

EXCESSIVE GRATUITOUS EVIL

Klaas J. Kraay
Toronto Metropolitan University

ABSTRACT: This paper draws together and engages with two recent – and independent – discussions of the problem of evil. Bruce Russell (2018) examines four arguments for atheism that appeal to suffering. He rejects the first three, but defends the fourth. Meanwhile, separately, William Hasker has discussed close variants of the third and fourth arguments. In an important but underappreciated series of papers, he criticizes the former (Hasker 1992, 1997, 2004b, 2008). More recently, he has deployed this criticism against the latter as well (Hasker 2019). The order in which Russell treats these four arguments is helpful and instructive, and so I will follow it. I will briefly discuss the first and second. I will then set out Hasker’s criticism of the third argument, and offer some resistance to his most recent defence of it. I then turn to the final argument, which I call the argument from *excessive gratuitous evil*. Russell and Hasker both think that it constitutes a formidable problem for theism. I agree. I do not discuss Russell’s (indirect) defence of it. Instead, I examine Hasker’s latest objections to it – including his new deployment of his earlier criticism – and I find them all wanting.

1. STAGE-SETTING

Bruce Russell (2018) examines four different arguments for atheism that invoke suffering. Here are the first premises of each:

- (1A) If God exists, there would be no suffering.
- (1B) If God exists, there would be no unnecessary suffering.
- (1C) If God exists, he would not allow unnecessary suffering unless allowing it is needed to bring about a greater good or prevent a greater bad.
- (1D) If God exists, he would not allow excessive unnecessary suffering.

The second premise of each argument asserts that the phenomenon mentioned in the consequent of the first premise actually occurs, and so, by modus tollens, these arguments all conclude that God does not exist.¹ Russell thinks that (1A), (1B), and (1C) are all false, and, accordingly, that the arguments that begin with them fail. He does not defend (1D); he simply assumes that it is “not open to criticism” (108). He then devotes most of his paper to defending, indirectly, the next premise: the claim that there is indeed excessive unnecessary suffering. The order in which Russell treats these four arguments is helpful and instructive, and so I will follow

it. Like Russell, I will say relatively little about (1A) and (1B). I will say rather more about (1C) and (1D). Along the way, I will propose modifications to these four claims. This is for two reasons. First, some are infelicitous as stated. Second, William Hasker has very important things to say about close variants of (1C) and (1D), and an important goal of this paper is to engage his arguments.

(1A) is widely rejected on the grounds that God might have sufficient reason to allow some suffering, in certain circumstances, to occur. This is often defended by appeal to analogy: good and loving parents, after all, have sufficient reason to allow their children to undergo some suffering in certain circumstances. For example, such parents are widely thought justified in allowing their children to endure the sort of low-level suffering typically caused by piano teachers and dentists in the course of their professional duties. Parents permit this in order for their children to obtain various goods. Two constraints can be added for greater precision. It is generally thought that good and loving agents – divine or otherwise – can permit suffering in order to bring about some good *provided that* that (i) the good cannot be obtained in some other way involving less suffering; and that (ii) the good is sufficiently valuable.² Constraint (i) seems obvious enough: if the agent could secure the desired good with less suffering, or no suffering at all, it's reasonable to expect just that. Constraint (ii) captures the intuitive idea that the suffering can't be permitted for the sake of just any old good: the good must be *good enough*. But this still leaves much open. Perhaps a good that *just barely* outweighs the suffering is insufficient. Perhaps what counts as *good enough* is person- or context-relative. And so on.³ But so long as it is plausible to think that that God can satisfy whatever reasonable constraints are offered, then (1A) can be deemed false.

Here is an independent line of attack against (1A): it is too broad in one respect and too narrow in another. It is too broad because not all instances of suffering count against theism. Consider, for example, the suffering that one voluntarily undergoes during athletic training. (1A) is also too narrow because what is *really* thought to count against theism is *evil* – and not all evil involves suffering. Suppose that someone steals all the money from your bank account. But there's a twist: due to a bank error in your favour, your balance is immediately restored, and neither the theft nor the error is ever detected. It's natural (or at any rate traditional) to think that the thief's action is morally wrong, and indeed that it is *evil*, even though it causes no suffering. Moreover, the experiences typically caused by piano teachers and dentists, while they may constitute *suffering*, are just not properly considered *evil* (no matter what my children tell you!) All of this prompts replacing (1A) with:

(1A') If God exists, there would be no evil.

Nevertheless, (1A') can still be resisted on the sort of grounds mentioned above: it seems that God can be justified in permitting the occurrence of some evil in order to bring about a good. Of course, this good must meet a slight variant of constraint (i) – let's call it (i') – namely, that it cannot be obtained in some other way that involves less evil. And this good must meet constraint (ii) – it must be sufficiently

valuable – and whatever other reasonable constraints are offered. But if it's plausible to think that God can satisfy these constraints, (1A') can be deemed false, and this argument for atheism can be set aside.

It is perhaps more difficult to see why (1B) can be rejected. Let's look at it carefully. Russell says that suffering is *unnecessary* "if and only if it is not needed to bring about a greater good" (106).⁴ Russell means that the *occurrence* of the suffering is not needed to bring about a greater good. So why might God allow such suffering to occur? Russell points out that God's *allowing* the suffering might be needed to bring about a greater good, even if the *occurrence* of the suffering is not. Russell defends this with the following analogy. A good and loving mother might allow her son to go out with his friends, even though it would be better for him to rest his injured ankle. The son's *suffering* might not be needed to obtain some greater good, but, plausibly, the mother's *allowing* the suffering (by allowing the son to go out) is needed to obtain a greater good: namely, the son's freedom to make (ir)responsible decisions. We can redeploy constraints (i) and (ii) to add more precision here. The resulting thought is that even if the *occurrence* of the suffering fails to meet (i) or (ii), a good agent can justifiably allow the suffering only if the good for which the suffering is *permitted* also meets (i) and (ii). Again, more constraints could be added, but so long as it is plausible to hold that God can satisfy whatever reasonable constraints are offered, then (1B) can be deemed false, and this argument for atheism can also be set aside.

