The Argument from Particular Horrendous Evils

PETER VAN INWAGEN

Suppose that by an "argument from evil" we understand an argument that attempts to show that there is no God by some sort of appeal to the evils we find in the world. I want to suggest that there is not one argument from evil but two. In saying this, I am not appealing to any of the usual distinctions—I am not, for example, appealing to the familiar distinction between the logical argument from evil and the evidential argument from evil. I have in mind an entirely different distinction, a distinction that I believe philosophical theologians have been insufficiently attentive to. I will call the two arguments I wish to distinguish "the argument from the fact of evil" and "the argument from particular evils." Very roughly speaking, the former appeals to the fact that there is evil in the world—or, if you like, to the fact that there is a vast amount of truly horrendous evil in the world—and it goes roughly as follows: if there were a God we should not find ourselves in a world like this, a world that contains vast amounts of horrendous evil. The latter asks us to consider some *particular* horrendous evil, and consists in an attempt to establish that, if there were a God, that particular horrendous evil would not exist. (I call this argument the argument from particular horrendous evils—plural—because, although any given presentation of the argument will mention only one evil, the choice of that evil will be arbitrary; there will be vastly many other evils that would have served the dialectical purposes of the author of the argument equally well.) I want to argue for several theses. First, that a refutation of the argument from the fact of evil, however successful it might be, would not be a refutation of the argument from particular evils. Secondly, that if there were a refutation of the argument from the fact of evil, it would *provide the materials for* a refutation of the argument from particular evils. It can hardly be evident at this point, however, what I am talking about. I counsel perseverance.

I will begin by asking you to imagine that some philosopher has refuted the argument from the fact of evil. But what am I asking you to imagine?

What would it be to refute the argument from the fact of evil? It would not, of course, be to show that it is false that the world contains vast amounts of horrendous evil, for that is not false and so can't be shown to be false. It would certainly be a step in the direction of a refutation to show that it does not *follow*—does not follow logically—from the proposition that there is a vast amount of horrendous evil, that there is no God. But to show that would not be to refute the argument from the fact of evil, since it might be that although the non-existence of God does not follow logically from the proposition that there is a vast amount of horrendous evil, the non-existence of God *does* follow logically from that proposition together with certain other propositions, all of which seem evidently true. And that, indeed, is just what any of our contemporaries who offers a version of the argument from the fact of evil will claim for that version of the argument: It will have as one premise the proposition that there is a vast amount of horrendous evil, and it will have various other premises that—at least according to the author of the argument—are evidently true; and the argument will be logically valid. In my view, if one would refute the argument from the fact of evil, one must do this: one must tell a story that contains both God and vast amounts of horrendous evil and one must try to show that this story has the following feature: it is true for all anyone knows; or, at any rate, it is true, for all anyone knows, given that there is a God. In other words, if one would refute the argument from the fact of evil, one must tell a story that contains both God and vast amounts of horrendous evil, and one must attempt to produce a reaction in one's audience that could be put into words like these: "Given that there is a God, the rest of that story might well be true, too." I will call any story that has this content and that has the power to produce this reaction in an audience a *defense*.

I want to say as little as possible about the content a defense might have, but I assume that any defense will have the following general form. It will assert that evil did not enter the world by God's will or decree; evil, a defense will tell us, had some other source than God's will (perhaps the creaturely abuse of free will, which is, of course, the "source" favored by most writers who attempt to answer the argument from the fact of evil). A defense will ascribe to God some *reason* for allowing the possibility of evil in his creation (for example: creaturely free will is a very great good, a good so great that its existence justifies the risk of its possible abuse). It will go on to say that this source, whatever it may have been, produced not just *some* evil but vast amounts of horrendous evil, and it will, finally, ascribe to God another reason, a reason for not simply removing from his creation by fiat the vast amounts of evil that issued from the Source of Evil, a reason for allowing the vast amounts of horrendous evil produced by the Source to

continue to exist. And I assume that the reason the story ascribes to God will be of this general sort: God has in mind some great good, a good that significantly outweighs—perhaps because it is infinite—even a vast amount of horrendous evil; this great good does not yet exist, but God has a plan that, if acted upon, will one day bring it into existence; the temporary existence of a vast amount of horrendous evil is an essential part of this plan; even an omnipotent being cannot bring about this good otherwise than by acting on a plan that has the existence of a vast amount of horrendous evil as a part.

