

The Problem of Evil

In this chapter, I'm going to look at *the* argument against the existence of God, the Problem of Evil. Why am I calling the Problem of Evil 'the' argument against the existence of God? Have I already covered some? Did you miss me doing so? Have you inadvertently skipped a chapter or two?

Just as Kant divided arguments for the existence of God into three kinds: those that begin from determinate experience; those that begin from indeterminate experience; and those that begin from pure categories a priori, so one could divide all arguments against the existence of God into three kinds too: those that begin from determinate experience, which will be versions of the Problem of Evil, if evil is understood in a broad enough sense; those that begin from indeterminate experience, from the mere fact that there is a universe; and those that begin from pure categories a priori, i.e. those that seek to show that there is some incoherence in the concept of God.

I say that the Problem of Evil is *the* argument against the existence of God because I think I've already covered arguments from pure categories a priori, i.e. any argument that would seek to establish the incoherence of the concept of God, in my first five chapters, where I argued that 'There is a God' made sense. One might argue for the non-existence of God from indeterminate experience, i.e. from the mere fact that there is a universe, by relying on the principle that if there were a God, he would have good reason not to create any universe at all. However, given our analysis of what it would mean to be perfectly good, one may dismiss such arguments very quickly: it does not seem at all plausible to say that God (were he to exist) would have been under an obligation to create no universe whatsoever or that it would have been good for him to create no universe whatsoever. To whom could he have been under this obligation? For whom could it have been good? *Ex hypothesi*, there would have been nobody else around and he could hardly be said to be obliged to himself not to create or to harm himself by bringing others into existence.

The only sort of argument against the existence of God that's left is thus some version of an argument that starts from determinate experience, that starts from some feature of the world that there is *prima facie* reason to suppose the theistic God would not have wanted to create, a feature that is, in other words, what we might call 'evil' if we allow the word 'evil' a rather stretched sense, to include

anything that is in any sense bad. It is important to stress the breadth of this use of the word 'evil', since evil in the everyday sense suggests malevolent intention, something that does not follow from the wider sense operative here, where as well as moral evils (bad things for which agents other than God are morally culpable, e.g. murders) there might be natural evils (bad things for which no agent other than perhaps God is morally culpable, e.g. deaths due to disease). Taking evil in this the broadest of senses then, the existence of evil in the world seems—at least *prima facie*—good evidence that there is no God, indeed it seems overwhelming evidence that there is no God.

1. **God is by definition omnipotent and perfectly good.**
2. **Evil is by definition that which is to some extent and in some respect bad.**
3. **God, being omnipotent and perfectly good, could never be compelled or have any reason to bring about or allow to be brought about something that was to any extent and in any respect bad, i.e. evil.**
4. **So, if there were a God, then there would be no evil.**
5. **There is evil.**
6. **So there is no God.**

Presented thus, the Problem of Evil is a deductively valid argument. The premises don't just make the conclusion—number 6—probable; they make it certain. So the theist—committed as he or she is to denying number 6—must deny one or more of the premises.

Numbers 1 and 2 are definitional claims; the first is—as we have seen—true of the theistic God: omnipotence and perfect goodness are constitutive of the theistic conception of God. The second reports the rather stretched sense of 'evil' operative in the argument. In an argument one may define one's terms however one wishes, so there's nothing to be argued with there. The theist can't deny 1 or 2.

Number 3 looks very plausible, at least initially. Evil things are precisely things that there is good reason not to bring about or allow to be brought about, they are in some respect and to some extent bad. We've seen already, in discussing God's perfect moral goodness, that God always does that which he has most reason to do. Surely then the definition of evil assures us that he'll never find himself with good reason to bring about anything evil and his omnipotence assures us that he'll never find himself having to allow any evils to occur.

Number 4 is a sub-conclusion: it follows from 1, 2, and 3. So the theist can't deny 4 unless he or she has more basically denied one or more of 1, 2, and 3.

Number 5 is pretty obviously correct. If you think you don't believe it, ask someone to assist you with your philosophy of religion by punching you as hard as they can in the most sensitive bits of your body: that'll soon change your mind. Remember that we're taking evil in a broad sense to include anything that is in any sense bad and, as such, suffering physical pain is certainly an evil.

Given numbers 1–5, the conclusion that there is no God, number 6, drops out deductively.