The same considerations that prompted replacing (1A) with (1A') also motivate replacing (1B) with:

(1B') If God exists, there would be no unnecessary evil.

But again, (1B') can be resisted on the grounds mentioned above. God's *permission* of some evil can be justified, provided that the permission serves to bring about a good that meets conditions (i'), (ii), etc. So, provided that it's plausible to think that God can meet these conditions, (1B') can be set aside.

Let's now turn to (1C). This claim might seem far more secure than (1A) and (1B), and than their variants. After all, if neither the *occurrence* nor the *allowing* of some instance of suffering is needed to bring about a greater good, it is natural to expect God to prevent it. And yet, some *very* able philosophers have sought to block (1C), and claims that are similar in spirit.⁵ Russell very briefly endorses one such argument due to Peter van Inwagen, and then turns his attention to (1D).⁶

I will not discuss van Inwagen's argument here, since I have criticized it elsewhere (Kraay 2014). I will instead discuss a different argument in this vein due to William Hasker. In my view, Hasker's argument is important, intuitive, and rather more difficult to resist than has typically been appreciated. Hasker's argument targets this claim, which is similar in spirit to (1C):

(1C') If God exists, there would be no gratuitous evil.

It is often said that an evil is *gratuitous* if neither its *occurrence*, nor God's *allowing* it to occur, is needed to bring about a greater good. Given this definition, (1C') is structurally similar to Russell's (1C). But Hasker himself offers a different definition: a token or type of evil is *gratuitous* if and only if God, if he exists, antecedently knows it to be certain or extremely probable that he could prevent it in a way that would make the world overall better (2010a, 308).⁷ I take it that the final clause is motivated by considerations similar in spirit to (i') and (ii). After all, with respect to (i'), if God permits an evil for the sake of some good, but this good could have been obtained with less evil, or no evil at all, then, *ceteris paribus*, God's permission of this evil fails to make the world overall better. And the requirement that preventing the evil would make the world overall better is one way to specify the generic requirement in (ii) to the effect that the resulting good must be *good enough*.

In section 2, I will set out Hasker's argument against (1C'), using his preferred definition of gratuitous evil. In section 3, I will discuss a promising way to resist Hasker's argument, and criticize some recent rejoinders from Hasker. Ultimately, however, I will grant – for the sake of discussion – that Hasker's argument against (1C') succeeds. All of this will set the stage for section 4, in which I consider an argument for atheism that begins with:

(1D') If God exists, it is false that gratuitous evil occurs *far in excess* of what must be permitted by God in order to achieve his goals.

As we will see, Hasker has recently urged that his arguments against (1C') can be redeployed against (1D'). In section 4, I argue that this move fails. Then, in sections 5 and 6, I respond to Hasker's criticisms of the claim that there is indeed *excessive* gratuitous evil.⁸

2. HASKER'S ARGUMENT AGAINST (1C')

Hasker does not exactly argue that (1C') is false. Instead, he maintains that (1C') can reasonably be resisted.⁹ Hasker's argument proceeds in two stages.¹⁰

The first is about the *epistemic effects* of God's preventing all gratuitous evil. Hasker thinks that if God were to do this, we would find out. He offers two related reasons. He claims that one of God's chief purposes in creating rational beings is to bring them to knowledge of his nature. This knowledge, Hasker supposes, would surely include the proposition that God, due to his essential moral goodness, prevents all gratuitous suffering (1992, 39; 1997, 391; 2004b, 92; 2008, 197). Hasker also claims that if God were to prevent humans from gaining this knowledge, this would amount to a "pervasive policy of deception" (1992, 37) or a "massive disinformation campaign" (1992, 39). Hasker takes this to be morally objectionable; unworthy of God. Officially, this stage of Hasker's argument turns on what we would come to *know* about God's nature and policies, given theism, but it could also be expressed – more modestly – in terms what we would come to *reasonably believe*.¹¹

The second stage of Hasker's argument involves the *morally demotivating effects* of our learning that God prevents all gratuitous evil. In the literature on the problem of evil, a distinction is standardly drawn between *moral* evil and *natural* evil. Roughly, the former is wrongdoing perpetrated by moral agents, and the latter is pain and suffering not caused by such agents.¹² Hasker considers each in turn. If we came to believe that God prevents all gratuitous *moral* evil, he argues, our motivation to behave morally would be undermined. This is because "an important part of what leads human beings to attribute great significance to morality is the perception that pointless harm and suffering very often result from morally objectionable behaviour" (2004, 82). Absent that perception, we would be rather less likely to deem morality significant, and as a result, our motivation to act morally would be seriously impaired. After all, why refrain from permitting – or, for that matter, *causing* – moral evil when you are confident that God will ensure that any evil that results is *non-gratuitous*?

Hasker next imagines what would happen if we learned that God prevents all gratuitous *natural* evil. Hasker thinks that our motivation to prevent or minimize natural evil would be compromised too. In particular, our motivation to acquire or develop goods such as "knowledge, prudence, courage, foresight, cooperation, and compassion" in the struggle against natural evil would be reduced or eliminated (1992, 38-9; see also 2004, 88 and 2008, 193). This is because we would be confident that any natural evil that occurs would also be non-gratuitous – so why bother trying to minimize or prevent it?

Pulling all this together, we can say that Hasker thinks that God can allow gratuitous evil (of both kinds) in order to *prevent moral demotivation*. Hasker believes that God deems it very important for human beings to place a high priority on fulfilling moral obligations, and, in particular, for them to assume major responsibility for the welfare of their fellow human beings (2004, 82; 2008, 191). Hasker also thinks that God deems it very important that human beings respond to natural evil by acquiring and developing the goods mentioned above. Indeed, according to Hasker, the downsides of moral demotivation are worse than the downsides of God's allowing some gratuitous evil. If Hasker is correct, then God and gratuitous evil are compossible, in which case (1C') can be resisted.