When I say that to produce a defense is what one must do to refute the argument from the fact of evil, I mean, of course, that this is the *minimum* that one must do. One might attempt more; one might, for example, tell a story that contains both God and vast amounts of horrendous evil, and attempt to convince one's audience that if God exists the rest of story *is* true. That, from the point of view of the theist, would be even better than convincing one's audience that the story *might well be true*. But, desirable as it would be to explain to the world why God actually does allow evil, one need not do that very desirable thing to refute the argument from evil. For, or so I would suppose, to refute a logically valid argument, it is not necessary to show that one or more of its premises is false; it suffices to show that one or more of its premises is open to reasonable doubt. And that is just what a defense, a story having the content and degree of plausibility I have specified, would accomplish: anyone who was willing to grant that a story had these features would also have to grant that at least one of the premises of the argument from the fact of evil was open to reasonable doubt.

Let us imagine that someone has done this. Against all the odds, some great philosopher—I'll call him Theodore—has contrived a defense and has thereby got even the atheist to admit that one or more of the premises of the argument from the fact of evil is, if not demonstrably false, at least open to reasonable doubt. (It would, of course, be a premise other than 'The world contains vast amounts of horrendous evil'; that premise is not open to reasonable doubt, and my proposal about what it would "take" to refute the argument from the fact of evil presupposes that that premise of the argument is true.) Theodore has done this, let us suppose, by composing a story called "the second-order free-will defense." (I will say almost nothing about the content of "the second-order free-will defense." The phrase is meant only as a place-holder for the name of a particular defense, and is, in fact, quite meaningless. Despite the name I have given it, we need not suppose that Theodore's defense has anything to do with free will. We will suppose only that it ascribes to God some reason for allowing the possibility of evil; that it asserts that this possibility became an actuality owing to factors other than God's will; that the possibility was realized in the form of vast amounts of

horrendous evil; and, finally, that God has a reason for allowing the vast amounts of horrendous evil that came to exist in this way to continue in existence, and that this reason is of the general sort I imagined a moment ago—that is, that this evil is a means necessary to a future good that outweighs it. I will feign that Theodore, in telling his story, has produced a plausible candidate for this great future good, that he has given us a nontrivial description of it, and that he calls it the Summum Bonum Creatum.) If Theodore has done all that I have imagined, he has, I believe, refuted the argument from the fact of evil. But this is not the end of the matter. The argument from the fact of evil is, as I have said, not the only argument from evil. There is also the argument from particular evils, to which I now turn. I will put the argument from particular evils into the mouth of an atheist called Athena, who, I imagine, has just heard Theodore's refutation of the argument from the fact of evil and has been suitably impressed. I will assume that she addresses the argument from particular evils to Theodore, and I will go on to consider how Theodore might reply to it. Athena argues as follows.

> Well, Theodore, your "second-order free-will defense" is very clever—it does seem to have the consequence that there is a very plausible reason a being like God might have for allowing the coming-to-be and continued existence of vast amounts of horrendous evil. He might, as you have said, allow the continued existence of vast amounts of horrendous evil because the continued existence of vast amounts of horrendous evil is a necessary condition for the achievement of the very great good you call the Summum Bonum Creatum. And I must admit that we atheists had never thought of that possibility. Nevertheless, there is more to be said about God and evil. I will try to say it. Let me begin by noting that everything you have said is very abstract. That is, one component of your defense is the existence of vast amounts of horrendous evil, but the defense says nothing about any *particular* horrendous evil. (Let us, for the sake of concision, call a particular horrendous evil a "horror.") There are, as you and I agree, many horrors in the world, vastly many. You have shown that the fact that the world is well supplied with horrors does not by itself show that there is no God. But I want to call your attention, and the attention of everyone who has heard and has been as impressed as I by your clever defense, to a certain feature that belongs to at least a significant proportion of the horrors of the world. It seems evident to me that many, many horrors have the following property: no discernible good results from them, I mean results from any of them individually—and certainly no good, discernible or not, that an omnipotent being couldn't have got without the horror; in fact, without any suffering at all. But let us descend from the abstract to the concrete. Let us look at a particular case. Here is a story that is, unfortunately, no philosopher's example. This story is true. A man came upon a young woman in an isolated place. He overpowered her, chopped off her arms at the elbows with an axe, raped her, and left her to die. Somehow she managed to drag herself on the stumps of her arms to the side of a road,

