a rough intuitive grasp of the phenomenon. Fortunately, as I mentioned, these impure social evils play only a peripheral role in Poston's own analysis (rightly in my view), and so I will wait until much later in the paper to address them. Like Poston, I will dedicate the bulk of my attention to pure social evils.[9]

Before getting down to business, though, I should offer the qualification that I am not aiming to offer a demonstration that the world contains very few social evils (at least pure ones) or that whatever social evils remain are scant evidence against the existence of God. Even coming close to offering a demonstration of this conclusion would involve thoroughly examining the myriad of complexities associated with the diversity of sociological and economic interactions the world contains, as well as detailed investigation

[7] Incidentally, there are also cases of "social goods"—situations where game-theoretic interactions among bad actors produce beneficial results. Various market mechanisms are famous examples—particularly noteworthy are cases where colluders in a cartel are incentivized to defect and ruin the cartel's exploitative advantage. For a non-technical discussion, see John Cassidy, *How Markets Fail: The Logic of Economic Calamities*, New York: Farrar, Straus, and Giroux, 2009.

[8] Poston (2014), 180.

[9] Social evils are not reducible to natural evils because they involve decisions which are free *ex hypothesi*. They are not reducible to moral evils because the individuals are not blameworthy, or at least not in proportion to the evil. In any case, the classification of evil in terms of natural evil and moral evil needs sharpening, independently of the issues posed by social evil.

into strategies for addressing evidential arguments from evil. This is just not possible in anything short of a monumental tome. My effort should instead be seen as an attempt to render my conclusion plausible and offer helpful guidance for future systematic investigation.

SIPHONING: EGOISM

It is no secret to anyone that Christianity, in virtually all its forms and throughout its entire history, has been unremittingly hostile to egoistic moral views. One of the Gospels' most famous exhortations is to "love your neighbor as yourself."[10] The rest of scripture is filled with admonitions to love others and to make sacrifices for their good, and Christian tradition has followed suit univocally.[11] While certainly not news, it is important to keep this in mind in the context of addressing social evil, since many classic examples of game-theoretic conundrums involve purely rational agents, where "rational" action is stipulated to involve only self-interested motivation.[12] The original prisoner's dilemma is a case in point. The parties involved are concerned only with their own individual welfare, and this concern with their own individual welfare causes each of them to betray their co-conspirator, resulting in a worse outcome for both players than if they had cooperated with one another.[13]

This analysis can also plausibly be extended to many real life cases of prima facie social evil. Consider many instances of pollution, for instance, where an individual's particular polluting actions are unnoticeable in the grand scheme of things, but all of the individuals together collectively make the problem far worse. Often, the agents involved are motivated by nothing more than trivial gains in personal convenience or tiny financial benefits, with little by way of noble motives anywhere in the vicinity.

[10] Matthew 19:19 and 22:39, Mark 12:31, and Luke 10:27, with precursor in Leviticus 19:18. Exhortations to love others are also common in John—see for instance 13:34. (See also Romans 13:9.) Some of these particular passages do make references to oneself, but of course they only make reference to oneself as a way of instructing the reader not to act in a merely self-interested way.

[11] Both the Islamic and Jewish traditions also have strong prohibitions against egoistic moral reasoning.

[12] Poston is of course careful to avoid these by stipulating in his examples that the agents involved are rational *moral* individuals. But since these kinds of self-interested cases are so common in the literature, it is worth addressing them.

[13] This is made clear by Poston's Jonathan Edwards-inspired suggestion in section 4.2 of Poston (2014).

The devastating recent financial crisis is another case in point. While replete with potential sources of social evil, many of those scenarios manifestly involved actors who were at least largely selfishly motivated, and whose motivations were clearly wrong according to Christian moral principles—they involved disregard for the welfare of others, and thus were not demonstrations of love. Take, for example, the market analysts and decision-makers who appear to have seen disaster on the horizon but adjusted their forecasts in an optimistic direction to bring them into line with conventional wisdom, thus seeking crowd protection (often successfully) and insulating themselves and their jobs from the wrath of their superiors, investors, and boards. Often this strategy allowed them—particularly high-level decision-makers—to reap hefty financial rewards for longer than they would have by telling the truth.[14]

All of these kinds of examples (with perhaps a scant few exceptions) violate the above definition of pure social evil, because all of them involve actors who are not morally blameless, at least if a Christian moral theory is correct. Their selfishness is a clear violation of basic tenets of Christian morality. One might object, however, that although their behavior is clearly *wrong* by Christian lights, it is not necessarily morally blameworthy by those same lights. These individuals may not know that Christianity is true or even believe in it, after all, and their intuitions may tend in an egoistic direction. Even if these individuals are Christians, they might be blamelessly ignorant of what their religion teaches about such ethical matters. Consequently, their actions may be examples of blameless wrongdoing, which would leave the scenarios they are involved in as serious candidates for pure social evil status.

This is an important objection that we will have occasion to consider in other contexts as we go along. In fact, it is important enough that I will give it a name—henceforth, I will refer to it as the "Blameless Wrongdoing Objection." It is not particularly plausible in the egoism case, though. It is hard to believe that a significant number of the individuals involved in the candidate scenarios under discussion really believe that their selfish actions are morally permissible, or at the very least have come to their warped moral views without previous blameworthy actions (either their own or someone else's) that have put them in an unfavorable epistemic position.

[14] See the discussion in Cassidy (2009), 177–9. In a famous example, Angelo Mozilo, the CEO of Countrywide, was not dishonest about future prospects (at least privately), but he did make decisions which were clearly designed to serve his own interests, not those of his shareholders or the broader economy (Cassidy (2009), 246–7).

