

Exploring Evil and Philosophical Failure: A Critical Notice of Peter van Inwagen's *The Problem of Evil*¹

Penultimate Draft
(final version in *Faith and Philosophy*, Vol. 24, No. 4, (October 2007): 458-474)

JOHN MARTIN FISCHER AND NEAL A. TOGNAZZINI

Introduction

This book presents somewhat revised versions of the Gifford Lectures, delivered by Peter van Inwagen at the University of St. Andrews in 2003. The lectures provide a systematic response to various versions of the problem of evil. The book is quite simply delightful to read. It is beautifully written, and full of ingenious, subtle, and insightful argumentation. The book displays Peter van Inwagen's signature combination of elegant and plain prose with philosophical sophistication and rigor. Anyone interested in the problem of evil should read this book. In our view, it contains much good sense, not a little wry (and slightly curmudgeonly) humor, and some novel insights into this central problem of philosophy of religion. In what follows, we shall begin by discussing van Inwagen's fascinating and suggestive views about philosophical methodology. We shall then turn to the global and local arguments from evil. We argue that although van Inwagen's responses to the global and local arguments from evil may succeed, his discussion points us to a third type of argument from evil (in addition to the global and local varieties) against which neither of his responses to the arguments he considers will work. The upshot will be that van Inwagen has yet to put the problem of evil to rest, since there is an important argument from evil he has yet to address.

I. Philosophical Methodology and Philosophical Failure

Chapter 3 is entitled "Philosophical Failure," but it is really about philosophical methodology in general, or perhaps more precisely, about how best to assess and also to develop philosophical argumentation. Van Inwagen begins by rejecting a model of

philosophical argumentation quite similar to what Robert Nozick has called “coercive philosophy”; on this model, one seeks to provide knockdown arguments from indisputable premises. As regards this sort of model, van Inwagen points out (p. 37) that Nozick said that when he was young he thought that a philosophical argument is adequate only if anyone who understood the premises but did not accept the conclusion would die! (One imagines the brain exploding.)

Van Inwagen goes on to consider a second model. On this view, philosophical argumentation can be thought of as a kind of debate between two parties who have opposite views about the issue under consideration, where the goal of each is to convince the other to give up his position and adopt the competing view. (A weaker requirement would be to convince the other party to the debate to switch from accepting the competing view to agnosticism; van Inwagen doesn’t explicitly consider this possibility. An even weaker requirement – or family of requirements – would be to get the other party to decrease to some degree his confidence in his view. These models might be fruitful to consider.) So suppose the issue is whether causal determinism is compatible with moral responsibility. Here we are to envisage an idealized debate (some of the conditions of idealization are discussed on pp. 42-3) between a compatibilist and an incompatibilist, where the compatibilist seeks to convince the incompatibilist to adopt compatibilism, and the incompatibilist seeks to convince the compatibilist to adopt incompatibilism. The argument offered (say) by the compatibilist is deemed successful only if the incompatibilist is persuaded to adopt compatibilism.

Van Inwagen rejects this second model for the same reason he rejects the first, saying that it places the bar too high: “I very much doubt whether any argument, or any set of independent arguments, for any substantive philosophical conclusion has the power to turn a determined opponent of that conclusion, however rational, into an adherent of that conclusion.” (p. 43)

Van Inwagen prefers a third model, according to which we understand philosophical argumentation as like an idealized debate between proponents of competing positions, where the goal is not to convert the other debater, but to convince an idealized “agnostic” – a person who is “neutral” in the sense that he has no particular antecedent inclination to accept the relevant position. Here the two debaters address a third party, rather than each other; and the goal is not “conversion”, but securing “conviction”, as it were. Van Inwagen contends that this weaker model is more reasonable, and that it has some advantages with respect to understanding issues about begging the question and the burden of proof.

We find van Inwagen’s proposal highly attractive, and we agree with him that it is preferable to the first two. In previous work, one of us (Fischer) has suggested something similar. That is, Fischer has suggested that (for example) a compatibilist understand his arguments as directed to a fair-minded, reasonable person without a prior commitment to either compatibilism or incompatibilism.² (In recent work Haji and McKenna and also Mele hold a similar view.³) In certain debates one often reads or hears the following sort of reply to a proposed argument or consideration: “but an incompatibilist would reject that...” It is as if a compatibilist may only present considerations and arguments that would not be rejected by someone with an antecedent and iron-clad commitment to incompatibilism! But this really sets the bar too high; clearly, if we accepted this sort of requirement, no philosophical progress could be made with respect to substantive and contentious philosophical issues. We think it is quite helpful to point out that the arguments and considerations should be addressed to a neutral agnostic, rather than a committed opponent. As van Inwagen points out, adopting this model might well result in different views about what is argumentatively permissible, as opposed to question-begging or otherwise illicit, and where the burden of proof lies.