where she was discovered. She lived, but she experienced indescribable suffering, and although she is alive, she must live the rest of her life without arms and with the memory of what had been done to her. No discernible good came of this, and it is wholly unreasonable to believe that any good could have come of it that an omnipotent being couldn't have achieved without employing the raped and mutilated woman's horrible suffering as a means to it. And even if this is wrong and some good came into being with which the woman's suffering was so intimately connected that even an omnipotent being couldn't have got the good without the suffering, it wouldn't follow that that good outweighed the suffering. (It would certainly have to be a very great good to do that.) I will now construct a version of the argument from evil, a version that makes reference only to this one event, this one horror, and not to the general or global fact that the world contains vast amounts of horrendous evil. (The argument is modeled on the central argument of William Rowe's "The Problem of Evil and Some Varieties of Atheism."[1]) I will refer to the events in the story I have told collectively as 'the Mutilation'. I argue:

(1) If the Mutilation had not occurred, if it had been, so to speak, simply *left out* of the world, the world would be no worse than it is. (It would seem, in fact, that the world would be significantly *better* if the Mutilation had been left out of it, but my argument doesn't require that premise.)

(2) The Mutilation occurred and was a horror.

(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—or at any rate he would have left it out if he had been able to.

(4) If an omnipotent being created the world, he was able to leave the Mutilation out of the world (and was able to do so in a way that would have left the world otherwise much as it is).

There is, therefore, no omnipotent and morally perfect creator.

I do not claim that this argument is formally valid (Athena continues), but it could obviously easily be made formally valid by the addition of suitably chosen additional premises. Since is obvious that all the additional premises that would be needed to make the argument formally valid would be true, it would be pedantic actually to search them out. You, Theodore, must deny at least one of the four premises I have explicitly stated; or at any rate you must show that serious doubts can be raised about at least one of them. But which? Your story of the Summum Bonum Creatum does not tell us which of these premises to deny. It gives us no reason to doubt any of them. God, you say, plans to bring the Summum Bonum Creatum to realization, and his plan can succeed only if the world contains vast amounts of horrendous evil. Well and

> good. But note this: if God had left the Mutilation out of the world, if, to speak in terms of time, he had prevented the Mutilation by some small miracle shortly before the moment at which the Mutilation would have occurred (and had left the world otherwise much as it was), the world would *still* have contained vast amounts of horrendous evil; leaving the Mutilation out of the world would therefore not have frustrated the plan for the realization of the Summum Bonum Creatum that you describe in your second-order free-will defense.

So speaks Athena. How might Theodore reply? Athena has said that her argument was modeled on an argument of William Rowe's. If Theodore models his reply on the replies made by most of the theists who have responded to Rowe's argument, he will attack the first premise. He will argue that, for all anyone knows, the world (considered under the aspect of eternity) is a *better place* for containing the Mutilation. He will argue that, for all anyone knows, God has brought, or will at some future time bring, some great good out of the Mutilation, a good that outweighs it, or else has employed the Mutilation as a means to the prevention some even greater evil; and he will argue that, for all anyone knows, the great good achieved or the great evil prevented could not have been, respectively, achieved or prevented, even by an omnipotent being, otherwise than by some means that essentially involved the Mutilation—or something else at least as bad. (And this great good, great as it is, will not be anything on the scale of the Summum Bonum Creatum, a *global* good that requires the existence of vast amounts of horrendous evil, but requires the existence of no particular evil; and, as Athena has pointed out, there would still exist vast amounts of horrendous evil if the Mutilation were removed from the world. Athena's argument has this form: that the Mutilation can be justified only if it is necessary for some great "local" good, a good for which it, that particular horror, is a necessary condition; and it is wholly unreasonable to believe that there is any such good.)

I am not going to have Theodore reply to Athena's argument in this way. I find (1) fairly plausible, even if I am not as sure as Athena is (or as sure as most atheists who have discussed the issue seem to be) that (1) is true. And I find it even more plausible, very plausible indeed, to suppose that the following existential generalization of (1) is true:

> There has been, in the history of the world, at least one horror such that, if it had not existed, if it had been, so to speak, simply *left out* of the world, the world would be no worse than it is.

If this generalization is true, then, even if (1) is false, there has been at least one horror in the history of the world that Athena could use to show that the world was not created by an omnipotent and morally perfect being—given, of course, that the other three premises of her argument, suitably adjusted, are true.