Hasker's story can sound paradoxical, even self-defeating. If God must allow gratuitous suffering in order to keep our moral motivation from being undermined, then it seems that this suffering is no longer gratuitous: indeed, it is needed to bring about a greater good!¹³ But Hasker has a good response: the *class* of gratuitous suffering is not gratuitous: it's needed to prevent moral demotivation. Each individual *member* of that class, however, really is gratuitous.¹⁴

3. WOULD MORAL MOTIVATION BE UNDERMINED?

Elsewhere, I have considered several objections to Hasker's argument (Kraay 2019). I won't rehearse them all here. But I would like to offer some resistance to Hasker's most recent response to one such objection. Ultimately, however, I will grant – for the sake of discussion – that Hasker's argument against (1C') succeeds.

As we have seen, Hasker's central claim is that God permits gratuitous evil in order to prevent moral demotivation. The most direct way to resist this idea denies that moral demotivation would result from our coming to learn that God prevents all gratuitous evil.

Here is one way to make this case: one could argue that God would *intervene to block us* from making whatever inferences culminate in moral demotivation. On this view, God could (so to speak) have his cake and eat it too: God could prevent all gratuitous evil without undermining moral motivation. Hasker, however, thinks that theists should eschew this move, because its core idea (that God would block the relevant inferences) could then be harnessed into a *new* argument for atheism. Hasker, after all, thinks that there are *many theists* whose "moral motivation has in fact been adversely affected by the belief that in God's providence 'all is for the best'" (2019, 59). But, given the objection at hand, there should be no such people – and so their existence counts against theism.

Perhaps Hasker is right that theists would be prudent to eschew this move for this reason. But this is at most a strategic point; it is not a *philosophical* reason for rejecting the idea that God would block us from making the relevant inferences. Moreover, Hasker offers no argument against this idea: he merely says that there is "something troubling" about it, and that "it has at best only a certain mild plausibility going for it" (58, 59).

Here are two substantive philosophical considerations that one *might* offer against the idea that God would intervene to block us from making the inferences that lead to moral demotivation: one could argue that (a) this would constitute morally objectionable mental meddling on the part of God; or that (b) God's doing so could compromise some other important goals that God is reasonably taken to have. But, against (a), it could be responded that God's interference is justified in light of the immense benefits of the total eradication of gratuitous evil. And to make good an argument along the lines of (b), one would need to offer a plausible account of these goals, and of the ways in which they would be compromised by God's interference, and one would then need to claim that it would be overall better for God to achieve these goals than for him to intervene in this way to keep our moral motivations intact. This would be a highly complex response involving rather controversial large-scale value judgments. Given all this, I find it difficult to assess the overall probative force of this way to resist Hasker's central claim that God's prevention of gratuitous evil would be morally demotivating.

But here is another way to resist this claim. Several authors have held that the motivation to follow a *deontological* moral system would not be compromised by the knowledge or reasonable belief that God prevents all gratuitous evil.¹⁵ This challenge can be made vivid by appealing to actual theists, many of whom hold *both* that God prevents all gratuitous evil *and* that they are indeed subject to a deontological moral system.

Hasker's most recent response to this challenge denies that these counterexamples are significant. He says that "there are relatively few such individuals", whereas it "is evident" that "a great many individuals" experience moral demotivation upon accepting that God prevents all gratuitous evil (2019, 60). These assertions strike me as insufficient, in the absence of supporting

evidence, and so I think that the burden of proof remains on Hasker. In particular, I think he needs to offer reasons for thinking that, if God exists, moral motivation would in fact be undermined pervasively, and (just as importantly) that the downsides of this would be greater than the upsides of God's eradicating all gratuitous evil.

All that said, in what follows I will grant, for the sake of discussion, that Hasker's argument against (1C') succeeds.¹⁶ I aim to show that even if we grant this, his attempt to redeploy this strategy against (1D) fails. I will do this in section 5. In section 6, I will criticize two other objections due to Hasker. First, however, I set the stage in section 4 by introducing the final argument that Russell considers.

4. THE ARGUMENT FROM EXCESSIVE GRATUITOUS EVIL

The final argument that Russell considers begins with (1D), and runs as follows:

- (1D) If God exists, he would not allow excessive unnecessary suffering.
- (2D) There is excessive unnecessary suffering.
- Therefore,
- (3D) God does not exist.

Russell calls this "the strongest version of the argument from evil" (2018, 108). As I mentioned, he does not argue for (1D); he simply assumes that it is "is not open to criticism" (108). Most of Russell's paper is devoted to an indirect defence of (2D): he criticizes two well-known responses to it. The first is an epistemic objection to (2D) known as *skeptical theism*. The second is an instance of what has been called the *G.E. Moore shift*: it consists in offering reasons for the *denial* of (3D) that are held to be stronger than the reasons offered in *favour* of (3D). I will not consider these criticisms, or Russell's responses, here, since they are beyond the scope of this paper. I will focus, instead, on what Hasker has recently said about this sort of argument.

I have already argued that (1C) should be replaced with

- (1C') If God exists, there would be no gratuitous evil.

Parity suggests making (1D) explicitly about gratuitous evil. Here is my proposal:

- (1D') If God exists, it is false that gratuitous evil occurs *far in excess* of what must be permitted by God in order to achieve his goals.

One advantage of this formulation is that it encodes the core idea driving Hasker's rejection of (1C'): God's goals require the permission of *some* gratuitous evil. The argument for atheism that begins with this premise then continues as follows:

- (2D') Gratuitous evil occurs far in excess of what must be permitted by God in order to achieve his goals.

Therefore,
(3D') God does not exist.