Van Inwagen goes on to argue that even on the somewhat more permissive third model, not only is the atheist’s argument based on evil a failure, but all arguments for substantive, interesting philosophical theses will turn out to be failures. We wonder why van

Inwagen thinks it is a fatal criticism of the first two models that they would have it that no philosophical arguments are successes, but not a problem for his favored model that it also has this implication. This is at least a *prima facie* problem – or puzzle – for van Inwagen. Perhaps van Inwagen believes that the first two models set the bar unreasonably high so that it would be unreasonably hard to meet their criteria for success, whereas the third model sets the bar at just the right height – but that even so no philosophical argument for an interesting, substantive thesis is successful. But it does seem at least to vitiate the interest of van Inwagen’s claim – that the various versions of the problem of evil are failures as philosophical arguments – that all philosophical arguments turn out to be failures (on his account).⁴ So perhaps it would be worth considering another model (or family of models), by reference to which it would turn out that some (but not all) philosophical arguments for the relevant sorts of theses are indeed successful (or at least successful to certain degrees).

Consider a model that adopts the assumptions of van Inwagen’s third model – we have a debate addressed to a neutral agnostic (with respect to the view under dispute). But here success does not require that the neutral agnostic actually adopt the relevant view; what is required is simply that the agnostic’s inclination to adopt the thesis increases (perhaps to some threshold amount). Here we can again imagine a family of different views; the weakest view would be that success simply requires some increase in the agnostic’s inclination to accept the relevant view, whereas stronger views would require a range of more significant increases in the strength of the inclination to accept the view in question. On this sort of model, we presumably could say that some (but certainly not all) philosophical arguments for contentious, substantive positions are successful.

Van Inwagen is aware of the possibility of models similar to the one we have suggested above; indeed, he says:

I would suppose that most real agnostics, most actual people who do profess and call themselves agnostics, are not neutral agnostics. Most agnostics I have discussed these matters with think that it’s pretty improbable that there’s a God. Their relation to the proposition that God exists is very much like my relation to the proposition that there are intelligent non-human beings inhabiting some planet within 10,000 light-years of the Earth. And this

consideration suggests a possible objection to my definition of philosophical success. Call those agnostics who think that it's very improbable that there is a God weighted agnostics. An argument for the non-existence of God, the argument from evil for example, might be a failure by my criterion because it lacked the power to transform ideal (and hence neutral) agnostics into atheists. But it might, consistently with that, have the power to transform neutral agnostics into weighted agnostics. If it does, isn't it rather hard on it to call it a failure? (p. 50)

Van Inwagen says that he would not object to revising his model to allow for success in the case described (in which the neutral agnostic becomes a weighted agnostic). But he goes on to contend that if the considerations he invokes in seeking to show that the argument from evil is incapable of turning neutral agnostics into atheists are persuasive, they would be equally persuasive in showing that the argument from evil is incapable of turning neutral agnostics into weighted agnostics. (p. 51)

But note that even van Inwagen's permitted revision is significantly different from the model we sketched above, insofar as our model does not require for success that a neutral agnostic become a weighted agnostic; after all, a weighted agnostic, according to van Inwagen, believes that it is "very improbable" that God does not exist. All that is required on our suggested model is that a neutral agnostic become inclined to some degree (which may fall considerably short of the threshold posited by van Inwagen) to accept the relevant position (say, atheism). Given this important difference, even if van Inwagen's conditional is true (that if the argument from evil is incapable of transforming a neutral agnostic into an atheist, then it is also incapable of transforming a neutral agnostic into a weighted agnostic), it would not follow that if the argument from evil is incapable of transforming a neutral agnostic into an atheist, then it is also incapable of increasing the inclination of a neutral agnostic to accept atheism (to the relevant degree).

Why exactly would we want a model by reference to which we could say that some (but not all) philosophical arguments for interesting, substantive theses are actually successful (as opposed to "close to successful")? We suppose this may depend to some extent on one's temperament or philosophical personality. Some philosophers might well think that it is accurate to describe all philosophical arguments for substantive, contentious

theses as failures, although perhaps they may invoke a model like ours to say that some arguments are “better” than others, or come closer to success than others, or perhaps are “partial” successes (to one degree or another). Others will be a bit less dour about the possibility of philosophical success. One might think that the interesting philosophical issues are so incredibly difficult, we should consider ourselves successful if we could change the inclination of an idealized neutral agnostic even a little bit. (And note that this change is in the context of a debate where the proponent of the opposing view can mount a vigorous argument of his own.)

There are of course various problems for the very sketchy proposal we have made. We simply note that it applies most naturally to highly contentious philosophical theses, where it is plausible that an ideal agnostic would in fact be neutral. If one considers theses which are in dispute but whose objective probability of truth is relatively high, then even a very uninspiring argument would presumably change (significantly) the inclination of an idealized unweighted agnostic. This suggests that perhaps in general the requirement should be that the argument move an ideal agnostic who has an inclination to accept the relevant view that matches (in a suitable sense) the objective probability of truth of the relevant view. But we will simply follow van Inwagen in noting this problem without addressing it in detail (p. 49).