I am going to represent Theodore as employing another line of attack on Athena's argument. I am going to represent him as denying premise (3), or, more precisely, as trying to show that, if there is a story having the features we are supposing the imaginary second-order free-will defense to have, that story casts considerable doubt on premise (3). In order to enable you the better to follow what Theodore says, I will attempt to fix the essential content of premise (3) in your minds by restating it in terms of a rather fanciful metaphor. Consider a morally perfect creator who is taking a final look at the four-dimensional blueprint of the world he is about to create. His eye falls on a patch in the blueprint that represents a horror. He reflects a moment and sees that if he simply erases that patch, replaces it with something innocuous, and does a little smoothing around the spatiotemporal edges to render the lines of causation in the revised blueprint continuous (or nearly so), a world made according to the revised blueprint will contain at least as favorable a balance of good and evil as a world made according to the unrevised blueprint. He perceives a moral obligation to revise the blueprint in the way he has thought of and to incorporate the revision into his creation, and, being morally perfect, necessarily revises and creates accordingly. Premise (3) simply says that this is what *must* happen when a morally perfect creator perceives in his plan for the world a horror that can be "edited out" without significantly altering the balance of good and evil represented in that plan.

Now that we have, I hope, got the content of premise (3) into our minds in an intuitive and memorable form, we are ready to hear Theodore's reply to Athena's argument.

> Why should we accept premise (3) of your argument? I have had a look at Rowe's defense of the corresponding premise of his argument, the entirety of which I will quote:
>
> > [This premise] seems to express a belief that accords with our basic moral principles, principles shared both by theists and non-theists.[2]

This is not much of a defense. Still, let us consider it. We must ask, what are these "basic moral principles, shared both by theists and non-theists"? Rowe does not say, but I believe there is really just one moral principle that it would be plausible to appeal to in defense of premise (3). It might be stated like this.

> 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.

It should be noted that this principle does *not* say that if allowing the evil to occur would result in some good that would outweigh the evil or preventing it would result in some other evil at least as bad, then one *should* allow the evil to occur—or even that it was morally permissible to allow the evil to occur. Here's a simpler moral principle that illustrates this point: One should not lend one's car to someone unless that person is sober. It doesn't follow from this principle that one should lend one's car to just anyone who is sober or even that it would in all cases be morally permissible to; one should not, for example, lend one's car to a sober acquaintance who has been twice convicted of vehicular homicide.

A word about the phrase 'in a position to'. I mean these words to imply both "is able to" and "is morally permitted to." As to the latter implication, perhaps—no matter what the utilitarians may say—it is sometimes simply *not one's place* to prevent certain evils. Some threatened evils may be such that to prevent them would constitute officious meddling in the lives of one's fellow citizens, a disregard of everyone's right to go to hell in his own custom-made hand-basket. Or to prevent certain evils might be to presume a legal or moral authority one does not have (consider the case of a police officer who secretly murders a serial killer whom the law cannot touch). Insisting on the moral component of 'in a position to' is probably needed to make the principle that (I suggest) Rowe is appealing to plausible. But, having said this, I can proceed to ignore it, since, it would seem, it is never morally impermissible for *God* to prevent an evil; not at any rate on the ground that to do so would be to interfere in a matter that is really none of his business, or on the ground that it falls outside the scope of his moral authority. God is not our fellow citizen but our Maker, and all moral authority is his.

Now: Is the moral principle true? I think not. Consider this case. 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. (Or consider the probable reaction of a prisoner who is, by bureaucratic foot-dragging, kept in prison one day longer than the term of his sentence.) Let's suppose that the only good that could result from someone's being in prison is the deterrence of crime. (This assumption is made to simplify the argument. That it is false introduces no real defect into the argument.) Obviously, nine years, three hundred and sixty-four days spent in prison is not going to have a significantly different power to deter felonious assault from ten 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 therefore 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. But the principle is in more trouble than this simple criticism suggests.

It would seem that if a threatened punishment of n days in prison has a certain power to deter felonious assault, $n - 1$ days spent in prison will have a power to deter felonious assault that is not significantly less. Consider the power to deter felonious assault that belongs to a threatened punishment of 1023 days in prison. Consider the power to deter felonious assault that belongs to a threatened punishment of 1022 days in prison. There is, surely, no significant difference? Consider the power to deter felonious assault that belongs to a threatened punishment of 98 days in prison. Consider the power to deter felonious assault that belongs to a threatened punishment of 97 days in prison. There is, surely, no significant difference? Consider the power to deter felonious assault that belongs to a threatened punishment of one day in prison. Consider the power to deter felonious assault that belongs to a threatened punishment of no time in prison at all. There is, surely, no significant difference? (In this last case, of course, this is because the threat of one day in prison would have essentially *no* power to deter felonious assault.)