Those who reject (1C'), like Russell and Hasker, hold that God can permit gratuitous evil. But, of course, to concede that God can allow *some* gratuitous evil is not to concede that God can allow *any amount* of gratuitous evil whatsoever. This distinction gives the final argument its purchase. (1D') asserts, in effect, that there is a limit to how much unnecessary evil God can allow.

This argument from *excessive* gratuitous evil has seen some discussion in the literature.¹⁷ Recently, William Hasker has urged that his case against (1C') can be deployed against (1D'), and he has also criticized (2D'). I now respond to these arguments.

5. HASKER'S CRITICISM OF (1D'): PANDEMICS

As we have seen, Hasker takes himself to have shown that it can be reasonable to think that God must allow gratuitous evil in order to prevent our moral motivation from being undermined. Hasker then imagines how someone who thinks this would reason, if she *also* thought that (1D') were true.¹⁸ Such a person, of course, believes that while God has to allow *some* gratuitous suffering in order to prevent moral demotivation, God should not permit *too much*. In other words, such a person thinks that there is an acceptable "zone" or "range" (2019, 65) of gratuitous evil, and that God will act as needed to keep gratuitous evil within it.

Hasker thinks that such a person will not typically experience any moral demotivation in light of her commitment to (1D'). After all, when she considers whether to undertake or refrain from some morally significant action, she can reason that her own choices have relatively modest effects in the grand scheme of things, and that, accordingly, they are very unlikely to affect whether the *overall* level of gratuitous evil is in the acceptable "zone" or "range".

But Hasker thinks that such a person will find her moral motivation undermined with respect to actions or omissions that are expected to have very significant, wide-ranging salutary effects. By way of example, Hasker imagines someone who is considering taking *demonstrably effective large-scale measures* to combat a pandemic. (Hasker's paper pre-dates the COVID-19 pandemic, but this ongoing tragedy has furnished us with vivid examples of such measures, such as the implementation of vaccination mandates, mask-wearing requirements, and quarantine rules across large populations.) Hasker thinks that someone considering taking such measures, who also believes (1D'), will reason as follows:

If I *don't* implement these measures, there will be a great deal of gratuitous evil as a result of the pandemic. But, then, given (1D'), God will have to ensure that there is much *less* gratuitous evil elsewhere or elsewhen "in the world-system" (66), in order to keep the level of gratuitous evil within the "acceptable" zone or range (65). On the other hand, if I *do* implement these measures, there will be far less gratuitous evil related to the pandemic. But,

then, given (1D'), God will have to permit much *more* gratuitous evil elsewhere or elsewhen in order for the level of gratuitous evil to remain acceptable.

Hasker thinks that a person who reasons like this will find her moral motivation (to respond to the pandemic with such measures) undermined. Accordingly, since God wants to prevent this demotivation, Hasker thinks it can be reasonable to hold, not just that God can permit gratuitous evil, but that God can permit *excessive* gratuitous evil too. Indeed, Hasker thinks, God *must* do both in order to prevent (this sort of) moral demotivation.

Here are two preliminary points about this argument. The first echoes something I said earlier about Hasker's argument against (1C'). This argument may sound paradoxical or self-defeating: if God needs 'excessive' gratuitous evil in order to prevent moral demotivation, then one might be tempted to think that it isn't gratuitous after all! But of course, Hasker can again respond by saying that while the existence of the *class* of 'excessive' gratuitous evil is needed in order to prevent moral demotivation with respect to large-scale actions, and so is non-gratuitous, each *member* of this class is indeed gratuitous.

Second, it may be a mistake for Hasker to construe this optimal level *globally*. Granting Hasker's argument against (1C'), the amount of gratuitous suffering that God must permit is more plausibly determined *locally*, so to speak, by features of particular circumstances or agents that are relevant to their moral (de)motivation.¹⁹

I now offer four replies to Hasker's argument. The first begins with the observation that the scope of Hasker's argument against (1D') is significantly narrower than the scope of his argument against (1C'). Against (1C'), Hasker urged that if we came to believe this claim, our moral motivation would be undermined *pervasively*. Hasker's new argument for moral demotivation, in contrast, targets only those actions or omissions that are reasonably expected to have *very* wide-ranging effects. Now, relatively few people ever have the opportunity to make such choices. And given this, even if Hasker is right that such people will find their moral motivation undermined, with respect to a certain narrow class of actions, one could argue that the benefits of permitting excessive gratuitous evil are outweighed by the costs, in which case his argument against (1D') fails.²⁰

The next three replies deny, in different ways, that moral demotivation for such actions will ensue from reasonably believing (1D'). Hasker doesn't clearly spell out how this moral demotivation is supposed to work, but I take it that he means something like this. The person realizes that whether or not she undertakes these large-scale measures to combat the pandemic, God will ensure that there is pretty well the same amount of gratuitous evil in "the world-system" (66). So, her actions will make no significant overall difference to the global amount of gratuitous evil – and *this* realization is supposed to be morally stultifying. "Why should I bother enacting these measures," the person thinks, "when God will ensure that the overall level of gratuitous evil is roughly the same no matter what?" The underlying thought appears to be that moral demotivation results from the realization that one's actions make **little or no** global difference. But this can be

resisted. Setting aside the particulars about gratuitous evil, a few moments' reflection on mind-boggling size of the actual world – which is extended immensely, *stupendously*, in both space and time – readily prompt the thought that our individual morally-relevant actions and omissions make **little or no** overall moral difference at the 'global' level. And yet, for most people, this realization is just not morally stultifying: we simply – and rightly – concern ourselves with making a *local* difference where we can. Likewise, it seems to me, the agent in Hasker's scenario can properly focus on reducing or eliminating what *local* suffering she can, even while she accepts that her actions make **little or no** difference to the overall *global* amount of gratuitous evil.

Third, here is an entirely different way to reject Hasker's claim that moral demotivation (for large-scale actions) will ensue from accepting (1D'). For all that his character knows or reasonably believes, her wide-ranging measures to combat the pandemic will *themselves* reduce the global amount of gratuitous suffering to an acceptable level, and so no increase in suffering elsewhere or elsewhen at all will be needed. Accordingly, she needn't believe that her actions will make no overall moral difference, and so her moral motivation needn't be undermined.