The premises and the deductive validity of this argument are more obviously correct than is its conclusion, that there is no God. In other words, it looks as if the Problem of Evil is a good argument for the non-existence of God. Is there any way for the theist to show that it's not good after all? I shall argue that there is. In fact, I shall argue that the existence of evil does not even support the claim that there is no God. My strategy will be to look at what God's perfect goodness requires of him in his creation and show that this is much less than the proponent of the Problem of Evil (as an argument supporting atheism) supposes.¹



The traditional theistic picture has God entirely unconstrained—perfectly free—in what world he creates. But, as we have seen, God's perfect freedom differs from our imperfect freedom in that it entails that he cannot do that which he ought not to do and, further, he must always do the best thing for his creatures (whenever there is a best) or one of the joint best (whenever two or more are equally good and none better). It has seemed to some to follow from this that if there were a God, he would have created the best world that is logically and metaphysically possible (if there is a best of all possible worlds) or one of the joint best (if there are two or more equally good and none better). And if we were to accept the principle that one cannot be morally justified in doing a particular thing if there is something better that one knows about and could equally well do, we'll have to conclude that if there were no best (or joint best) of all possible worlds, God—to preserve his perfect goodness—would have to do nothing, create no world at all.² If we accepted this argument, then we'd have to conclude that the theist is committed to this world's being the best or joint best of all possible worlds. But in fact we shouldn't accept this argument, though some theists (notably Leibniz) have accepted it. It doesn't work because God's perfect goodness entails only that he do the best (or joint best) *if there is one for his creatures*.

Prior to the creation of a universe, there were, *ex hypothesi*, no creatures around for whom the question of God's doing the best or joint best could even arise; there was no creature who could either benefit or suffer from the continuing absence of a universe or from its creation. In particular we, as not yet existing, were not in a better or worse state than we are now—a state the improvement or diminution of which God could effect by bringing us into existence. We were not in a better or worse state prior to the creation of the universe *not* because we were in the same state, but because we were not in *any* state—we did not yet exist. And although it might be good or bad for those who do not yet *but will* exist if one does or fails to do certain things for them (e.g. put some money aside or fail to put some money aside for their education), it cannot be good or bad for them to bring them into existence. God cannot then be said to have had a reason for creating a world stemming from his perfect goodness towards his creatures. His perfect goodness is a matter of his perfectly fulfilling

the demands of love *towards his creatures* and, prior to his creating a universe, there were no creatures who could make any such demands of him.³

The most analogous earthly situation I can think of is that of the choice, which I imagine most couples face, of whether or not to have a child. If one were to specify carefully various conditions (that there are no health risks involved in the potential mother conceiving; that having a child would not be financially ruinous to the couple or in some other way reduce their ability to meet their obligations; etc.), then it seems reasonable to suppose that there is no obligation and nor would it be better either to produce or to refrain from producing a child: they ought to be morally indifferent. They cannot show love to their 'possible child' by their decision to conceive or not to conceive, to make that possible child actual. It is not supererogatory if they do or if they do not have a child.

Imagine now a drug becoming available. It costs nothing; has no side effects; and the consumption of it affects one's gametes such that the more of the drug one takes, the healthier, more intelligent, etc. any child conceived is. With the arrival of this drug, no couple comes under an obligation and nor does it become better for them—a supererogatory act—to refrain from having any child at all just because it is now true of any child that they do have that they could always have had 'one better' by taking more of this drug. So, by analogy, even if it were true that for any possible world, God could always create a better, it would not follow that his perfect goodness would compel him not to create any world. It might be helpful in driving this latter point home were I to introduce you to Leibniz's ass, a hypothetical donkey that is a close cousin of a more famous donkey, Buridan's ass.

Buridan's ass was a donkey that, finding itself equidistant from two equally nourishing bales of hay, reasoned correctly that it had no more reason to eat one rather than the other. It then went on to conclude that the only reasonable thing for it to do was eat neither; it thus starved to death. Leibniz's ass was a donkey which found itself equidistant from an infinite number of bales of hay, such that for each of these bales of hay there was one more nourishing. It thus reasoned correctly that of any particular bale of hay it might eat it had less reason to eat that bale than it did to eat another. It then went on to conclude that the only reasonable thing for it to do was to eat none; it thus starved to death.

So, if there is a best of all possible worlds, God is not under an obligation and neither is it supererogatorily good for him to create it, for prior to his creation there are no creatures to whom he can have obligations or be supererogatorily good. If there is no best of all possible worlds, God is not under an obligation and neither is it supererogatorily good for him to create nothing just because for any world he does create it is *ex hypothesi* true that he could have created one—indeed, an infinite number—better. So far, it's looking as if God's perfect goodness doesn't constrain him in what world he creates at all. May we conclude at least that God's perfect goodness would have compelled him to create any creature he did create in the best of all possible worlds for it (if there is a best) or

in one of the joint best (if there are two or more that are equally good and none better)? An affirmative answer to this question is much more plausible, at least initially (in a moment, I'll argue against it).