A moment's reflection shows that if this is true, as it seems to be, then the moral principle entails that Blodgett ought to spend no time in prison at all. For suppose that Blodgett had lodged his appeal to have his sentence reduced by a day not shortly before he was to be released but before he had entered prison at all. He lodges this appeal with you, the official who accepts the moral principle. For the reason I have set out, you must grant his appeal. Now suppose that when it has been granted, clever Blodgett lodges a second appeal: that his sentence be reduced from ten years minus one day to ten years minus two days. This second appeal you will also be obliged to grant, for there is no difference between ten years less a day and ten years less two days as regards power to deter felonious assault. I am sure you can see where this is going. Provided only that before Blodgett enters prison, he has the time and the energy to lodge 3648 successive appeals for a one-day reduction of his sentence, he will escape prison altogether.

This result is, I take it, a *reductio ad absurdum* of the moral principle. As the practical wisdom has it (and this is no compromise between practical considerations and strict morality; it *is* strict morality), "You have to draw a line somewhere." And this means an *arbitrary* line. The principle fails precisely *because* it forbids the drawing of morally arbitrary lines, which is a thing we are sometimes morally required to do. There is nothing wrong, or nothing that can be determined *a priori* to be wrong, with a legislature's

setting ten years in prison as the minimum punishment for felonious assault—and this despite the fact that ten years in prison, considered as *a precise span of days*, is an arbitrary punishment. Let us suppose that a certain legislature has indeed passed a law to this effect. It seems obvious that if the terrestrial day were about twenty-five seconds shorter, that law (which would no doubt be stated in the same words) would be neither more nor less effective than it is in fact; the interval of time denoted by the phrase 'ten years' would, however, be about one day shorter than the interval of time actually denoted by that phrase. It is obvious, when you think about it, that the lengths of the prison sentences written into our laws depend on accidents of astronomy, on the (to some degree) accidental fact that we use a decimal rather than a binary or duodecimal system of numerical representation, and on many other arbitrary factors—such as the prejudice, widespread among those who draft legislation, in favor of numbers that can be concisely specified. And there is nothing wrong with this. Since, however we "draw the line," its exact position will be an arbitrary matter, we might as well let its exact placement depend partly on the set of (morally speaking) arbitrary preferences that nature has dropped into our collective lap.

So, the moral principle is false—or possesses whatever defect is the analogue in the realm of moral principles of falsity in the realm of factual statements. What are the consequences of its falsity, of its failure to be an acceptable moral principle, for Athena's argument for the conclusion that God would not have allowed the Mutilation to occur? Let us consider Theodore's story, the imaginary story I have fancifully dubbed the second-order free-will defense. This story accounts for the existence of horrors—that is, that there are horrors (and in fact a vast number of them) is a part of the story. The story explains why there are such things as horrors, although it says nothing about any particular horror. But to explain why there are horrors is not to meet the argument from particular horrendous evils—or, as we may now call it, the argument from horrors.

A general account of the existence of horrors does not constitute a reply to the argument from horrors because, as we have seen, it does not tell us which premise of the argument to deny. It might well be that, although theists have an excellent reply when their critics say, "There's just no conceivable reason why God would allow the existence of vast quantities of horrendous evil," they have no reply when their critics say, "There's just no conceivable reason why God would allow the existence of the Mutilation." But if Theodore cannot simply present his second-order free-will defense and claim thereby to have refuted Athena's argument, his story does, as I have said, provide the materials for a refutation of her argument. We are now in a position to imagine how Theodore might reply to Athena's invita-

tion to say which of the premises of her argument he is going to declare "not proven." Here is what Theodore might say.