There is no doubt that van Inwagen provides us with a striking and illuminating model of philosophical argumentation. We have suggested some possible revisions to this essentially “forensic” approach. It should perhaps be noted that this family of models might capture only part of the picture here.⁵ That is, it might be that philosophical argumentation is more heterogeneous than what is suggested by the forensic model (or even the family of such models) sketched above. For example, perhaps some philosophical arguments are “successful” in virtue of making us (or many of us – or some of us) see a certain debate in a different way – as structured in a different way, or as requiring different presuppositions from what we had antecedently believed. Or maybe an argument can be successful by

helping us to tease out certain distinctions, or clarify certain concepts, or encouraging us to think in new directions about an old problem. We suppose that an argument could be successful simply in virtue of helping us to see the force of some puzzle or problem. It is not clear that it is helpful to seek to explicate all philosophical argumentation by reference to one model or family of models – to do so might seem to force a disparate range of phenomena into a Procrustean bed.⁶

II. The Global and Local Arguments from Evil

Let's move on to consider van Inwagen's treatment of the problem of evil in particular. In his book he departs from the usual taxonomy that distinguishes the logical threat from the evidential threat and replaces it with one that distinguishes between the global argument from evil and various local arguments from evil. Although it isn't altogether clear why van Inwagen finds the logical/evidential distinction useless, he makes it clear that he does find it useless (p. 8).⁷ We agree that the distinction he uses instead is indeed useful. Roughly speaking, the global argument from evil purports to disprove the existence of God by pointing out that there is much evil in the world, whereas any particular local argument from evil purports to disprove the existence of God by pointing out that our world contains a particular instance (and most often a particularly egregious instance) of evil. We can regiment these rough characterizations as follows:

The Global Argument from Evil (p. 56)

- (1) We find vast amounts of truly horrendous evil in the world.
- (2) If there were a God, we should not find vast amounts of truly horrendous evil in the world.
- (3) Therefore, there is no God.

The Local Argument from Evil (p. 97-98)⁸

- (1) A particular horror, H, in fact occurred.⁹
- (2) If H had not occurred, the world would be no worse than it is.
- (3) If a morally perfect creator could have left a certain horror out of the world he created, and if the world he created would have been no worse if that horror had been left out of it than it would have been if it had included that horror, then the morally perfect creator would have left the horror out of the world he created if he had been able to do so.
- (4) If an omnipotent being created the world, then he was able to leave H out of the world.

- (5) So, if an omnipotent and morally perfect being created the world, then he would have left H out of the world. (by 2, 3, and 4)
- (6) Therefore, there is no omnipotent, morally perfect creator. That is, there is no God. (by 1 and 5)¹⁰

Interestingly, as van Inwagen points out, a satisfactory response to the global argument will not automatically be a satisfactory response to any particular local argument. To see this, note that even if one is able to tell a plausible story according to which God is blameless for allowing vast amounts of evil, this story will not automatically double as a plausible story according to which God is blameless for allowing a particular horror H to occur. Perhaps the world had to contain vast amounts of evil. But did it have to contain H? (pp. 8-9) Indeed, van Inwagen responds to the global argument in a much different way than to the local argument. In response to the former, he offers a particular version of the well-known free-will defense to cast doubt on premise (2).¹¹ His response to the local argument is particularly novel. He denies premise (3) by appealing to considerations of vagueness.

III. The Global Argument from Evil

Although we will focus our discussion on the local argument from evil, it will be useful to give a brief overview of van Inwagen's response to the global argument from evil so as to better situate our criticism. Recall that the global argument from evil starts from the undeniable fact that there is a vast amount of evil in the world.¹² It then asserts that there should not be a vast amount of evil in the world if God exists, and ends with the atheistic conclusion. Many philosophers have addressed this general sort of argument from evil by appealing to considerations of free will, and van Inwagen's response belongs to this tradition.

Van Inwagen offers a defense against this argument, which is to be distinguished from a theodicy. Whereas a theodicy is a story that one thinks expresses the truth about why God allows a vast amount of evil, a defense is a more modest project. The proponent of a defense need not actually believe that the story he tells expresses the truth of the matter. Rather, the proponent of a defense needs only to tell a plausible and epistemically possible

(i.e., true for all anyone knows) story that explains why God would allow a vast amount of evil. Van Inwagen says that the reaction he hopes to elicit in an ideal agnostic with his defense is the following: “Given that God exists, the rest of the story might well be true. I can’t see any reason to rule it out.” (p. 66) Such a reaction is tantamount to an admission that for all one knows, one of the premises of the global argument from evil is false. And for the ideal agnostic to admit this much is, as we have seen, enough to conclude that the global argument from evil is a failure.

So van Inwagen tells us a story: God created the world and “guided the course of evolution so as eventually to produce certain very clever primates, the immediate predecessors of *Homo sapiens*.” (p. 85) God miraculously raised these primates to rationality and gave them the gift of free will, not least of all because free will is necessary for love. But these creatures abused their free will and separated themselves from God. As van Inwagen says, “The result was horrific.” (p. 86) Gradually more and more evil came into the world as a result of this separation from God. A god who was merely merciful would have just wiped out human beings altogether, to put us out of our misery. But God is more than merciful; he is also loving. So God came up with a plan for our salvation that includes this essential feature: “Its object is to bring it about that human beings once more love God.” (p. 87) But since love requires free will, humans must freely choose to love God once more. And humans will only be motivated to freely turn to God if they realize what it means to be separated from God. So God must leave in place a vast amount of evil. If the amount of evil in the world were less than vast, so to speak, God’s plan of atonement would be frustrated.