If a donkey found itself equidistant from any number of bales of hay but one of those bales was the most nourishing bale possible (or two or more were joint 'most' nourishing), the donkey would be less than fully reasonable if knowing of this it then chose to eat any bale other than this one (or one of these ones). By contrast, some have held that even if there is a best or joint best of all possible worlds for a certain creature, God's perfect goodness necessitates only that if he chooses to create that creature, he must choose which world to create that creature in from among those worlds that are 'good enough', a world's being good enough if in it that creature leads a life that's not so bad that it would have been better for it if it had never existed. But this does not seem plausible to me because, for the reasons sketched previously, the notion of a creature's being potentially better off (or worse off) if it had never existed seems confused. It's as confused as speaking of the brother that I never had being taller, less tall, or the same height as the sisters that I do have. The brother that I never had is not on the height scale at say zero feet and zero inches; he is not on the height scale at all. Thus his height cannot be compared with the heights of the sisters that I do actually have and who—being actual—do indeed have particular heights on this scale. Of course I can say things like, 'Had I had a brother, the chances are that he would have been taller than either of my sisters', but, as it is, the brother that I never had is not taller than, less tall than, or the same height as my sisters for he doesn't exist at all. Similarly, the brother that I never had is not less well off than, better off than, or as well off as the sisters that I do have. So any creature that does exist is not better off, worse off, or enjoying the same level of wellbeing as if he, she, or it had never existed.

If there is a best of all possible worlds for a particular creature, while God would have been morally indifferent about whether or not to create that creature at all, it seems then as if we should say that if he does create that creature, he has reason to create it in that world rather than any other; and if there are joint best worlds for it, if he creates that creature, he has reason to create it in one of those rather than any other.

Of course, even if we were to say this, we still could not directly conclude that theism is committed to this world's being the best or joint best of all possible worlds for each of us. Perhaps for any creature (that actually exists), whatever world it might exist in, there's always a possible world that that creature could have found itself in instead and that would have been better for it than that world. As we've already seen from considering the case of parents who could take a drug to 'improve' what sort of child they conceived and Leibniz's ass, if this were the case, then God's perfect goodness would not dictate that he not create this creature. His perfect goodness only dictates that he do the best or joint best for his creature where one is possible. It might be then that there is no best or

joint best of all possible worlds for us and thus God's perfect goodness left him with *carte blanche* not just over whether or not to create us at all but over what world to create us in having decided to create us. Matters would be rather as they would be with a more acute version of Leibniz's ass. Although sadly Leibniz's ass did not realize it, he had *carte blanche* over which bale of hay to eat. But if there is a best or are joint bests of all possible worlds for us, then while God would still be able, without deviating from perfect goodness, not to create us at all, it would *prima facie* seem that we should say that he did not have *carte blanche* over the issue of in which world to create us. Why? Because, as we have just seen, he has good reason to create each of us in the world that is the best (or a joint best) of all possible worlds for us if there is one (or more than one joint best). That theists view the world as God's creation might therefore seem to commit them to its being the best or joint best of all possible worlds for each of its creatures or there being no best or joint best of all possible worlds for those of its creatures for whom it's not. But in fact, I'm about to argue, it doesn't commit theists to even this. We can start to see this by asking ourselves this question:

Does the fact that God's perfect goodness entails that if there is one, he must create the best or joint best of all possible worlds for his creatures mean that he must create each and every creature that is actually created in what is for each of them considered in isolation the best or joint best of all possible worlds (if there is such a world), or does it entail that for possible worlds each of which contains a given set of creatures, if he creates that set of creatures he must create them in the best or joint best of all possible worlds for this set of creatures considered as a totality (if there is such a world)? I shall maintain that it is only the latter and that he could do the latter without doing the former. This being so, we may conclude that God's perfect moral goodness does not dictate that he create any creature he does create in the best (or one of the joint best) of all possible worlds for it even if there is a best world or are joint best worlds for it. His good reason to create a particular creature in what is for it the best or joint best of possible worlds could be outweighed or balanced by good reasons to do the same for other creatures.