> God (according to the defense I have presented) did not decide that the world was to contain vast amounts of horrendous evil—or any evil at all. But he did have a reason for allowing the *possibility* of evil, and this possibility was, unfortunately, realized, and in a very horrible way: in the existence of vast amounts of horrendous evil. God could, being omnipotent, purge the world of this evil, but he has a reason for not doing so; he plans eventually to bring a very great good, a good that significantly outweighs the horrors of the world, the Summum Bonum Creatum, out of this these horrors, and his plan requires the existence of vast amounts of horrendous evil. But it is *consistent* with the second-order free-will defense that God does remove much of the evil that would otherwise exist from the world. We have imagined that a source of possible evil is inherent in God's creation; we have imagined that this source does indeed produce evil. But suppose that God does act to prevent much of the evil that would have issued from the source of evil. That is, once evil has entered the world, God sees that if he interacts with the world simply by sustaining the existence and normal causal propensities of its inhabitants, a horror will occur, and he does more than let matters take their course; he makes specific local changes in the world in such a way that what would have happened doesn't and the threatened horror is prevented. And suppose he does this not once but an enormous number of times: he acts to prevent an enormous number of horrors that otherwise would have occurred. He may prevent an enormous number of horrors, but he cannot prevent *all* of them, for that would frustrate his plan for achieving the Summum Bonum Creatum. And if he prevents only some horrors, how shall he decide which ones to prevent? Where shall he draw the line?—the line between threatened horrors that are prevented and threatened horrors that are allowed to occur? I suggest that wherever he draws the line, it will be an arbitrary line. That this must be so is easily seen by thinking about the Mutilation. If God had added that particular horror to his list of horrors to be prevented, and that one alone, the world, considered as a whole, would not have been a significantly less horrible place, and conditions necessary for the eventual achievement of the Summum Bonum Creatum would not have been significantly different from what they are. Therefore, as Athena has pointed out, preventing the Mutilation would in no way have interfered with his plan for the realization of the Summum Bonum Creatum. If the second-order free-will defense is a true story, God has made a choice about where to draw the line, the line between the actual horrors of history, the horrors that are *real*, and the horrors that are mere averted possibilities, might-have-beens. And the Mutilation falls on the "actual horrors of history" side of the line. And this fact shows that the line is an arbitrary one, for if he had drawn it so as to exclude the Mutilation from reality (and left reality otherwise the same) he would have lost no good thereby and he would have allowed no greater evil. Therefore, he had no reason for drawing the line where he did.

What, then, justifies him in drawing the line just where he does? What justifies him in allowing the Mutilation to occur when he could have prevented it without losing any good thereby? Has the victim of the Mutilation not got a moral case against him? He could have saved her and he did not and he does not even *claim* to have achieved some good by not saving her. It would seem that God is in the dock, in C. S. Lewis's words. I, Theodore, seem to be playing the part of the counsel for the defense. And what can I say in God's defense but this: There was no non-arbitrary line to be drawn? Wherever he drew the line, there would have been countless horrors left in the world—his plan *requires* the actual existence of countless horrors—and the victim or victims of any of those horrors could bring the same charge against him that we have imagined the victim of the Mutilation bringing against him. But I see Athena stirring in protest; she is planning to tell you that, given the terms of the second-order free-will defense, God should have allowed the *minimum* number of horrors consistent with his project of achieving the Summum Bonum Creatum, and that it is obvious he has not done this—or the Mutilation would not have occurred. She is going to tell you that there *is* a non-arbitrary line for God to draw, and that it is the line that has the *minimum* number of horrors on the "actuality" side. But there is no such line to be drawn. There is no minimum number of horrors consistent with God's plan for the realization of the Summum Bonum Creatum, for the prevention of any one particular horror could not possibly have any effect on God's plan. For any n, if the existence of n horrors is consistent with God's plan for the achievement of the Summum Bonum Creatum, the existence of $n - 1$ horrors will be equally consistent with his plan. To ask what the minimum number of horrors consistent with his plan is, is like asking, What is the minimum number of raindrops that could have fallen on England in the nineteenth century that is consistent with England's having been a fertile country in the nineteenth century? England *was* a fertile country in the nineteenth century, and if God had prevented any one of the raindrops that fell on England in the nineteenth century from reaching the earth, England would still have been a fertile country. And the same, of course, goes for any two raindrops or any thousand raindrops or any million raindrops. But, of course, if God had allowed *none* of the raindrops that in fact fell on England in the nineteenth century to reach the earth, England would have been a desert. And England would have been a desert if he had allowed only one or only a thousand or only a million of those raindrops to reach the earth. And no one, I expect, thinks that there is some number n such that (1) if God had prevented n or fewer of the raindrops that fell on England in the nineteenth century from reaching the earth, England would have been fertile, and (2) if God had prevented $n + 1$ or more of those raindrops from reaching the earth, England would not have been fertile. (I *expect* no one thinks this. But the operative concept in this example is vagueness—the vagueness of fertility—and vagueness is *such* a puzzling topic that philosophers who have banged their heads against it for a while have said some very startling things. *Very* able philosophers, for example, have been heard to say that there is a perfectly sharp cut-off point between "being tall" and "not being tall," a height such that someone of that height is tall and