SIPHONING: PARTICULARISM

While many apparent pure social evils can be disqualified due to the inappropriately selfish motives of the individuals involved, not all cases are so easy to handle. This is because, in many real-world instances, the decision-makers do not seem to be motivated by their self-interest, but rather the interest of some particular group that may not even include them. In some cases, the tendency to promote the welfare of the group is ultimately explicable by self-interest—the manager who looks out for her company's interest so she can get a promotion and buy a fancy new car, for instance, or the man who helps his next door neighbor in anticipation that his aid will be reciprocated. But many real world examples seem to resist such facile assimilation to self-interested explanation. What is to be said about them?

While not as vocal in its condemnation of this kind of preferential treatment as in its condemnation of egoistic moral reasoning, the Christian tradition has also tended to strongly oppose giving special moral treatment to particular groups, especially when one is a member of the group or the group enjoys privileged status in some important way.[15]

Scripturally, perhaps the clearest expression of the theme is in the tendency to use familial and neighbor metaphors to describe all of humanity.[16] In addition to specific sayings and the occasional use of metaphors, there is widespread scholarly agreement that a prominent theme in the Gospels, particularly in Luke, is Jesus' desire that his disciples transcend destructive and oppressive social arrangements that rely on the special treatment given to members of exclusive groups. For instance, this desire is thought to underlie much of the Lucan Jesus' opposition to the agendas of the Pharisees.[17] Consider also Jesus' exhortations to love one's enemies—the ultimate outsiders.

As I mentioned above, the Christian tradition has recognized more instances where special treatment to groups is warranted than where special treatment to self is warranted. Christians typically believe that special obligations accrue to family members in virtue of being family members and in virtue of their role in the family, for instance, perhaps the most prominent

[15] There are certainly more caveats and qualifications in the particularism case than in the egoism case, however.

[16] Relevant material includes the story of the Good Samaritan (Luke 10:29–37), the warnings against attachment to family (e.g., Luke 14:26), and Jesus' reactions to his own family (Mark 3:31–5, and softer versions in Luke 8:19–21 and 11:27–8).

[17] See, for instance, the discussion of Jesus' table fellowship with outcasts in Joel B. Green, *The Gospel of Luke*, New International Commentary on the New Testament, Grand Rapids, MI: Eerdmans, 1997, 244–50.

being the duty on the part of children to honor father and mother and the corresponding duty of father and mother to care for children and give them spiritual instruction.[18] In the same vein, Christians commonly hold that we need not show total altruistic indifference to friends and family members as compared with complete strangers, though there are certainly strong limits on how much special favor is tolerable.[19]

It is clear, however, that in many real world cases of prima facie pure social evil where the parties avoid selfish motivation, they do not avoid motivations that are objectionable on the above grounds. The particularity of the interests they take into account transcends the strong limits of appropriate favoritism (or appears in contexts where any particularity is inappropriate), and consequently the cases fail to qualify as pure social evils in spite of initial appearances.[20]

There remains the familiar issue of the Blameless Wrongdoing Objection to deal with, however. And this objection is more plausible in many scenarios where particularistic rather than egoistic motivations are in play, precisely because favoritism for things like family, ethnicity, or nation is more widely believed to be morally appropriate in general than favoritism for self (and arguably more often innocently believed to be such). Hence, it would be much easier for a non-Christian or an uninformed Christian to fall into this sort of epistemological trap without being to blame for it. (Though there would of course be limits—presumably many cases of particularistic motivation would involve giving in to temptation in a blameworthy way, rather than being blameless expressions of the agent's genuine moral convictions.) Let us bracket this worry for the moment, though, and consider what is perhaps the starkest challenge to the claim that pure social evils are much rarer than we might initially think: impartial altruistic cases.

[18] See, for example, Ephesians 6:1–4, citing Exodus 20:12 and Deuteronomy 5:16. A recent example of a philosopher affirming that parents have special obligations to children can be found in Alexander Pruss, *One Body*, Notre Dame, IN: University of Notre Dame Press, 2013, especially 184–5 and 381–92.

[19] It may well be that the tradition—even most particular segments of the tradition—does not have precise views about exactly how much special favor is tolerable. But the key is that, while there may be uncertainty or quibbling over the details, there is widespread agreement about many kinds of special treatment going too far. This is all I will rely on.

[20] Included are cases where individuals' behavior mimics that of self-interested parties, but where the motive behind the behavior is protection of family, friends, or some other favored group. Parents who work at jobs that contribute to "social evil" but who do so not out of self-interest but out of a desire to give greater opportunities to their children can easily be examples of this phenomenon.

SIPHONING: CONSEQUENTIALISM

As Derek Parfit pointed out in his ground-breaking book *Reasons and Persons*, multi-person prisoner's dilemmas and other sources of what we are calling "social evil" are not restricted to cases where the parties involved are self-interested or acting on particularized motives—social evil can afflict even purely altruistic actors who favor no particular group or individual.[21] Poston's fictional drought example, discussed in the introduction, is such a case: each of the parties has the general welfare of the world in mind (and is well-informed about the consequences of different courses of action), yet each makes decisions that collectively lead to a much worse outcome than some other set of decisions would have led to.