There is much more to this story in van Inwagen’s book, and it is certainly worthy of consideration, as it is a considerably expanded and sophisticated version of the traditional free-will defense. But we’ve told enough of it to serve our purposes. The global argument from evil is a failure because, upon hearing this story, an ideal agnostic would say to herself, “Given that God exists, this story might well be true. And thus premise (2) of the global argument from evil might well be false.” In the next two sections we will see how van

Inwagen responds to the local argument from evil, and how his response paves the way to a third, and equally troubling, argument from evil.

IV. The Local Argument from Evil

Recall that according to the local argument from evil, it is the occurrence of some particular horror, H, that gets us to the conclusion that God doesn't exist. In philosophical treatments of the local argument, for example, there has been much discussion of William Rowe's example of a fawn that died an agonizing death in a forest fire that was not caused, in any way, by human beings and whose death is not observed by any human and leaves no trace.¹³ The following quotation provides us with a concise explication of van Inwagen's strategy of response, although it will take some work to unpack:

Even if no good came of [H], the occurrence of that event does not tell against the existence of an omnipotent, morally perfect being; for it may be that the omnipotent morally perfect Creator of the world was morally required to draw a morally arbitrary line through the set of threatened evils, and that [H] fell on the 'actuality' side of the particular line he chose. (p. 124-5)¹⁴

The basic idea here is that even if it is true that God could have left H out of the world without thereby making the world a worse place than it actually is, it doesn't follow that he would have left H out, because he had to draw some line to separate the horrors that made it into the world from those that didn't, and where he drew the line was arbitrary. And if the place at which the line gets drawn is arbitrary, then it's inappropriate to blame the person who draws the line on the grounds that this particular horror falls on this side of the line, rather than on that side. He had to draw it somewhere, and no better there than here. But let's unpack this some more.

According to van Inwagen, the best way to respond to any particular local argument from evil is to reject the moral principle that underlies premise (3). Recall premise (3) from above:

- (3) If a morally perfect creator could have left a certain horror out of the world he created, and if the world he created would have been no worse if that horror had been left out of it than it would have been if it

had included that horror, then the morally perfect creator would have left the horror out of the world he created if he had been able to do so.

Van Inwagen says that there is only “one moral principle that it would be plausible to appeal to in defense of premise (3)”, and he states it as follows:

- (P) If one is in a position to prevent some evil, one should not allow that evil to occur – not unless allowing it to occur would result in some good that would outweigh it or preventing it would result in some other evil at least as bad. (p. 100)

The problem, however, is that (P) is surely false. Van Inwagen first presents a simple counterexample to the principle, and then goes on to extract even deeper problems that the principle leads to. It will be helpful for our purposes to go through both of these steps. Van Inwagen presents the following counterexample to (P):

Suppose you are an official who has the power to release anyone from prison at any time. Blodgett has been sentenced to ten years in prison for felonious assault. His sentence is nearing its end, and he petitions you to release him from prison a day early. Should you? Well, the principle says so. A day spent in prison is an evil – if you don’t think so, I invite you to spend a day in prison...Let’s suppose that the only good that results from someone’s being in prison is the deterrence of crime. Obviously 9 years and 364 days spent in prison is not going to have a significantly different power to deter felonious assault from 10 years spent in prison. So: no good will be secured by visiting on Blodgett that last day in prison, and that last day spent in prison is an evil. The principle tells you, the official, to let him out a day early. This much, I think, is enough to show that the principle is wrong, for you have no such obligation. (p. 101)

And of course nothing in this counterexample turns on the specific numbers van Inwagen chose. If his choices don’t sound convincing, think of a criminal sentenced to 50 years who petitions to be released a day early. Or even a criminal sentenced to 50 years who petitions to be released six hours early. If, in any of those circumstances, you aren’t obligated to release the criminal early, then (P) is false.

Although this is enough to show that (P) is false and thus that premise (3) should be rejected, a closer examination of this counterexample will prove enlightening. What could explain our intuition that we aren’t required to release Blodgett from prison one day early? Here van Inwagen’s idea is that the length of time used for prison terms is largely arbitrary, and indeed must be largely arbitrary given the fact that there are many spans of time

available for the picking that are so similar that choosing one or the other could not possibly make a moral difference. Van Inwagen rightly points out that the lengths of our prison sentences are partly dependent on things like “accidents of astronomy” and “our preference for numbers that can be specified concisely”. (p. 102) We don’t sentence criminals to 9 years and 364 days in prison. Is this because there is some important moral difference secured by adding that extra day? Not at all. We just like our round numbers. In deciding the lengths of prison terms, there are many morally equivalent places at which we could draw the line. But – and here’s the important point – we must draw the line somewhere, and where we draw it will be to some extent morally arbitrary. As van Inwagen puts it, “the principle fails precisely because it forbids the drawing of morally arbitrary lines.” (p. 102)

Another way to put the point is in terms of a sorites paradox. Let n be the number of days Blodgett will spend in prison, and suppose that being in prison for n days has a certain power to deter felonious assault. Surely spending $n - 1$ days in prison will not be significantly different with respect to its power to deter felonious assault. So if (P) is true, then you are required to let Blodgett out of prison after $n - 1$ days in prison. But n could be any number at all. Suppose n is 1,000. Surely 999 days in prison is not significantly different with respect to its power to deter felonious assault. Again, if (P) is true, then you must release Blodgett from prison after 999 days. But given that 999 days in prison has sufficient power to deter felonious assault, surely 998 is not significantly different in this respect. Thus, (P) mandates that you must let Blodgett out of prison after 998 days. And so on. As van Inwagen puts it, “A moment’s reflection shows that...the moral principle entails that Blodgett ought to spend no time in prison at all.” (p. 101) But this is surely absurd, and thus (P) is false.