someone one millimeter shorter is not tall, even though no one knows and no one can find out precisely where that cut-off point lies. So perhaps there *is* someone—a philosopher, of course—who would say that there is and *has to be* a smallest number of raindrops that could have fallen on England during the nineteenth century consistently with England's having been a fertile country during that century. Well, if there is such a philosopher, that philosopher is wrong. Philosophers *can* be wrong. Descartes was wrong when he said that animals felt no pain, and anyone who says that there is an exact boundary, a boundary whose width is a single raindrop, between a fertile nineteenth-century England and an infertile nineteenth-century England is as evidently wrong as Descartes.) Here is a simple analogy of proportion: a given horror is to the achievement of the Summum Bonum Creatum as a given raindrop is to the fertility of England in the nineteenth century.

And this is why God did not prevent the Mutilation—insofar as there is a "why." He had to draw an arbitrary line and he drew it. And that's all there is to be said. This, of course, is cold comfort to the victim. Or, since we are merely telling a story, it would be better to say: if this story were true and known to be true, knowing its truth would be cold comfort to the victim. But the purpose of the story is not to comfort anyone. It is not to give an example of a possible story that would comfort anyone if it were true and that person knew it to be true. If a child dies on the operating table in what was supposed to be a routine operation and a board of medical enquiry finds that the death was due to some factor that the surgeon could not have anticipated and that the surgeon was not at fault, that finding will be of no comfort to the child's parents. But it is not the purpose of a board of medical enquiry to comfort anyone; the purpose of a board of medical enquiry is, by examining the facts of the matter, to determine whether anyone was at fault. And it is not Theodore's purpose to provide even hypothetical comfort to anyone. His purpose is to determine whether the existence of the Mutilation is a cogent argument for the conclusion that God is at fault—or, rather, since by definition God is never at fault, to determine whether the existence of the Mutilation is a cogent argument for the conclusion that an omnipotent and omniscient creator of a world like this would be at fault.

It is perhaps important to point out that we might easily find ourselves in a moral situation like God's moral situation according to Theodore, a 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% 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 one thousand dollars more than it in fact plans to spend. And the state can *always* find another thousand dollars, and can find it without damaging the economy or doing any other sort of harm. But this example takes us into the troublesome real world—troublesome because the real world is a world of all but infinite complexity, and, if we talk of the real world, we shall never make an end of it for there is always more to be said. I offer in place of this real example an artificially simple philosopher's example. If this example does not correspond very closely to anything in the real world, it can at any rate be discussed within the restricted scope of an essay like this one.

> One thousand children have a disease that is fatal if untreated. We have a certain amount of a medicine that is effective against the disease. Effective, that is, if the dose is large enough. If we distribute our store of medicine evenly, if we give one one-thousandth of the medicine on hand to each of the children, all one thousand of them will certainly die, for one one thousandth of the medicine is definitely too little to do anyone any good. We decide to divide the medicine into N equal parts (N being some number less than one thousand) and divide it among N of the children. (The N children will be chosen by lot, or by some other "fair" means.) Call each of these N equal parts a "unit." And where do we get the number N? Well, we get it somewhere—perhaps it is the result of some sort of optimality calculation; perhaps no optimality calculation is a practical possibility and some expert on the disease has made an educated guess and N is that guess. But we have *somehow* to come up with a number, for, of logical necessity, once we have decided to distribute the medicine in equal doses, a certain number of children are going to get doses in that amount. Now, since N is less than the number of children who have the disease, *whatever* we do *must* have the following consequence: at most 999 of the children will live; at least one of them will die.
>
> Now consider any one of the children who will die if this plan is carried out (the lots have been drawn but the medicine has not yet been distributed); suppose the child's name is Charlie. Our plan, as I said, is this: to give each of N children one unit of medicine. But suppose now that Charlie's mother proposes an alternative plan. She points to the N units of medicine laid out in N little bottles on the table, waiting to be distributed, and says: Take 1/N+1 units from each bottle and pour them all into one bottle and give the bottle to my Charlie—who will thus receive N/N+1 units. Then each of the N children who would have received one unit will instead receive 1- (1/N+1) units—which is (N+1/N+1) - (1/N+1), which is N/N+1. If this redistribution is carried out, each of the N children *and* Charlie will receive N/N+1 units. Now, thus algebraically represented, her plan is rather too abstract to be easily grasped.