The final way to resist Hasker's claim about moral demotivation invokes a move we saw in section 3, above: perhaps God would *intervene to block* those specific individuals from making whatever inferences culminate in moral demotivation. Given that only very few individuals ever find themselves in a situation where their actions make such a large-scale difference, it's difficult to believe that God's intervening to prevent them from being morally demotivated on these occasions would either (a) constitute morally objectionable mental meddling on the part of God; or (b) compromise some other important goals that God is reasonably taken to have.

6. HASKER'S CRITICISMS OF (2D'): CLIMATE CHANGE AND CANCER

The most obvious way to resist (2D') is by endorsing some version of *skeptical theism*. But this move is not open to Hasker, since he is adamantly opposed to this view.²¹ (Russell, by the way, is likewise opposed.²²) Hasker instead offers two different criticisms of (2D'); I will respond to both.²³

Hasker calls his first criticism of (2D') "the it-has-to-get-really-bad problem" (2019, 63) and it again involves moral motivation. Here is the basic idea. Sometimes, alas, we are not motivated to solve a problem until things get *really bad*. Hasker's example is plausible: climate change. Tragically, it seems that many of us will fail to take adequate individual or collective action until things get really bad – perhaps even far worse than they already are now. Since Hasker thinks that God wants *us* to respond appropriately to climate change (rather than, say, intervening miraculously to cool the earth), God is justified in permitting an enormous amount of gratuitous evil that results from climate change. And this, in turn, counts against (2D').

I have two objections to this move. First, it is insufficient against (2D'). Suppose that everything Hasker says about climate change and our moral

motivation is correct. Suppose, further, that there are other contexts in which God really does have to let things get really bad before we are motivated to do something about it.²⁴ This does not suffice to show that (2D') is false or implausible. After all, there may well be *still other* contexts in which it is reasonable to suppose that excessive gratuitous evil occurs. So long as the defender of (2D') can find just *one context* in which *just one instance* of excessive gratuitous evil can reasonably be thought to occur, this premise can be maintained. Hasker's criticism does not – indeed it cannot – show that no such context and no such instance exists.²⁵

Second, Hasker is wrong to suggest that God simply has to wait for things to get really bad, so that our lacklustre moral motivation finally kicks in and we begin to address the problem. God has another – and better – option. To see this, consider some famous remarks by Hume's character (and presumptive mouthpiece) Philo:

In order to cure most of the ills of human life, I require not that man should have the wings of the eagle, the swiftness of the stag, the force of the ox, the arms of the lion, the scales of the crocodile or rhinoceros; much less do I demand the sagacity of an angel or cherubim. I am contented to take an increase in one single power or faculty of his soul. Let him be endowed with a *greater propensity to industry and labour; a more vigorous spring and activity of mind; a more constant bent to business and application* (*Dialogue XI*, 208, emphasis added).

Of course, one might object to Philo that giving humans an across-the-board boost in industriousness would be a mixed blessing: after all, malefactors as well as benefactors would then more vigorously pursue their ends.

Still, Philo's point can be adapted to the present context as follows: God could have given us – or could still now give us²⁶ – just a bit more *moral* motivation and *moral* fortitude, so that we would not so often wait until things get really bad. This would surely help the climate, and it would doubtless help in many other contexts too. And one could argue that even if there are costs of God's doing so, the benefit – the significant reduction of gratuitous evil – outweigh them.

Hasker calls his second criticism of (2D') the "too much intervention" objection. The basic idea is that the only way for God to prevent unnecessary suffering is to intervene miraculously with great frequency. But, Hasker thinks, doing so would undermine God's goal of allowing the natural world, for the most part, to "function naturally, with the various entities it contains operating according to their own inherent natures and capacities" (2019, 64). His example is cancer. He presumes that God could cure *most* individual cases of cancer without endangering God's "overall goals for the world and for human life", but he claims that if God were to cure *all* of them, this would constitute "too much intervention", since it would threaten the regularity, and hence intelligibility, of the natural order. (2019, 64).

I again have two objections, and the first echoes what I said above: this move is insufficient to resist (2D'). Even if Hasker is right that the *only* way that God could prevent excessive gratuitous suffering due to *cancer* is to intervene

massively, there may well be other contexts or situations in which it is reasonable to think that God *could* intervene to prevent excessive gratuitous evil *without* compromising the regularity or intelligibility of nature. The defender of (2D') just needs to find just one such context or situation to support this premise. Hasker's argument does not – indeed it cannot – show that there is *no* such context or situation.²⁷

The second objection, as before, comes from Hume's Philo:

Health and sickness, calm and tempest, with an infinite number of other accidents, whose causes are unknown and variable, have a great influence both on the fortunes of particular persons and on the prosperity of public societies; and indeed all human life, in a manner, depends on such accidents. A being, therefore, who knows the secret springs of the universe, might easily, by particular volitions, turn all these accidents to the good of mankind, and render the whole world happy, without discovering himself in any operation. A fleet, whose purposes were salutary to society, might always meet with a fair wind. Good princes enjoy sound health and long life. Persons born to power and authority, be framed with good tempers and virtuous dispositions. *A few* such events as these, regularly and wisely conducted, would change the face of the world; and yet would *no more seem to disturb the course of nature, or confound human conduct*, than the present economy of things ... (*Dialogue XI*, 206-7, emphasis added).