While such cases may not be exceedingly common in the real world, they certainly do not appear on a first take to be an extreme rarity either. One important thing to notice about all paradigmatic mechanisms that lead to social evil is that they require the actors involved to think in a consequentialist fashion—these individuals must make decisions based on the perceived goodness of *outcomes*, allowing expected consequences to trump other factors.[22] Interestingly, though, there are prohibitions in Christianity against consequentialist motivation in many contexts, even when that consequentialist motivation is purely altruistic and based on the maximization of human well-being or the general good of the world. Part of the reason for this is that Christianity is a religion whose ethical framework is at least largely grounded in love for God and our fellow human beings. But loving our fellow human beings is often incompatible with doing things that make sense when one adopts a consequentialist worldview. Alexander Pruss offers such a case: imagine a misguided billionaire offers to make me a deal. If I approach a destitute stranger, spit in his face, and then spend two minutes verbally abusing him and denying his worth (trying hard to genuinely mean what I say), the billionaire will give the stranger one million dollars. Even suppose that after the episode is over, I will get to explain to

[21] Derek Parfit, *Reasons and Persons*, Oxford: Oxford University Press, 1984, 66. Parfit addresses the issue primarily for *pure* altruists—people who treat their own well-being as morally irrelevant and count the well-being of everyone else as of equal importance *ceteris paribus*. But the same issues can arise for people who are merely impartial altruists—i.e. those who count everyone's well-being as of equal importance *ceteris paribus*, including their own.

[22] I am grateful to an anonymous reviewer for pointing out that this problem can arise for anyone who allows consequences or expected consequences to trump other factors, whether or not the individual is a pure consequentialist. For ease of exposition, I will focus on pure consequentialist motivation in the subsequent discussion, but the lesson can be extended beyond such cases.

the stranger what is going on. One can easily imagine that the stranger might be psychologically constituted in such a way that my taking the billionaire up on this deal would give the best outcome overall for him—suffering the verbal abuse and getting the million dollars might easily be better for him, in some normal consequentialist sense, than getting nothing at all. It might even be better overall for the world. Yet, by Christian lights, it would clearly be wrong for me to do this, because verbally abusing the stranger and denying his worth would be inconsistent with loving him as I ought.[23] While this example may be fanciful, the core principle is likely to apply in many real life cases.

It is also worth keeping in mind the long and venerable tradition of non-consequentialist ethical reasoning embodied in the Doctrine of Double Effect, a favored casuistic tool in many branches of Christianity.[24] Even those Christians who do not subscribe to the Doctrine of Double Effect are typically unsympathetic to the idea that we should ignore scriptural injunctions against consequentialism, and for our purposes the rejection of consequentialism is the key aspect of the doctrine.[25]

It is also important to note the lessons that come from the celebration of martyrdom—especially early martyrdom—that one finds in Christianity. As is widely known, many early martyrs died because they would not offer sacrifices to the Roman imperial cult or perform other seemingly trivial acts of betrayal, in spite of the fact that they were often given generous opportunities to do so by authorities. Their refusal undoubtedly had powerful consequences in many cases (raising awareness of Christianity and of the level of commitment of its adherents, for instance), but their reasoning did not typically appear to be consequentialist in nature.[26] These martyrs were and continue to be celebrated, moreover, and the reasons for this have had little or nothing to do with any consequences their actions had or were expected to have (at least not consequences of a problematic sort).[27]

In Christian circles, the actions of martyrs have often been taken to justify strong deontic truth-telling requirements, and Christianity in various forms has been sympathetic to a wide variety of other purely deontic principles.

[23] Pruss (2013), 27.

[24] In typical formulations, an action is impermissible, regardless of its consequences, if it is intrinsically bad.

[25] See especially Romans 3:8's implication that we should not do evil "that good may come of it."

[26] Many of these consequences were highly negative, of course, and it was probably the negative consequences of suffering and death that were usually most salient in the moment.

[27] As illustration, consider the passage on martyrs from Latin Father Minucius Felix, *The Octavius*, 37.

While there is undoubtedly uncertainty and disagreement over specifics, within mainstream Christianity there is little controversy surrounding the claim that there are many purely deontic principles that should govern human moral decision-making.

What is crucial for our purposes, of course, is the issue of how many candidates for pure social evil status involve violations of these deontic requirements, in one form or another. (This is because, if deontic requirements are violated, then we have wrongs being done. And if wrongs are being done, individuals are likely to blame for those wrongs, thus ruling out the case from status as a pure social evil.) Obviously, it is difficult to offer a definitive analysis without being clear about exactly what the deontic requirements are and exactly what candidate scenarios exist. As I mentioned above, this would be too much for any single paper to address. Still, we can make significant progress by examining particular cases that appear to be representative and assessing whether there is a deontic principle in the vicinity that Christians will plausibly accept as part of the religion's overall ethical view.

Take, for instance, Poston's sanitized hypothetical case of water conservation. Although not explicitly stated, his description implies that there is an understanding among all the residents that the best way to prevent very bad consequences is to come together and conserve water,[28] and a shared commitment to prevent very bad consequences if possible. Hence, there appears to be some tacit set of mutual promises or at least a tacit agreement among the residents not to violate reasonable water use restrictions (perhaps an application of their obligation to obey legal authorities under normal circumstances). (We can suppose that the L.A. area authorities are trusted—this is a fictional case, after all!—and that they have publicized both the need for water restrictions and the nature of those restrictions.)[29] Thus, residents who decide to defect based on consequentialist reasoning—even if it is altruistic consequentialist reasoning—appear to violate commonly held

[28] If not, we may just be looking at a case of widespread ignorance—ignorance of the ways to handle a water shortage. This ignorance is likely to be either a moral evil or a natural evil, depending on the explanation of why it is present. (In itself, this doesn't guarantee that the present scenario will be disqualified from pure social evil status, but it will have a significant effect on its evidential weight if not.)