When construed in this way, the problem with (P) is that it fails to recognize the inherently vague nature of a predicate like ‘has the power to deter felonious assault’. The fact is that there is no precise number of days such that a prison sentence of exactly that length would have the power to deter felonious assault whereas a prison sentence that is one day shorter would not have the power to deter felonious assault. There just is no such sharp

cutoff. And without a sharp cutoff to stop the sorites march, (P) leads us to the conclusion that we should reduce Blodgett's prison sentence to 0 days. Again, this is surely the wrong conclusion to draw, and so (P) must be rejected. Even when making explicitly moral judgments and decisions, we sometimes need the freedom to draw morally arbitrary lines, and it is precisely this freedom that (P) denies us.

Given the plausible supposition that premise (3) of the local argument from evil relies on (P), we can resist the conclusion of the local argument from evil by rejecting (P), and (3) along with it. For any particular horror, H, God is not required to eliminate H because there was a morally arbitrary line to be drawn between the horrors God would allow and the horrors God wouldn't allow, and the place God chose to draw it happened to entail that H would be one of the horrors that God would allow. (Recall that part of van Inwagen's response to the global argument from evil is that God must allow a vast amount of horrors in any case.)

Although we have some residual worries about this defense, let's grant its success for now. That is, let's grant that in some sense, there was a morally arbitrary line to be drawn between horrors that God would allow to exist and horrors that God would preclude from existence, and it had to be drawn somewhere. And let's grant that this fact is reason enough to reject (P) and thus reason enough to reject premise (3) of the local argument from evil. The local argument from evil, therefore, is a failure. Moreover, let's grant that the global argument from evil is a failure, too, given van Inwagen's expanded version of the free-will defense.

It seems to us that even if we grant all of this, the problem of evil still hasn't yet been laid to rest. That is, we don't think that the global and local arguments from evil exhaust the possible atheistic arguments one could mount against the existence of God under the general rubric of the problem of evil. There is a third argument from evil that van Inwagen has yet to consider, and it is van Inwagen's discussion of the local argument of evil that brings this third argument to the fore.

V. The Range Argument from Evil

To see what this third argument looks like, let's consider what the theist has accomplished in establishing the failure of the global and local arguments from evil. This will help us see what the theist has yet to accomplish. A successful response to the global argument from evil provides an adequate explanation of why God would allow a vast amount of evil to exist in the world. And a successful response to any particular local argument from evil provides an adequate explanation of why, for any particular horror H, God would allow H to occur. So far so good. But even after all of this is said and done, something remains to be explained.

What remains to be explained is the particular amount of evil that exists in our world. Granted, it has to be vast. But must it be as vast as it is? Granted, if we pick any particular horror, there's no case to be made that this horror alone seals the case against the existence of God, because a morally arbitrary line had to be drawn. But what if we take a group of horrors considered together? Couldn't a case be made that the world could still have contained a reasonably vast amount of evil even if it didn't include these particular horrors? And might it be plausible to suppose that the inclusion of all of the horrors in this set takes us out of the range of reasonable vastness and into the range of overkill?

Let's begin to regiment the intuitions that these rhetorical questions attempt to capture. First, let us flesh out the notions of reasonable range and overkill. Consider a non-theistic example that includes the justified drawing of arbitrary lines, such as setting the speed limit on a particular stretch of road in a residential area. Typically (in the U.S. anyway), the speed limit on a residential road is set at 25 miles per hour. Given facts such as that children often play along residential roads and that the houses lining the road are quite close to it, 25 mph seems a reasonable choice. And we can easily think of other choices that would not be reasonable – for instance, 80 mph at one extreme and 2 mph at the other. But to choose precisely 25 mph as the speed limit is to some extent arbitrary. Why not 26 or 24? Perhaps there is no reason at all other than the fact that we all like numbers that end in '5'

much better than numbers that end in '6 or '4' (and the perhaps related fact that speed-limit signs are printed this way).

But to say that 25 mph is to some extent an arbitrary choice is not to say that it is a wholly arbitrary choice. After all, there are speeds that are clearly inappropriate choices for this particular stretch of road – 80 mph and 2 mph, for instance. The existence of unreasonable speeds is enough to show that the choice of 25 mph is not wholly arbitrary. But, as we have said, it is partly arbitrary because there are speeds that are clearly as appropriate as 25 mph for this particular stretch of road – 24 mph and 26 mph, for instance. That the people who legislate speed limits decided on 25 mph rather than 26 mph most likely has no interesting or significant explanation. So it looks like there is some reasonable range of speed limits such that each speed in this range is equally appropriate for this particular stretch of road as any other speed in the range. But the reasonable range has limits (though perhaps not precise limits) – 80 mph is clearly outside this range. 80 mph is overkill.

Now suppose that whatever committee decided on 25 mph as the speed limit for this road is asked to defend its choice against two challenges. The first challenge asks the committee to explain why the speed limit needs to be set at any high number at all instead of, say, 2 mph? The second challenge asks the committee to explain why the speed limit needs to be set at the particular high number of 25 mph instead of, say, 24 mph? It should be obvious what the committee will say in each case.