Philo's incisive point here is that it's too simplistic to suggest that a great number of miracles – enough to undermine the regularity or intelligibility of nature – are needed: a few key ones will do. Let's apply this point to Hasker's cancer example. Hasker thinks that “many millions per year of special divine interventions ... would be required to cure the many cases of cancer” in the United States alone (2019, 64). But Philo would (rightly) retort that this just is a mistake: all God needs to do is intervene in a few key situations – perhaps, for example, by giving greater insight or stamina to a few key cancer researchers on a few key occasions. Such interventions would be undetectable, Philo would say, and would surely not undermine the regularity or intelligibility of nature. And they would vastly reduce, if not entirely eliminate, the evils of cancer.²⁸ And, finally, even if a few more *ad hoc* miracles would be needed here and there, these wouldn't be enough to undermine the regularity or intelligibility of nature. Given all this, I think that Hasker's criticisms of (2D') fail.

CONCLUSION

Here is a brief synopsis of the dialectic so far – one that aims to itemize various points of agreement and disagreement between Russell, Hasker, and me. Russell and I agree that the arguments that begin with (1A) and (1B) fail, and I have argued that those that begin with variants (1A') and (1B') likewise fail. I am sure that Hasker would concur. Russell rejects (1C), and Hasker rejects the closely related

(1C'). Russell and Hasker both endorse Peter van Inwagen's 'no minimum' argument, while Hasker also offers his argument from moral demotivation. I did not engage with the former, having criticized it elsewhere. I have mixed feelings about the latter. I think it has been underappreciated in the literature, and I find it intuitively plausible, but I'm not convinced that it works. I offered some resistance to Hasker's most recent defence of it against an important objection, but I ultimately granted – for the sake of discussion – that it succeeds.

This brought us to the very important arguments that begin with (1D) and (1D'): the arguments from *excessive* suffering and *excessive* gratuitous evil, respectively. Russell, Hasker and I agree that these arguments constitute formidable problems for theism. Russell simply assumes that (1D) is not open to criticism. I think this that is a mistake: Hasker, after all, is a very able philosopher, and he has criticized the closely related claim (1D'). But I don't think that Hasker's criticisms of this premise succeed. The next premise holds that excessive evil exists. Hasker has criticized this too, but again, I don't think his criticisms succeed. Meanwhile, Russell spends most of his paper indirectly defending the closely related empirical premise that excessive unnecessary suffering exists, and I am sure that his arguments apply equally to the claim that excessive gratuitous evil exists. I did not engage these arguments, since they were outside the scope of this paper. Still, a complete assessment of the argument from excessive unnecessary suffering and the argument from excessive gratuitous evil will require doing just this. But that will have to wait for another day.²⁹

REFERENCES

- Adams, M.M. (1999) *Horrendous Evils and the Goodness of God*. Ithaca, NY: Cornell University Press.
- Almeida, M. (2020) "On Necessary Gratuitous Evils." *European Journal for Philosophy of Religion* 12: 117-135.
- Alston, W. (1991) "The Inductive Argument from Evil and the Human Cognitive Condition." *Philosophical Perspectives* 5: 29-67.
- Chrzan, K. (1994) "Necessary Gratuitous Evil: An Oxymoron Revisited." *Faith and Philosophy* 11: 134-137.
- Fischer, J.M. and Tognazzini, N. (2007) "Exploring Evil and Philosophical Failure: A Critical Notice of Peter van Inwagen's *The Problem of Evil*." *Faith and Philosophy* 24: 458-474.
- Hasker, W. (1992) "The Necessity of Gratuitous Evil." *Faith and Philosophy* 9: 23-44.
- _____ (1997) "O'Connor on Gratuitous Natural Evil." *Faith and Philosophy* 14: 388-394.
- _____ (2004) "Can God Permit 'Just Enough' Evil?" In *Providence, Evil, and The Openness of God*. New York: Routledge, 81-94.
- _____ (2008) *The Triumph of God over Evil: Theodicy for a World of Suffering*. Downer's Grove: InterVarsity Press.
- _____ (2010a) "Defining 'Gratuitous Evil': A Response to Rhoda." *Religious Studies* 46: 303-309.
- _____ (2010b) "All Too Skeptical Theism." *International Journal for Philosophy of Religion* 68: 15-29.
- _____ (2019) "God and Gratuitous Evil: A Response to Klaas Kraay." *Oxford Studies in Philosophy of Religion* 9: 54-67.
- Himma, K. (2011) Review of Hasker's *The Triumph of God over Evil*. *Religious Studies* 47: 129-140.
- Howard-Snyder, D. and Howard-Snyder, F. (1999) "Is Theism Compatible with Gratuitous Evil?" *American Philosophical Quarterly* 36: 115-130.

- Hume, D. (1776) *Dialogues Concerning Natural Religion*. Edited by Norman Kemp Smith, 1948. New York: Social Science Publishers.
- Jordan, J. (2004) "Divine Love and Human Suffering." *International Journal for Philosophy of Religion* 56: 169-178.
- Keller, J. (1989) "The Problem of Evil and the Attributes of God." *Faith and Philosophy* 26: 155-171.
- _____ (2007) *Problems of Evil and the Power of God*. Aldershot: Ashgate, Chapter 2.
- Kraay, K. (2014) "Peter van Inwagen on Gratuitous Evil." *Religious Studies* 50: 217-234.
- _____ (2016) "God and Gratuitous Evil (Part II)." *Philosophy Compass* 11: 913-922.
- _____ (2019) "Theism, Pro-Theism, Hasker, and Gratuitous Evil." *Oxford Studies in Philosophy of Religion* 9: 31-53.
- Mandelkern, M. (2020) "A Counterexample to Modus Ponenses." *The Journal of Philosophy* 117: 315-331.
- Mawson, T.J. (2011) "Theodical Individualism." *European Journal for Philosophy of Religion* 3: 139-159.
- Mooney, J. (2019) "How to Solve the Problem of Evil: A Deontological Strategy." *Faith and Philosophy* 36: 442-462.
- O'Connor, D. (1995) "Hasker on Gratuitous Natural Evil." *Faith and Philosophy* 12: 380-392.
- _____ (1998) *God and Inscrutable Evil: In Defense of Theism and Atheism*. Lanham: Rowman and Littlefield, Chapter 3.
- Stump, E. (2010) *Wandering In Darkness: Narrative and the Problem of Suffering*. Oxford: Oxford University Press.
- Reichenbach, B. (2010) Review of Hasker's *The Triumph of God Over Evil* and Keller's *Problems of Evil and the Power of God*. *Faith and Philosophy* 27: 212-218.
- Rhoda, A. (2010) "Gratuitous Evil and Divine Providence." *Religious Studies* 46: 281-302.