[29] If the authorities are not trusted—as may often be the case in real life analogs of this situation—we must ask why. If they are not trusted because they have a track record of lying to the public, for instance, then we may have a case of a moral evil, or at least a social evil of little evidential significance (because it is the product of a moral evil). If they are not trusted because of the laziness of residents, the take-home message is probably similar. If they are not trusted because people do not have the energy to pay careful attention to all of the information they are getting but are blameless in this, we may be looking at a natural evil or a social evil produced by a natural evil. More on these various possibilities later.

Christian tenets about either the importance of promise-keeping or the importance of obedience to legitimate authorities.

Particularly insidious, at least potentially, are cases where an agent is faced with a decision to act on an egoistic or particularistic motive that is inappropriate because of its egoism or particularism, but where the motive can easily be transformed into an altruistic one on a modicum of reflection. How can this occur? It often happens in situations where the individual has very limited options, but can achieve a Pareto improvement by defecting, benefiting himself or his favored group significantly while making no one else worse off. (A Pareto improvement is precisely a situation where one group or individual can be made better off without making anyone else worse off.) Consider Frank the Fisherman. Frank is a man of modest means and supports his children through the revenues from his fishing. Legitimate authorities acting legitimately to curb overfishing have placed restrictions on the size of a catch, but they do not have the resources to adequately police the restrictions (and so any threat of punishment can be safely ignored). Initially, Frank is tempted to defect, take more fish, and benefit his family. But he feels it would be wrong to show this kind of favoritism to his family, since after all they aren't starving or in any desperate need. (Imagine that others like him think the same way—let us suppose that there are many.) But then it dawns on him! If he takes a few more fish, this will do absolutely nothing to the catch of any other fisherman in the entire world.[30] Hence, by taking a few more fish, Frank is not merely improving the welfare of his family—he is improving the welfare of the entire world (the entire human world at least)! In this case, his distaste for particularism will not impede him, and—assuming the others reason similarly—the social evil will be well on its way to being done.

But again, Christianity has something to say about this case. Even though Frank's action produces a Pareto improvement and is justifiable on altruistic consequentialist grounds, it is still impermissible by the Christian's lights. Because Frank failed to obey a legitimate authority's regulations (or perhaps because he failed to abide by an agreement that he made), he has done

[30] Let us suppose that this assumption is true. If it is not—if Frank is actually affecting the probability that others will catch fish (either now or in the future)—things will get trickier. We (and Frank himself, if he is being careful) will have to assess the expected effects of his decision. If they are sufficiently small, then his decision will be exactly the same de facto as if his action had no effect. If the effects are large, then we are probably not looking at a social evil candidate anymore. Rather, this will just be a straightforward moral decision that can be analyzed as a potential instance of moral evil if a wrong decision is made, assuming that Frank is to blame for the wrong act. (If not, it is probably going to be a case of blameless ignorance, which is standardly classified as a natural evil.)

wrong.[31] As a result, his case will fail to satisfy the requirements for pure social evil, at least so long as the Blameless Wrongdoing Objection can be overcome. Before finally addressing this objection and other issues, we will consider a handful of miscellaneous siphoning strategies. Thus far, our siphoning strategies have concentrated on cases where agents fail to satisfy the requirements for blamelessness that pure social evil requires. Subsequent ones will focus on the failure of some of the problem-causing game-theoretic idealizations to apply to most real world cases.

SIPHONING: OTHER STRATEGIES

In addition to the strategies we have already seen for explaining away candidates for pure social evil status, there are several others. Three of them are especially important. The first involves noting that a major source of social evil in the abstract is the "simultaneous move game"—a variety of game-theoretic interaction where all of the players act in isolation from information about other players' decisions. Poston's drought example is like this—all of the players make a decision to either cooperate or defect without any specific information about how the other players are choosing. But such situations in the real world are rare. Typically, one is both receiving information about the choices of players who have already moved and broadcasting information about one's own choice for players who have not yet moved. This allows for more effective cooperation among players, particularly when this feature of the interaction is combined with the two other features discussed below.[32]

One other such feature—the source of the second important miscellaneous siphoning strategy—is iteration. Comparatively few real life social interactions involve only one "round" in the way that the classical prisoner's dilemma does (although such interactions may be getting increasingly common in a world where anonymous social encounters in big cities and online

[31] Although I will not pursue the suggestion here, it may be that Christianity's insistence on deontological principles even provides a bit of evidence in favor of Christianity, assuming that early proponents of the religion could not have foreseen the usefulness of deontological commitments in forestalling social evils.

[32] It will allow for more effective cooperation for two reasons. First, because there will be additional information that will allow individuals to coordinate their actions. And second, the additional information will alert individuals of the impending dangers and allow those individuals to take steps to confront the behavior that is giving rise to the problem. I am grateful to an anonymous reviewer for pointing out the need for this clarification.

are on the rise). But it is well-known that, when prisoner's dilemma-like scenarios repeat, it will often be possible for the players to cooperate and avoid social evil, because they are able to monitor one another's behavior and issue threats and punishments for defection. Of course, if the players know that the series of interactions is coming to an end, they will revert to the behavior that would be appropriate if there were only a single round. But often, under real-world circumstances, agents do not see this end point coming.