Against the first challenge, the committee will point out that part of the idea behind any sort of infrastructure is that it is meant to help people get to where they are going. If drivers went 2 mph on every road, the infrastructure would not be functioning properly. Speed limits need to be set at some relatively high number so that traffic can actually flow rather than crawl. Car traffic is meant to be faster than pedestrian traffic, and so speed limits must be set comparatively high.

Against the second challenge, the committee will point out that their choice was to some extent arbitrary. Why not 24? No particular reason other than, perhaps, the fact that most car speedometers are graduated in increments of 5 mph (and readily available speed-limit signs are printed this way). Why not 26? Same answer. But of course some choice had to be made, and so the committee selected 25 mph, despite its partly arbitrary nature. This is perfectly within their province.¹⁵

So the committee has easily defended its choice of 25 mph against both challenges. But isn't it clear in this case that there is a remaining challenge to be raised? Suppose a coalition of concerned residents were to get together and present the following argument to the committee:

We understand that the speed limit must be set at a relatively high number given the point of infrastructure. And we understand why you are justified in choosing 25 mph over 24 mph or 26 mph, given the partly arbitrary nature of the choice you had to make. But surely what makes your choice of 25 mph only partly (as opposed to wholly) arbitrary is the fact that you must have first selected (in some non-arbitrary way) a particular reasonable range of speeds from which to make your arbitrary choice. Given a particular reasonable range, say, from 20 mph to 30 mph, any choice you make is indeed wholly arbitrary. But this is on the assumption that you've already figured out what the reasonable range is. And what we are wondering now is: Why did you choose the particular range you did? And what makes that range reasonable?

Now, we suspect that the committee will have a response to this challenge, as well, and it will most likely have to do with some complicated statistics from which it can be shown that setting residential speed limits at, say, 60 or 50 or 40, is a pretty bad idea. But it is important to see that this challenge is indeed distinct from the first two challenges we discussed. An answer to one of the first challenges will not ipso facto be an answer to this current challenge.

Let's leave our story about speed limits, and head back into the theistic realm. We hope that the parallel is fairly easy to see. The global and local arguments from evil, as van Inwagen has presented them, can be captured by the following two questions (perhaps directed at God):

Global: Why is there a vast amount of evil in the world?
Local: Why does the world include this particular horror?

Van Inwagen has ably provided answers to these questions on God's behalf. But it should be clear that there is a third question he hasn't provided an answer to, namely: Why does the amount of evil in the world fall in the particular range that it in fact does? Call the argument that this question captures The Range Argument from Evil. The argument can be stated as follows:

- (1) If there were a God, the amount of evil in the world would fall within a reasonable range.
- (2) But the amount of evil in the world surely falls within the overkill range.
- (3) Therefore, there is no God.

Let's look closely at the premises here, starting with premise (1). The lesson that van Inwagen taught us with his response to the global argument from evil is that perhaps the world has to include a vast amount of suffering in order for God's plan of atonement to succeed. But of course there is vastness and then there is vastness. Given that God's goal in allowing evil is to ensure the success of his plan of atonement (as van Inwagen's defense has it), there will surely be a measure of vastness that is far beyond what is absolutely necessary for the plan to succeed. This is not to say that there will be some precise number of evils that separates the 'vast enough' and the 'too vast' categories, though. It is merely to say that there is such a thing as a 'too vast' category in the first place.¹⁶ And it is plausible to think that God would not include in the world any amount of evil that falls into that category. Not only would such an amount be unnecessary for the success of God's plan of atonement, it would be unreasonably unnecessary, as it were, or unnecessary in the extreme.

The need to draw a morally arbitrary line through those evils to allow into existence and those to preclude guarantees that there will be some amounts of evil that are strictly speaking unnecessary for God's plan of atonement. But some of those unnecessary amounts would have been just as reasonable for God to have chosen to include, and some would have been unreasonable for God to have chosen to include. Analogously, the need to draw an arbitrary line through the numbers at which the committee could possibly set the speed limit guarantees that there will be some speed limits that are strictly speaking unnecessary

for the success of the committee's plan to manage infrastructure efficiently. But some of those unnecessary amounts would have been just as reasonable for the committee to have chosen (such as 26 mph), and some would have been unreasonable (such as 80 mph). Premise (1) of the Range Argument merely states that an all-powerful, all-loving creator would have drawn the line at an amount of evil that, while perhaps strictly speaking unnecessary, was at least *reasonably unnecessary*.

Premise (2) goes on to assert that the actual amount of evil in the world is *so* vast that it must fall into the range of the *unreasonably unnecessary* – the overkill range. How might one defend this premise? At first it may seem that this is exactly where the Range Argument from Evil is much weaker than its siblings. After all, in both the global and local arguments from evil, the premise about the actual evil in the world is indisputable. The global argument merely asserts that there is a vast amount of evil, and the local argument merely points to an actual instance of evil. Neither claim is controversial. The Range Argument, on the other hand, makes a more substantive claim about the nature of the evil in the world, namely that it falls into the 'too vast' category. And it's hard to see how an atheist could plausibly argue that the amount of evil in the world is far more than is reasonably necessary for God's plan of atonement. Nevertheless, we think the atheist has a plausible way to make his case.