- Rowe, W. (1979) "The Problem of Evil and some Varieties of Atheism." *American Philosophical Quarterly* 16: 335-341.
- _____ (1991) "Ruminations about Evil." *Philosophical Perspectives* 5: 69-88.
- Rowe, W.L. (1996) "William Alston on the Problem of Evil." In *The Rationality of Belief and the Plurality of Faiths*. Edited by Tom Senior, 71-93. Cornell: Cornell University Press.
- Russell, B. (2004) "The Problem of Evil: Why Is There So Much Suffering?" In *Introduction to Philosophy: Classical and Contemporary Readings*. Edited by Louis P. Pojman, 207-13. New York: Oxford University Press.
- _____ (2018) "The Problem of Evil and Replies to Some Important Responses." *European Journal for Philosophy of Religion* 10: 105-131.
- Tooley, M. (1991) "The Argument From Evil." *Philosophical Perspectives* 5: 88-134.
- van Inwagen, P. (2001) "The Argument from Particular Horrendous Evils." *American Catholic Philosophical Quarterly* 74: 65-80.
- _____ (2006) *The Problem of Evil*. Oxford: Oxford University Press.

NOTES

¹ For those concerned about the validity of modus tollens involving a conditional premise whose consequent is in the subjunctive mood, the consequents of premises (1A)-(1D) can simply be re-written in the indicative mood. For a discussion of modus ponens with subjunctive conditionals, see Mandelkern (2020).

² The good could, of course, be the prevention of even worse suffering.

³ Further constraints can be, and have been, proposed. An important – and controversial – one holds that the goods for the sake of which the suffering is permitted must *benefit the sufferer*. Requirements in this vein are offered by Tooley (1991), Rowe (1996), Adams, M.M. (1999), and Stump (1990, 2010). For criticisms, see Jordan (2004) and Mawson (2011).

⁴ Russell writes about bringing about a greater good *or* preventing a greater evil, but I subsume the latter under the former for simplicity. Also, if conditions (i) and (ii) are thought to encode implausible expectations regarding *(fore)knowledge*, they could be expressed, more modestly, in terms of what the good agents *reasonably expect*.

⁵ See, for example, van Inwagen (2006), Hasker (1992, 1997, 2004, 2008, 2019), Mooney (2019), and Almeida (2020).

⁶ See Russell (2018, 107-8). Interestingly, Hasker (2019) appears to endorse van Inwagen's argument as well, while also defending his own independent argument against (1C).

⁷ I have replaced "could antecedently know" in Hasker's formulation with "antecedently knows", since if God could know *p*, God knows *p*. Hasker here means to refine a definition of gratuitous evil offered by Rhoda (2010, 287-289), which Rhoda, in turn, takes to be an improvement over William Rowe's (1979) conception. Other critics of Rowe's account of gratuitousness include William Alston (1991, 33-34) and van Inwagen (1991, 164, note 11). Rhoda's definition appears to be inspired by certain remarks by van Inwagen (2001, 69; 2006, 97).

⁸ Hasker does not explicitly address Russell here; his objections are found in his critical discussion of Kraay (2019). And Russell does not consider Hasker's paper, for the excellent reason that it appeared after his own paper was published.

⁹ Hasker says that the claim expressed by (1C') "...should be rejected by theists, since it comes into conflict with other, better-entrenched elements of the theistic worldview" (2004, 81; and see 2008, 189). It's worth adding that *non-theists* could also resist (1C') on Haskerian grounds. Such individuals would hold that if theism were true, then these "better-entrenched elements of the theistic worldview" would be more plausible than (1C').

¹⁰ My exposition of Hasker here is drawn from Kraay (2016) and Kraay (2019).

¹¹ Hasker does not insist that *all* people would come to know this, so perhaps his argument should be further restricted to theists (or, more modestly, to most theists).

¹² These definitions may be a bit too broad: perhaps it's not the case that *all* moral wrongdoing counts as evil, and perhaps it's not the case that *all* pain and suffering due to natural processes counts as evil. There's no need to hash this out here, though, since nothing turns on this for Hasker's purposes, or for mine.

¹³ Rowe (1991), Chrzan (1994), and Howard-Snyder and Howard-Snyder (1999) all press this criticism against Hasker.

¹⁴ Hasker (1992, 33; 2004, 89; and 2008, 195). An anonymous referee objects as follows:

Suppose some theodicy says that God needs to throw hardships into people's way in order to promote character development, say. The idea is usually not that God needs this particular broken leg or broken heart to achieve this purpose, but that God needs some hardships. In other words, if some evil is part of a class which is such that the permission of some member of that class is necessary for a greater good, then the evil is *not* gratuitous (emphasis added).

The referee's phrasing in the antecedent of the final sentence has the potential to mislead: it's not that *some particular* member of the class is necessary for a greater good (which would make the evil nongratuitous); instead, it's just that, *ex hypothesi*, the class must *not be empty*. In other words, in order to prevent moral demotivation, on Hasker's view, there must be some instances of gratuitous evil. Why are they gratuitous? Because, given Hasker's definition of this term, God knows that it's certain or likely that he can prevent

any of them in a way that would make the world overall better. But God can't prevent *all* of them from occurring, because, according to Hasker, this would lead to an overall worse outcome: widespread moral demotivation. He offers the analogy of applause after a fine musical performance (2008, 195). God could prevent any particular audience member from clapping without ending the applause, which shows that no particular individual instance is needed for the phenomenon to occur – each instance, so to speak, is gratuitous with respect to the phenomenon of applause. But, of course, if God were to prevent *all* the audience members from clapping, there would be no applause.