The third relevant feature is the tendency of real-world defection decisions to probabilize the bad effect. In Poston's drought example and other classic candidates for social evil status, defection is especially seductive because the cases are set up in such a way that, by defecting, one does not increase the likelihood at all that the bad effect will occur nor does one tangibly worsen the situation, even in a small way. In the drought scenario, for instance, any individual's choice is stipulated to have no effect on whether the negative outcome is realized (and presumably no effect on the probability that it or the alternative will be realized).[33] But in the real world, things are not generally this tidy, particularly when there are multiple iterations of the interaction and one is broadcasting information via one's choices. But even when there is only one round, often real life is not so simple. My story of Frank the Fisherman above may make this clear. Although I stipulated away this complication originally, a more realistic version of the story would have to acknowledge that if Frank takes more fish, there are likely to be consequences down the road, even assuming Frank and his colleagues are about to sail into the sunset of fishing retirement. These fish will no longer be around to reproduce and increase the fish population, which will probably have small but tangible effects in future fishing efforts.[34] (There may be a small but non-zero chance that it will have large effects.) In some cases, these complications may be enough to affect payoff structures and force moral actors into cooperation with one another, even setting aside worries about the appropriateness of consequentialist reasoning. (Hence, in these instances when

[33] I assume here that whenever one has an effect on the probability of an overall outcome (at least in this context), one counts as having an effect on the realization of the outcome (if the outcome in fact occurs).

[34] Because many classical game-theoretic puzzles presume that the parties are motivated by pure self-interest, they often do not have reason to consider such complications in their payoff structures. Even assuming he doesn't retire, Frank's chance of affecting his *own* future catch by taking more fish really is zero, for all intents and purposes. Given that there are billions of fish in the sea, what is the chance that he personally is going to run into this one again or one of its offspring? Interestingly, the U.S. Supreme Court backed legislative attempts to solve coordination problems of this sort in Wickard v. Filburn. (Thanks to Dan Tyman for alerting me to this example.)

cooperation fails to result, we know that we are not looking at pure social evil, at least not unless the parties involved are blameless wrongdoers.)

We have now examined numerous reasons why prima facie candidates for status as pure social evils fail to qualify as genuine pure social evils on a Christian framework. Many of these candidates involve motivation that is blameworthy by Christian lights, while others involve factors (such as iteration and information sharing) that make them disanalogous from the kinds of classic game-theoretic scenarios that those worried about the evidential weight of social evil take as paradigms.

Now that we have examined all the various siphoning strategies, there are likely to be a number of objections. I will examine what I consider to be the most pressing of these in the next section.

OBJECTIONS

(A) Your argument treats acting in self-interest (i.e, acting in an egoistic way) as though it were inherently wrong. But acting in self-interest under some circumstances is morally permissible—indeed, even morally obligatory. But then we wind up with a far greater number of serious social evil candidates than you are letting on.

Reply—I acknowledge that sometimes egoistic decisions—most obviously egoistic decisions made in financial contexts—are neither blameworthy nor wrong. This is because sometimes in these situations the pursuit of self-interest is morally justified and appropriate. (Or, at the very least, sometimes the pursuit of personal economic gain in financial and other markets is appropriate, and such action will often mimic the pursuit of self-interest.) In my view, some cases of this sort do survive scrutiny and do appear to be legitimate social evil candidates. I will address these at greater length later in the evidential section below.

(B) The deontic requirements that Christianity embraces are understood to be defeasible. If I am a doctor and I promise to have lunch with you, for instance, but on the way run into a man having a heart attack, I do not do wrong by stopping to help the man, even if this causes me to miss lunch. Now, one might think that defeasible requirements of the nature we are discussing can be defeated in a case where I produce a Pareto improvement by defecting (i.e. I create a situation where someone is made better off without anyone being made worse off). And, as per Poston's stipulations, by using extra water in his drought case I produce just such a Pareto improvement. Thus, the mere existence of deontic requirements within Christianity does not rule out many social evil candidates.

Reply—The trouble is that the defeasibility standards in these cases are generally taken to be at least fairly high. Turning on a hose to fight a fire that is about to consume my house is one thing; turning it on to water my fruit trees (Poston's example) or a community garden is another. Not just any old Pareto improvement is enough. One might object that this is unprincipled—why would a true moral code not allow an agent to produce the best outcome available to them in a situation like this? Note, though, that there is great wisdom in a moral code designed with prohibitions against bringing about certain kinds of Pareto improvements. This prevents morally upright agents from causing the kinds of widespread social evils that would arise easily if these restrictions were not in place.[35]

(C) Some of your examples rely on Christianity's alleged deontic requirement to obey legitimate authorities, under at least some circumstances. Many branches of Christianity do allow for individuals to disobey authorities when their conscience conflicts with the dictates of those authorities, however.

Allowances for conscientious objection to authorities are notoriously tricky within Christianity, and unfortunately it would take me too far afield to treat them in depth here. In any case, they are unlikely to apply to paradigmatic social evil candidates. The sorts of conscientious objection countenanced by Christianity always involve one of three things: (1) the individual has a conviction that the authorities have false or unjustified beliefs about morally relevant empirical phenomena, (2) the individual has a conviction that the authorities have false or unjustified beliefs about fundamental goods, or (3) the individual disagrees with authorities over the existence of a deontic moral principle. None of these rationales applies in the paradigmatic situations I discuss.

WHAT'S LEFT?

Now that we have completed our survey of siphoning strategies and examined a number of pressing objections, there appear to be several kinds of pure social evil candidates that survive. In this section, we will take stock of what remains, and in the process I will argue that pure social evils due to blameless wrongdoing are not evidentially important and can be safely discounted.