A suggestion we made above was that the Range Argument from Evil could be made vivid by considering a *set* of horrors, rather than by considering any single instance of a horror. If the set were sufficiently large, it seems that a case could be made for the claim that God's including all the members of that set of horrors brings the actual amount of evil out of the reasonable range and into the overkill range. Granted, the elimination of a single horror from that set wouldn't make a morally relevant difference, but perhaps the elimination of all of the horrors in that set would. So let's consider a particularly large set of evils: those that constitute the Holocaust. The Holocaust is not one horror, at least not on

any plausible way of individuating horrors. Rather, the Holocaust is a collection of millions, probably billions, of horrors. Couldn't someone plausibly address God as follows:

I know that in order for your plan of atonement to succeed, the world must contain a vast amount of evil. I also know that for any particular horror I choose, your choice to include that horror may just have been the result of your drawing a morally arbitrary line. But what about the Holocaust? Did the world really need to include the Holocaust, with its billions of horrors? Surely a world without the Holocaust would still have contained a vast amount of evil – plenty to ensure the success of your plan of atonement. So doesn't it seem that including the Holocaust is just overkill?

We submit that the theist needs a defense against this charge and, significantly, neither van Inwagen's defense against the global argument nor the defense against the local argument will work in this case.

To bring out the problem in a slightly different way, consider what van Inwagen says about the context of distributive justice and taxation:

... we might easily find ourselves in a moral situation... in which we must draw an arbitrary line and allow some bad thing to happen when we could have prevented it, and in which, moreover, no good whatever comes of our allowing it to happen. In fact, we do find ourselves in this situation. In a welfare state, for example, we use taxation to divert money from its primary economic role in order to spend it to prevent or alleviate various social evils. And how much money, what proportion of the gross national product, shall we – that is, the state – divert for this purpose? Well, not none of it and not all of it (enforcing a tax rate of 100 percent on all earned income and all profits would be the same as not having a money economy at all). And where we draw the line is an arbitrary matter. However much we spend on social services, we shall always be able to find some person or family who would be saved from misery if the state spent (in the right way) a mere £1000 more than it in fact plans to spend. And the state can always find another £1000, and can find it without damaging the economy or doing any other sort of harm. (p. 108)

It is indisputable that where we draw the line is an arbitrary matter in the sense relevant to van Inwagen's response to the Local Argument from Evil: as he points out, wherever the state draws the line, we can always save another family from misery by taxing people just a bit more and providing more money. But it should also be absolutely clear that it is not an arbitrary matter whether the income tax rate is set at 5 percent or 15 percent; that is, even granting van Inwagen's point, we could still think that it makes a non-arbitrary – indeed a significant – moral difference whether the tax rate is in one range (around 5 percent) or

another (around 15 percent). To suppose that the latter issue is an arbitrary matter would be to consign much of the philosophical debate about distributive justice to the flames, and it is just manifestly false that it is morally arbitrary whether the state engages in minimal or more substantial redistributive taxation. (This is of course not to prejudge the substantive issue of what is just in this context; it is simply to point out that it is a substantive moral issue.)

Conclusion

In this book van Inwagen offers a sustained and ingenious argument for the view that the problem of evil is a philosophical failure. There is much to be learned from the resourceful argumentation. One might think that the interest of the conclusion that the problem of evil is a philosophical failure is at least somewhat vitiated by van Inwagen's view that all philosophical arguments for interesting and substantive conclusions are failures; perhaps van Inwagen would say that it is important to see why the members of this particular family of arguments are failures, even if all philosophical arguments for the relevant sorts of conclusions are indeed failures. On this sort of approach, one could view philosophy in general as showing why exactly each philosophical argument for the various interesting philosophical theses fails; although this is perhaps a somewhat less than perky vision, it would still leave an interesting role for philosophical argumentation.

We have suggested some alternative models of philosophical argumentation, by reference to which not all philosophical arguments would be deemed failures. Also, we have suggested that van Inwagen has not put the problem of evil to rest. Even if he has pointed toward the inadequacies of the global and local arguments, the range argument remains. So even if it is an arbitrary matter whether to set the speed limit at 25 mph or 24 mph, it is obviously not an arbitrary matter whether to set it at 25 mph or 45 mph. And even if it is morally arbitrary whether to set the income tax rate at 5 percent or 4 percent, it is surely not morally arbitrary whether to set it at 5 percent or 15 percent. Similarly, one could surely wonder why God created a world with so much evil? Isn't this overkill? No doubt there are

various possible answers, but van Inwagen's defenses against the global and local versions of the problem of evil do not in themselves provide an answer.¹⁷

¹ Clarendon Press, Oxford: 2006; xiv + 183 pp. (We refer in parentheses to page numbers of this book throughout the text of this paper.)

² See, for example, John Martin Fischer, "Frankfurt-Type Compatibilism," in S. Buss and L. Overton (eds), *Contours of Agency: Essays on Themes from Harry Frankfurt* (MIT Press, 2002): 1-26; reprinted in G. Watson (ed.), *Free Will: Oxford Readings in Philosophy* Second Edition, (Oxford University Press, 2003): 190-211, and John Martin Fischer, *My Way: Essays on Moral Responsibility* (New York and Oxford: Oxford University Press, 2006): 124-142.