Consider two possible universes, A and B, in each of which live two creatures, P and Q. In universe A, creature P has freedom to do that which isn't the best that he could do for Q and freedom even to do certain things that he oughtn't to do to Q. We've already seen that having this sort of freedom is in itself a good thing for P. To have this freedom over Q necessitates P having more power than Q in certain respects and Q not having certain powers. For example, if P is going to have the freedom to insult Q, Q can't have the power to stop himself being insulted merely by willing it. In universe B, their roles are reversed: Q has power over P to the same extent and in the same respects as P has it over Q in world A. If these are the only differences between the universes, we may say that A is then a better universe for creature P than B to a certain extent, *e*, but it is not so good for Q and world B is better for Q than A but less good for P by the same extent, *e*. Of course God could create P and Q in two different universes, but he

cannot—of necessity—create them in two different universes yet give them this freedom to influence one another; their having the freedom to influence one another necessitates their being in the same universe. Let's now suppose for the sake of argument that A is the best of all possible universes for P and B is the best of all possible universes for Q. Given what we have said earlier, these facts give God a reason to create P in A rather than anywhere else and to create Q in B rather than anywhere else, but of course it is logically impossible that God act on both these reasons and these reasons are equally strong, their strength being determined by *e*. What would his perfect goodness dictate that he do then? Create neither? If it did dictate this, it would dictate that he ought never to create any universe in which one or more creature had freedom to affect one or more other creature for good or ill. But that seems wrong. Imagine this:

You are a donkey-herder. Your herd is small. You have only two donkeys, P and Q. One day you find yourself with your two donkeys equidistant from two bales of hay. Bale A would be slightly better (to a certain extent, *e*) for donkey P than bale B and slightly less good (to the same extent, *e*) for donkey Q than B; bale B would be slightly better (by *e*) for Q than bale A and slightly less good (by *e*) for P than A. You cannot herd P to bale A and Q to bale B; you have to choose between the two bales. Should you have never let yourself get into such a situation in the first place? It doesn't seem at all obvious that you were under an obligation not to allow yourself to get into such a situation. Indeed it seems obvious that you weren't. There's a weakness in the analogy, in that if you don't take your donkeys to some bale, they'll both starve. So let's remove that disanalogy by supposing you're a potential donkey-herder; you're about to choose whether or not to (non-ultimately of course) create a set of donkeys. You find yourself knowing that for any set of donkeys that contains more than one member that you create, you will one day face such a choice. Should you therefore create one or no donkeys, thus ensuring you never get into a situation such as that described, where you have to do something less good than you could do for one of your donkeys? It's not obvious that you shouldn't, especially if creating two or more donkeys would give each donkey a good—e.g. the possibility of donkey friendship—that it would not have been able to have had it been created on its own.

We've already seen that for finite creatures (though not for God) the freedom to be less than perfect in two ways (doing less than the best we could and less than we ought for someone) is a power; it's something that it's good for us to have. Of course, as we've also seen, it's not the only thing that is or would be good for us to have, the ability to avoid being insulted by others simply by willing it would be an ability it would be good to have. So, God's perfect goodness would not dictate that he not create either world A or world B, even though each of these worlds has as a feature that it is not the best of all possible worlds for all of the creatures in it and there are *ex hypothesi* best of all possible worlds for each of its creatures. If the good of freedom to be less than perfect to

one another required creating creatures in a world with some evils, then God's perfect goodness might allow him to create a world in which there were the evils necessary for it. (It could not compel him to do so as he'll always remain perfectly free to create no world whatsoever, but remain instead the sole existent thing.)

Not all evils are brought about by agents acting freely in blameworthy ways. As well as cold-blooded murders, there are deaths due to disease or accident. How are these, what we might call 'natural evils' in contrast to 'moral evils' to be explained on theism?



My argument is that natural evils are a necessary result of there being free creatures living in a world governed by natural laws and that natural laws are necessary for there to be a world with agents who enjoy the freedom to be less than perfect to one another. Natural evils are the inescapable accompanying features of natural laws, natural laws being the necessary means to the good of this sort of freedom.