[35] I am grateful to an anonymous reviewer for spurring me to expand my discussion of this objection.

Let us first consider scenarios where the Blameless Wrongdoing Objection is plausible—cases where the bad outcome is produced by a collection of free decisions that are wrong (according to Christianity), the bad outcome would not be produced (at least not in full) by any small proportion of the free decisions, but where the agents involved all manage to avoid blameworthiness for their wrong choices. While there are likely to be some cases of egoistic motivation that fit into this category, most will be cases of particularistic or altruistic motivation. (Among moral attitudes that conflict with Christianity, these tend to be the ones that non-Christians or poorly informed Christians are most likely to adopt. Many instances of particularistic or altruistic motivation will not be blameless, though, even if particularistic and altruistic cases are more common among blameless wrongdoing cases. There will be plenty of instances, in other words, where the agents involved are clearly giving in to some temptation, not acting on genuine moral convictions.)[36]

While cases that qualify as pure social evils due to blameless wrongdoing do exist, they are not evidentially important. To understand why, we must classify them further. Blameless wrongdoing cases will be situations where either (A) the parties are blameless in the present decision at least partly because they or someone else is to blame for putting them in a deficient epistemic position,[37] (B) the parties are blameless in the present decision because of natural obstacles that put them in a deficient epistemic position, but whose effects are in no way due to the blameworthy choices of agents,[38] or (C) the parties are blameless in the present decision because of previous pure social evils.

Some cases falling under (A) may not even count as satisfying the definition of pure social evil, because my definition (following Poston's more informal characterization) does not specify whether the blamelessness can be restricted to the present choice considered in complete isolation, or whether it requires blamelessness in relevant choices leading up to the present one. In general, even if they exist, cases falling under (A) are evidentially untroubling to theism, or at least add little or no new troubling evidence. This is because they are examples of moral evils, and moral evils (even horrendous ones) are unfortunately already all too common and familiar.[39]

[36] If nothing else, the material above should convince readers that social evils among informed Christians are quite a bit rarer than one might think.

[37] Included are cases where agents deceive themselves about probabilities or payoffs in the scenario.

[38] There are subtleties about the definition of "moral evil" and "natural evil" that are important in the grand scheme of things, but can be ignored for our purposes.

[39] We can understand a moral evil roughly as a bad outcome of a blameworthy decision, and hence the consequences of these social evil situations are also moral evils

Similar things can be said about cases falling under (B)—the seriousness of natural evil, including natural evil that results in the moral ignorance of agents, is already recognized.[40] Such moral ignorance already produces devastating effects in a large number of familiar cases where no problematic game-theoretic interactions occur, so it seems doubtful that piling on some additional negative effects of naturally produced moral ignorance will make much difference.

Cases falling under (C) appear to be non-existent, at least where the social evil traces all the way back to an ultimate source in social evil (i.e. when we continue to trace our way back through the causes of this particular scenario, we never wind up with an explanation that ultimately rests on something other than moral evil, natural evil, or some combination thereof). How plausible is it that the blameless ignorance of the agents involved would be due only to social evil? We have seen strong reasons to believe that pure social evils are rare to begin with—even when we require only synchronic blamelessness—so what are the chances that we are going to find complicated diachronic social evils going all the way back in the explanatory sequence?

While this is possible, it is sufficiently obscure that the burden of proof clearly lies with the atheologian to produce real-life cases where it operates. To put the matter another way, the challenge for the atheologian is to describe a possible world where the only evils are pure social evils, and to provide some reason to think that the actual world shares some of the relevant features with this world.[41]

Aside from instances of blameless wrongdoing, what other kinds of scenarios slip through the cracks? There are cases where egoistic, particularistic, or altruistic consequentialist motivation are morally appropriate, and can lead to suboptimal outcomes. Arguably, there are economic scenarios of this sort. Take, for instance, situations where the "Paradox of Thrift" arises. In these cases, the sensible frugality and caution of all the individual actors in an economy leads to very negative consequences for the economy as a

because they are indirect consequences of blameworthy past decisions. If one prefers to restrict moral evils to *immediate* consequences of blameworthy decisions, then these social evils can be described as effects of moral evils rather than moral evils themselves. The substantive point is the same. (The introduction of impure social evil may require a complication of the basic definition of moral evil, but my arguments below are designed to show otherwise. In any case, there are other reasons to make the classification of evils more subtle than typical discussions of the problem of evil do, but as I mentioned previously these are not relevant for present purposes.)

[40] Some of these problems are likely related to the problem of divine hiddenness, another evidential issue that is already prominent in the theism debate.

[41] I am grateful to Jonathan Kvanvig for suggesting a formulation along these lines.

whole.[42] No one would challenge the wisdom or moral appropriateness of individual frugality in situations of economic stress (especially not Christianity, which has always valued responsible stewardship of resources), but it can collectively lead to serious suffering. An altruistic example that may slip through the cracks—though not necessarily one with many close analogues in real life—is Poston's saintly orphanage case, which he presents as his purest social evil example.[43]