¹⁵ See Keller (2007, 13), Rhoda (2010, 291), and Himma (2011, 132).

¹⁶ An anonymous referee objects to Hasker's argument from moral demotivation in the following way:

... if God is justified in allowing Adam to kick Eve for the great good for free will etc, then Adam still has a reason to avoid kicking Eve. The world in which he freely refrains from kicking Eve is better than one in which he freely does kick Eve. If he were to so refrain, God would not need to cause some other person to kick his wife to make up for it. God's perfect goodness is compatible with a wide range of actual amounts of evil. So, humans still have motivations to avoid or prevent evil.

Here are a few comments in response. First, nothing in Hasker's argument requires that if Adam refrains from kicking Eve, God will *cause* someone else to kick his wife – or even that God will *permit* someone else to kick his wife. Second, the referee makes it sound like Adam's motivation must flow from the recognition that, *ceteris paribus*, refraining from kicking Eve will make the world overall better. But this needn't be the case. His erstwhile motivation could be grounded in more local consequentialist considerations, or in deontological considerations, etc. Third, and most importantly, this comment seems to miss the point of Hasker's thought experiment. Hasker thinks that if Adam knows or reasonably believes that God will ensure that kicking Eve is *non-gratuitous*, then his motivation to *refrain* from kicking her will be undermined – and he might even take himself to have a reason *to* kick her.

¹⁷ This argument is discussed briefly by Keller (1989, 163-6; 2007, 14-16); Rowe (1991, 88, note 20); Chrzan (1994, 135); O'Connor (1995, 391; 1998, 69-70); and Reichenbach (2010, 214). It is discussed in a bit more detail by Russell (2004), Fischer and Tognazzini (2007), Kraay (2019), and, of course, Hasker (2019).

¹⁸ Hasker presumably thinks that if (1D') were true, we would come to know (or at least reasonably believe) it. After all, this is what he argued about (1C').

¹⁹ A related point can be found in Keller (2007, 13).

²⁰ An anonymous referee notes that while there may only be few individuals in such a position, their moral demotivation might have significant disvalue, given the wide-ranging effects of their actions or omissions. This point is presumably meant to count against my suggestion that the costs of God's permitting excessive gratuitous evil might well outweigh the benefits. I agree that the moral demotivation of such individuals would be particularly disvaluable, but I'm not sure that this is enough to tilt the scale in Hasker's favour. Presumably only theists would be so demotivated, and, moreover, only those theists who

reason along Haskerian lines – and this might be quite a small number. Now, their moral demotivation would be bad, to be sure, and it would lead to evil. But on the other hand, in the imagined scenario, God will work indefatigably to ensure that the overall level of gratuitous evil is in the acceptable ‘zone’ or ‘range’, so as to prevent widespread moral demotivation – and this is surely a very good thing. Ultimately, it seems far from obvious that Hasker is entitled to think that it is clearly better to permit excessive gratuitous evil. But for those who remain unconvinced, I next offer three different responses to Hasker’s argument.

²¹ Hasker thinks that skeptical theism ineluctably leads to an untenable skepticism about induction (2010b, 19-21), and to a repellent moral skepticism (2004a, 51-52; 2010b, 21-27), and that it courts incoherence (2010b, 27-29).

²² See Russell (2018, 111-114), where he similarly argues that the same considerations offered to support skeptical theism equally license absurd forms of skepticism, like Young Earthism (the view that the earth was created very recently).

²³ Hasker does not assert that his criticisms of the claim expressed by (2D’) are decisive; on the contrary, he thinks that this claim “certainly has not been disproved”, and ultimately concedes that he does not know whether it is true or false (2019, 65). Still, he offers them as considerations against (2D’), and so it is fair game to object to them.

²⁴ Hasker says that “additional examples can easily be found” (64).

²⁵ Hasker might retort that his move should be *presumed* to apply to all other contexts – but he offers no such argument, and it’s difficult to see how one might go. For one thing, there’s just no reason to suppose that *every* context is such that things have to ‘get really bad’ in order for our moral motivation to kick in. Climate change is a special case. For one thing, many people, due to their circumstances, are significantly shielded from its most deleterious effects. Moreover, since it is a large and complex phenomenon, it can be very difficult for people to understand how individual and collective choices, over many generations, have led to the current crisis. In contrast, the defender of (2D’) can appeal to other contexts in which it seems highly plausible that gratuitous evil occurs far in excess of what God needs to prevent moral demotivation. Consider, for example, the enormous gratuitous suffering of innocents caused by warfare, racism, sexual exploitation, and so forth. It is extremely difficult to believe that God needs to permit *all* of this suffering in order to prevent moral demotivation.

²⁶ I add this qualification for those who think, for example, that the Fall damaged our moral motivation.

²⁷ An anonymous referee presses me for an example. I reply that Philo’s examples in the passage I quote next seem entirely apposite. The Humean recipe for generating further examples is simple: God’s interventions need to be high-impact, but sufficiently infrequent or subtle, so that neither the regularity nor intelligibility of nature is compromised.

²⁸ In reply, one might surmise that God is already doing everything he can in this regard – but that is not likely to persuade Philo, Russell, or other defenders of this sort of argument.

²⁹ I presented ancestors of this paper at the Vrije Universiteit Amsterdam (May 2020); to the Toronto Philosophy of Religion Work-in-Progress Group Mini-Conference on God and Gratuitous Evil (May 2023); at the 45th Anniversary Meeting of the Society of Christian Philosophers (September 2023) and at the Central Division Meeting of the American Philosophical Association (February 2024). Thanks to all my interlocutors on those occasions, and, in particular, to my commentators at the last two events: Libby Regnerus and Emily McCarty, respectively. Special thanks are due to William Hasker, Bruce Russell, Dean Zimmerman, and to two anonymous referees for OSPR, for their insightful feedback.