Suppose, for example, that P wishes to exercise his freedom to choose to do what he knows is less than the best he could do for Q. In fact he knows it's something that he shouldn't do to Q. It's causing Q to suffer ten minutes of excruciating agony just because he doesn't like the cut of Q's jib. Either P will get his wish—in which case Q will find that natural facts less than perfectly serve his interests; Q will find himself without enough power to stop P—or Q will be able to block P's malevolent intention—in which case P will have his interests less than perfectly served by natural facts; P will find he doesn't have enough power to harm Q as he'd like. If one agent is to have the freedom to choose to do evil to another, then that agent must have more power than the other in the relevant domain. And the fact that one agent has more power than another must be the result of facts that are not themselves within the power of those agents to determine, i.e. they must be natural facts. One can say then that natural evils are a foreseen but unintended necessary consequence of creating free creatures in a world with natural laws, natural laws being necessary for there to be free creatures able to choose to affect one another for good or ill. Natural laws provide the arena within which significantly free agents may interact, and a necessary feature of that arena is natural evil.⁴

So, our freedom to be less than perfect requires natural evils as well as moral evils. This in itself might not be of comfort to the theist. That there are those who suffer in any system of interrelating free creatures God might create doesn't mean that he's morally justified in creating a system of interrelating free creatures who suffer to the extent that some creatures in this world suffer. Perhaps, while his creatures having the freedom to affect one another in this way is itself a good, it's not a good that's good enough to justify the sorts of evils that we actually have in this world, the level of suffering that some of its inhabitants

undergo. And perhaps—even if it is good enough to justify this—God doesn't have the right to create a system where some suffer to this extent.

Let's turn to address ourselves to these worries. To do so it will be helpful to introduce the notion of a good 'compensating' for an evil.

The notion of a good compensating for an evil is a rather tricky one and not just for the epistemic reason that it may not always be obvious whether a good really does compensate for an evil. It's also tricky because the compensating good may not be the same sort of good as the evil is evil and thus may not be said to outweigh it in any even-in-principle-quantifiable way. This will be easier to understand if I give another example.

Suppose you have a choice to make. You can become a sculptor or you can become a painter. Suppose also that you know (don't ask me how you know this) that if you choose to become a sculptor, you will become a truly great sculptor—on a par with Phidias or Henry Moore—but you'll suffer from the occasional bruised finger as your hammer goes awry during your chiselling. You also know that if you choose to become a painter, you will become a truly mediocre painter, with slightly less bruised fingers than you'd have had if you'd become a sculptor. You also know that apart from these differences each life will be the same for you.

If this was the choice that faced you, I think we would all agree that the physical pain of a few bruised fingers would be outweighed by the greater good of your being a truly great sculptor and not just outweighed for others: it would be outweighed for you too. Your being a great sculptor, even though it would mean being someone who suffered the physical pain of an above-average number of bruised fingers, would be a better life for you to lead than your being an average painter with a lesser number of bruised fingers. So the good of being a truly great sculptor is a greater good than the evil of a few extra bruised fingers is bad, but being a great sculptor isn't a physical pleasure which can be straightforwardly weighed against the physical pain of the bruised fingers. So, although there is a sense in which being a truly great sculptor compensates for the physical pain of some extra bruised fingers, this is a sort of compensation that can't be represented as an outweighing on some common scale.

Now one could in principle become a great sculptor without bruising any fingers and even if that never happens in practice, it's not the bruising of the fingers that makes one a great sculptor anyway—it's having a set of skills that one develops while, as a matter of fact, bruising a few of one's fingers along the way. The bruising of one's fingers is a contingently accompanying feature of a contingent means to the end of becoming a great sculptor. Let me suppose for a moment though that actually having had a few more than average bruised fingers is physically *necessary* for being a great sculptor for some reason—perhaps you just can't hold your tools properly unless your body has in some sense instinctively learnt how by doing actions that must bruise its fingers more than most. If some reason like that did obtain, then having bruised fingers would be a

physically necessary accompanying feature of what was as a physically necessary matter of fact the only means to the end of becoming a great sculptor. In that case then we'd say that the good end of being a great sculptor would justify the bad features accompanying the means; it would compensate for them. Of course, even then, the accompanying evil to the only means of becoming a great sculptor wouldn't be a *logically* or *metaphysically* necessary accompanying feature to the only logically or metaphysically possible means to that end—God could have miraculously given you this instinctive ability to hold your tools without your having learnt it in the 'School of Hard Knocks' as it were. Nevertheless, the example serves to illustrate the point that a certain good can compensate for a certain evil when that evil is either a means to it or an accompanying feature of that means and that this compensation need not be a matter of giving one a greater amount of the same sort of thing that the evil has deprived one of.