Before continuing on to discuss evidential implications of the cases that remain, it is time to return to the issue of impure social evils. Recall that impure social evils are social evils where not all the parties involved are blameless, but where the bad consequences are "disproportionate to the individual choices in the game," in Poston's words.[44] While game-theoretic interactions (or at least scenarios closely analogous to the interactions game theorists discuss in the abstract) are a very real part of human life, it is exceedingly difficult to understand what it would be for bad consequences to be "disproportionate" to individual choices in them. Surely we cannot readily compare a decision made in a game-theoretic scenario with "the same" one made outside a game-theoretic scenario, since after all it is typically of the very essence of the decision that it was made in a game-theoretic scenario. (The dilemma I face in the drought situation, for instance, is not at all the same kind of choice that I would face in a similar scenario where I was the only potential consumer of water in the area. The only way the decision in the scenario could be qualitatively the same as one outside it is if I mistakenly believe that the scenario I am involved in does not include any other individuals whose actions could affect outcomes.) It could be that Poston has in mind that each decision has a level of praiseworthiness and blameworthiness associated with it, and that we can compare different levels of praiseworthiness and blameworthiness in different kinds of decisions—specifically, in our case, the praiseworthiness/blameworthiness of decisions in game-theoretic scenarios with ones outside game-theoretic scenarios. But this won't do either. Regardless of the finer points of one's theory of praiseworthiness or blameworthiness, it seems clear that there is no normal

[42] A famous real life mini-example of this paradox occurred in the "Capitol Hill Babysitting Co-Op" in the 1970s. See the description in Paul Krugman, *Peddling Prosperity*, New York: W.W. Norton, 1994.

[43] Poston (2014), 184. In this case, an individual running an orphanage is faced with the decision to make extra phone calls to solicit donations to buy a Christmas gift for an unfortunate child, but where individuals running other orphanages are dealing with similar issues. Even here, there are reasons to be suspicious, on some of the same grounds discussed above.

[44] Poston (2014), 180.

general consequence for any given blameworthiness level which we can use to calibrate what is "proportionate." This is because the same level of praiseworthiness or blameworthiness in different contexts often produces dramatically different consequences, and no context can be identified as *the* standard one. (A small child making a decision with a specific level of blameworthiness and the President of the United States making a decision with the same level of blameworthiness will likely have very different outcomes, even when no game-theoretic machinery is involved.)[45]

I am inclined, then, to conclude that there is nothing mysterious or idiosyncratically problematic about the way that game-theoretic machinery causes bad consequences in the context of blameworthy decisions. These simply are the consequences of the kinds of blameworthy decisions that occur in those scenarios. We must learn to identify them as such in the same way that we learn to identify the different kinds of consequences that occur with blameworthy moral decisions in the diversity of non-game-theoretic contexts we are familiar with from history, politics, and daily life. For those unsatisfied with this response by itself, I should also point out that impure social evils are clearly more corrupted by moral evil than are pure social evils, since moral evil is directly associated with the decisions in the scenarios themselves.

For all of these reasons, I am not inclined to see impure social evils as evidentially threatening, especially if it turns out that pure social evils are not (as I will argue in the next section). After all, the main threat associated with pure social evils is that they introduce a novel kind of evil, supposedly unaddressed by standard theodicy strategies. But there is nothing novel about the kind of evil associated with impure social evil, once we have taken account of pure social evil along with moral and natural evil.[46]

[45] Perhaps there is some way of averaging the outcomes of all the decisions with this level of blameworthiness in every possible world, but such a process is likely to be enormously difficult. Moreover the result of the analysis is far from clear, so it will be hard to say which decisions will turn out to have disproportionate effects. There are also cases of natural bad luck, where unforeseen natural circumstances magnify a bad effect of a decision, but the magnification in these cases is much easier to quantify.

[46] Quantity clearly does play some role, and I admit that I cannot fully address messy issues surrounding how the overall quantity of relevant evil in the world is affected by impure social evil. It may also often be more difficult to foresee the consequences of one's choices in game-theoretic cases, but of course it is also often difficult to foresee these consequences in non-game-theoretic scenarios as well. As I mentioned above, this paper should not be seen as an attempt to offer a once and for all demonstration of the evidential insignificance of social evil.

EVIDENTIAL IMPLICATIONS

Now that the siphoning process is complete, we have seen that there are a small handful of candidates that do appear to qualify for bona fide pure social evil status. These include some cases of blameless wrongdoing and some cases where morally appropriate consequentialist reasoning is employed (whether of an egoistic, particularistic, or altruistic sort). The crucial issue now is of course their evidential importance for atheism. In this section, I will argue that the remaining social evils do not constitute significant evidence against theism, because they are rare and ultimately trace back to some combination of moral and natural evil.

The evidential weight of the blameless wrongdoing cases was addressed above. Pure social evils of this sort that are ultimately caused only by other pure social evils don't appear to exist. Other kinds of pure social evils based on blameless wrongdoing appear to be much rarer than we might have initially supposed. Consequently, it is hard to see how they could be of significant evidential weight. Ones that are ultimately explicable in terms of moral evil don't appear to add a great deal to the quantity of moral evil in the world (since moral evils, including severe ones, are unfortunately very common and these kinds of social evils are not). Nor does the quality of the new evil appear very significant. After all, moral evils are already known to produce devastating consequences in all sorts of different ways. If we were to discover that there is a previously unnoticed way that moral evils cause devastating consequences—namely, by creating conditions that give rise to pure social evils—we would not learn something of much evidential significance, because there is no reason to think that God would be more likely to act to prevent indirect bad consequences of moral evils than direct ones. The same basic point goes for such social evils that have their ultimate cause in natural processes—again, natural processes are already known to produce devastating consequences, including natural processes leading to ignorance and cognitive failure. Piling on a few new cases of a subtly different kind does not appear to be of great evidential significance.