³ Ishtiyaque Haji and Michael McKenna, "Dialectical Delicacies in the Debate about Freedom and Alternative Possibilities," *Journal of Philosophy* 101 (2004): 299-314; and "Defending Frankfurt's Argument in Deterministic Contexts: A Reply to Palmer," *Journal of Philosophy* 103 (2006): 363-72; and Alfred Mele, *Free Will and Luck* (New York and Oxford: Oxford University Press, 2006).

⁴ In footnote 5, Chapter (Lecture) 3, on page 160, van Inwagen says:

[...my model] of philosophical success and failure has the same consequence as the two criteria I have rejected: most if not all arguments for substantive philosophical conclusions are failures. And this was the consequence that was my reason for rejecting those criteria. Should I not, for just this reason, reject the proposed criterion and look for some more liberal criterion? Alas, there is no more liberal criterion. The criterion I have proposed is the most liberal possible criterion. ... My position, then, is that, sadly, every known argument for a substantive philosophical position is a failure – by the most liberal (by the most possibility-of-success-friendly) possible criterion of success and failure.

We are not sure why van Inwagen believes that his suggested model is the most liberal possible. In the text above, we suggested a liberalization according to which proponents of competing positions may change the level of certainty or credence they place in their beliefs (rather than requiring them to give up their antecedent views and adopt the alternative

view). Additionally, we develop another sort of liberalization of van Inwagen's model in the text below.

⁵ Here we have benefited from very helpful discussion on the blog, "The Garden of Forking Paths": <http://gfp.typepad.com>; see the post, "So A Mysterianist and a Semi-Compatibilist Walk Into a Bar...", January 18, 2007.

⁶ Here is an amusing (and also very realistic) view about success in philosophical argumentation:

Recently I had lunch with a graduate student at the APA. During the lunch I presented an argument for compatibilism. The student works in ethics and epistemology but does no work on free will. By the end of the lunch she claimed that she no longer thought that compatibilism was a crazy view. I thought that my argument was a huge success! I generally think that I've done a good job if I can convince a few students in my undergraduate metaphysics class that compatibilism is not a crazy view. (Joseph Campbell, "So A Mysterianist and a Semi-Compatibilist Walk Into a Bar," <http://gfp.typepad.com>, discussion thread)

⁷ Van Inwagen is never explicit about why he finds the traditional distinction useless, but what is interesting is that part of his response to the global argument from evil actually doubles as a response to the traditional logical argument from evil. Indeed, the initial response he has his theist make to the global argument from evil includes this statement: "The conclusion that evil does not exist does not, therefore, follow logically from the premises that the non-existence of evil is what God wants and that he is able to bring about the object of his desire – since, for all logic can tell us, God might have reasons for allowing evil to exist that, in his mind, outweigh the desirability of the non-existence of evil." (p. 65) Perhaps since this is included in the theist's very first response to the global argument from evil, van Inwagen thinks that the logical argument from evil so obviously fails that it isn't worth discussing on its own.

⁸ Strictly speaking, what follows is just the general structure of any particular local argument from evil, abstracted from van Inwagen's presentation in the text.

⁹ A horror, as van Inwagen uses the term, is just a "certain particular very bad event". (p. 95)

¹⁰ This presentation of the local argument follows van Inwagen's in the text, though we have added the "suitably chosen additional premises" needed to make the argument formally valid. See p. 98.

¹¹ On the free-will defense, see Alvin Plantinga, *God, Freedom and Evil* (Grand Rapids, MI: Wm. B. Eerdmans, 1977).

¹² Van Inwagen tells us that the word 'evil' in this context just means 'bad things'. (p. 4)

¹³ William L. Rowe, "The Problem of Evil and Some Varieties of Atheism," *American Philosophical Quarterly* 16 (1979): 335-41. See also the similar example of the baby bat due to John Perry in his *Dialogue on Good, Evil, and the Existence of God* (Indianapolis, Ind.: Hackett Publishing Co., 1999: 49).

¹⁴ We have lifted this quotation from the lecture in which van Inwagen discusses the "sufferings of beasts", although this particular quotation clearly expresses his strategy of responding to the local argument from the sufferings of humans, as well.

¹⁵ Perhaps this also explains – what is now a matter of common knowledge, whether true or not – the alleged fact that a police officer won't pull you over for speeding unless you are traveling at least 3 or 4 mph over the limit.

¹⁶ A word of warning. We need to be careful when talking about ranges, lest we fall away from "the bright world of good sense" (p. 108) into thinking that there are sharp cutoffs governing vague predicates, as epistemicists about vagueness think. Our discussion of reasonable and overkill ranges may make it seem that we are illegitimately supposing that there will be a sharp cutoff between, say, the amount of evil at the upper limit of the reasonable range and the amount of evil at the lower limit of the overkill range. But our argument requires no such supposition. Let the upper and lower limits of the ranges be fuzzy, if you like. All the argument requires is that there are clear cases of amounts that fall within the boundaries of each range.

¹⁷ For helpful discussion of these issues, we are grateful to Mary Amschel, Chris Franklin, James Gibson, William Hasker, Robert Sanchez, Patrick Todd, and Howie Wettstein.