Let us look at a particular number. Suppose N is 100. Then here is what will happen if the "original" plan of distribution is carried out: Each of 100 children will get 1 unit of medicine and live (at least if 100 was a small enough number); 900 children, Charlie among them, will die. And here is what will happen if Charlie's mother's plan is carried out: Charlie and each of 100 other children get 100/101 units of medicine—about 99% of a unit—and live; 899 children will die. Or, if you like, we can't say that this is what would have happened; we can't assert this counterfactual without qualification. (For that matter, we weren't able to predict with certainty that all 100 children would live in the actual case.) But we can say this: this is almost certainly what would have happened. "So," Charlie's mother argues, "You see that you can avert the certain death of a child at very small risk to the others; perhaps no risk at all, for your guess that N should be set at 100 was just that, a guess. One hundred and one would have been an equally good guess." We can make her argument watertight if we assume that for any determination of what number to set N equal to, that number plus 1 would have been an equally reasonable determination. That seems plausible enough to me; if you don't find it plausible, we can always make it plausible by making the number of children larger: suppose there were not a thousand children but a thousand million, and that everyone's best guess at N is somewhere around a hundred million. You will not, I think, find it easy to deny the following conditional: if a certain amount of some medicine or drug has a certain effect on someone, then that amount minus one part in a hundred million would not have a significantly different effect.

"Well," someone may say, "Charlie's mother has a good point. But the fact that she has a good point just shows that 100 wasn't the optimum value for N. The authorities should have chosen some larger number." But I had my fictional authorities choose the number 100 only for the sake of having a concrete number with which to illustrate Charlie's mother's argument. She, or some child's parent, could have presented essentially the same argument no matter what number the authorities chose, and they had to choose *some* number. And what shall the authorities say to Charlie's mother? They must either accept her proposal or reject it. If, on the one hand, they accept it, they will have to deal with *Alice's* father, who will say, "You have 101 bottles of medicine on the table, each of which contains the same amount of medicine. Call that amount a 'dose'. I want you to take 1/102nd of a dose from each bottle and give what you collect by this method to Alice." If, on the other hand, they reject Charlie's mother's proposal, they will have to condemn Charlie to death without achieving any good thereby. We cannot evade this conclusion: No matter what the authorities do, they will have to permit the death of a child they could have saved, or almost certainly could have saved, without achieving any good by permitting that child's death. It seems clear, therefore, that there can be cases in which it is morally permissible for an agent to permit an evil that agent could have prevented, despite the fact that no good is achieved by permitting that evil.

And, it would seem, this is exactly the moral structure of the situation in which God finds himself in Theodore's story, when he sets out to achieve the Summum Bonum Creatum. The argument from horrors, therefore, fails.

Let me put this question to you: Has Theodore successfully replied to the argument from particular horrendous evils? Well, much depends on what further things Athena might have to say. Perhaps Athena has a dialectically effective rejoinder to Theodore's reply to the argument from horrors. But one must make an end somewhere. The trouble with real philosophical debates is that they almost never come to a neat and satisfactory conclusion. The Dutch historian Peter Geyl once said that history—the academic discipline of history—was argument without end. This is a more appropriate description of philosophy than of history—for many historical facts are beyond dispute, and in philosophy very little indeed is beyond dispute. I do think this much. If Athena has nothing more to say, then a disinterested audience who have witnessed her debate with Theodore should render a "Scotch verdict"—"not proven"—as regards premise (3) of the argument from horrors and the moral principle on which it is based, namely that, if it is within one's power to prevent some evil, one should not allow that evil to occur 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.

If this much can be established, then, it would seem, the argument from particular horrendous evils can be answered if the argument from the fact of evil can be answered. Despite the fact that these two arguments are importantly different, theists considering arguments from evil may turn their full attention to the argument from the fact of evil.

Peter van Inwagen, University of Notre Dame

Note

1. See William Rowe's "The Problem of Evil and Some Varieties of Atheism, *American Philosophical Quarterly* 16 (1979), pp. 335–341.

2. Ibid., p. 337.