What about cases where legitimate consequentialist reasoning is used? These are trickier to deal with, but similar considerations point us toward the same conclusion for them. To understand these cases, we must understand what gave rise to the scenarios *themselves*, as opposed to any moral ignorance that is present in them. (These scenarios do not contain moral ignorance *ex hypothesi*—the individuals in them are not mere blameless wrongdoers, but individuals who positively do right.) Again, there are three basic options (ignoring subtle complications): (1) they are ultimately produced

by blameworthy moral choices (even though the scenarios themselves involve no blameworthy choices), (2) they are ultimately produced by natural processes, or (3) they are ultimately produced by further social evil unrelated to either blameworthy moral choices or natural processes.

As with the blameless wrongdoing classification above (but with even more plausibility here), I conjecture that cases of (3) do not exist. While I can offer no demonstration of this thesis, the material I have presented thus far in the paper should have convinced the reader that pure social evils are quite a bit harder to produce than one might initially think, whereas moral evils and natural evils are ubiquitous and occur easily. Thus, it would be surprising if the world contained pure social evils that are social evils to their core, depending not at all on previous blameworthy moral decisions, natural scarcity of resources, or other phenomena already familiar from mainstream discussion of the problem of evil (or natural failure to discern effective solutions to coordination problems).[47] It is hard to fathom what such evils would even look like. Consequently, they can be safely ignored until someone comes forward with plausible examples of them.

Cases of (1) likely comprise a substantial portion of pure social evils that involve no blameless wrongdoing in the social evil scenario itself. But again, these kinds of cases are not of great evidential significance. The world already contains substantial amounts of moral evil, and such a small number of genuine cases of pure social evil remain after siphoning that we are likely looking at little more than a drop in the bucket. And although social evils are different in quality from moral or natural evils, their quality difference does not appear to be of great enough significance to justify a drastic solution on God's part. After all, removing from the world the social evils ultimately caused by blameworthy free choice would require God either to eliminate all free choice (or at least the tangible consequences of free choice) or to selectively eliminate free choices that lead to pure social evils. But, given that pure social evils appear to be rare, it is inelegant for a God who so values free choice that he allows it even in (non-game-theoretic) scenarios where it leads to tremendous destruction to eliminate it as soon as an occasional negative game-theoretic interaction threatens. It is hard to fathom what sort of principle such a God could be operating under.

Similar points go for cases under (2). Many pure social evils will fit in this category, produced by the combined natural processes of scarcity of some valuable resource plus naturally produced cognitive limitations. (The natural cognitive failure could manifest itself in failure to see workable solutions for

[47] Arguably, all financial examples of social evil ultimately have resource scarcity as at least part of their explanation.

coordination problems.) But again, severe natural evil is already well known and its evidential implications well explored. It may be that the effects of natural evil are magnified because some natural evils give rise to social evils, but given the wide range of catastrophic consequences of many regular natural evils that are unassociated with game-theoretic scenarios, it is unlikely that the evidential significance of natural evil will be greatly enhanced by the fact that some natural processes cause social evils. The plausibility of commonly produced justifications for natural evil (soul-making, the benefits of stable laws of nature for moral learning or aesthetic beauty, etc.) are not likely to be affected by a few additional cases of suffering thrown on top of the already familiar ones, even if those cases are of a novel sort.[48]

Nearing the end of our journey, the reader might raise one final practical objection. If God's ultimate justification for allowing social evils is just that he *really* likes the kinds of things theodicists always talk about (significant free will, soul-making, natural law stability, etc.), why not write a very short paper saying that?! Why spend so much time "siphoning off" social evil candidate scenarios? My response is that God may really like the kinds of things theodicists talk about without *really really* liking them. At some point, the quality and quantity of suffering may reach a threshold that any God worthy of the name would be unwilling to tolerate, and so such a level would provide clear evidence against the existence of such a God. Adding a high quantity of suffering caused by the choices of blameless moral agents (often acting rightly no less!) might very well push the world over that threshold. Hence, it is important to do our best to show that genuine pure social evils are rare.

As we have already discussed, a key unresolved issue in deciding exactly how much pure social evil exists is settling on whether pure social evil requires diachronic blamelessness or only synchronic blamelessness. But this issue is not important for evaluating the evidence social evil provides against theism. This is because the more permissive synchronic definition will allow for some additional social evils, but the further social evils it permits will be among the least interesting from an evidential perspective—the ones that are brought about as a result of previous blameworthy moral decisions.

[48] Poston (2014) does discuss soul-making and natural law stability as theodicy strategies for addressing social evil, but he merely points out that they can't play any role in directly explaining the existence of social evil. This is true, but it discounts the possibility that they play an indirect role, by justifying God in allowing processes that have social evils as consequences.

CONCLUSION

We have now seen a host of siphoning strategies that show that many candidates for social evil status don't qualify. In addition, we have examined a host of reasons to think that the candidates that remain (both the ones that involve blameless wrongdoing and the wholesale avoidance of wrongdoing) do not provide significant evidence against the existence of God, because they appear to be produced by processes that are already familiar from mainstream discussions of the problem of evil. These processes are already known to produce significant suffering in other contexts and are already the targets of much attention by theodicists. The central argument of this paper has been non-demonstrative, however, since offering anything close to a systematic case would require surveys of vast swaths of both empirical and conceptual territory. It is my hope that the ball has been put squarely in the court of the atheologian, though. If there are collections of social evils that are problematic for theism, they have not yet been produced or identified.[49]

[49] Thanks to editor Jonathan Kvanvig, Jamie Hebbeler, Paul St. Amour, Dan Tyman, Brandon Gergel, Becky Germino, and several anonymous reviewers from *Oxford Studies in Philosophy of Religion* for their helpful feedback on previous versions of this paper.