Consider now this story: Once upon a time, a little fawn—let's call him 'Bambi'—got caught in a thicket in a forest. Bambi struggled for a while, but in the end realized that he could not get out of the thicket on his own. Not to worry. He waited for his friend, the rabbit Thumper, to bounce playfully along (as was his habit) and help him out. Unfortunately, that day Thumper was bouncing happily in another part of the forest and a forest fire had broken out near to Bambi. Bambi yelped as loudly as he could, trying to summon Thumper's help; but sadly Thumper was far away and the fire was getting closer. There was nothing the panicking Bambi could do; he struggled wildly to escape, but to no avail; the fire reached him and slowly started to burn him alive. Eventually, Bambi died in excruciating pain; nobody ever discovered his body. All the other animals in the forest lived such carefree lives that they didn't even think about where Bambi might have got to; even Thumper simply bounced playfully about as he always had done. The End.

The fire certainly wasn't the means to an overall good end for Bambi and if nobody ever discovers what's happened to Bambi and nobody even thinks about his absence, then it can't produce any effect on anybody else; *ipso facto*, it can't be a means to any good effect for anybody. So the forest fire burning Bambi to death isn't a means to any good end that compensates for it. But this does not mean that it itself cannot be compensated for.

Just because Bambi suffered an evil that was uncompensated for in this world does not mean that he suffered an evil that was uncompensated for full-stop. As we've seen, on theism, there is another world, after this one, in which these loose ends are tidied up. There's something rather pleasingly airtight about this move. Unless there is some conceptual absurdity in maintaining that God could arrange for Bambi to enter a heaven after this life, on theism there's every reason to think it's true that he does arrange for him to do so; and surely then we cannot have any reason to believe that in this heavenly realm there couldn't be compensating goods. Heaven is, after all, of infinite duration. Whatever someone suffers in a finite pre-mortem life, it has to be possible that they be compensated

for it eventually in an infinitely extended post-mortem life. I conclude then that any evil a creature suffers in this world could be compensated for by God in the next. Of course, to establish that any evil that befalls a creature in a finite life could be compensated for by God in an infinite afterlife is not to show that all the evils in the world are necessary as means to those compensating goods. In our example, God could have arranged for Bambi to get into Heaven—and so have the goods that, as it is, constitute the compensation—after a quick, painless death, these heavenly goods thus not being needed to compensate him for anything, they being bonuses. Wouldn't that have been better for him? It pretty obviously would have been. If the theist were committed to there being no evil in the world that is not necessary as a means to a good end that compensates for it, theism would thus be untenable. However, the theist is—as we have seen—not committed to this. He or she is committed to there being no evil in the world that is not necessary as a means to a good end that compensates for it or as an accompanying feature of such a means. Bambi's death was not a means to any good end, for Bambi or anyone else, but it was a consequence of the laws of nature operative in the universe in which he lived, the laws of nature being necessary as means to the good of the freedom of the creatures in that universe to be less than perfect to one another. To have this sort of freedom requires, as we have seen, natural laws—i.e. laws that operate independently of any creature's will—and these laws must thus give rise to natural evils, suffering for which no agent (other perhaps than God himself) is responsible. The only question that can remain, then, is whether God has the right to create a universe where creatures like Bambi suffer to the extent that Bambi does as a result of natural laws, the operation of which is for the greater good of creatures as a whole. Would God's moral perfection compel him not to create a world where creatures suffer in this way as a result of the system? We are not talking of whether he has the right to create a universe in which he himself uses some creatures merely as a means to the end of the freedom of others (for Bambi isn't used as a means to anybody's end) but of whether he has the right to create one where he allows nature to 'take its course' and thus generate the suffering of Bambi, suffering that is a foreseen but unintended consequence of the laws of nature that he creates as necessary as the means to the end of the freedom of some of his creatures.⁵



Consider this situation: you are a teacher in charge of a group of schoolchildren at playtime. We've established that it is good for these children to have freedom to do less than the best that they can for one another and indeed to do what they should not do to one another. That being so, you have reason to stand in a corner of the playground and let them invent and play their own games with one another, rather than ceaselessly stop them from interfering with one another and organize them 'for their own good'. Let's suppose you allow that reason to guide you. You stand to one side. Now and again, you notice that some of the children

are choosing to use the autonomy you've generated to invent games that involve some of them suffering to some, limited, extent. There are, one might say, 'victims' of your laissez-faire system. Let's take an example: one of the children is chosen by mob rule to be piggy in the middle for some game. This is a role that is considerably less fun than the other roles, indeed it involves positive suffering; the child thus chosen suffers to some extent as a result of your system. Perhaps the child's character is developed in helpful ways by this experience; but, then again, perhaps it is not. Let's suppose that it is not and that neither is there any other greater good for the child or the children in general that comes from his or her suffering in this way on this occasion. You watch this happen. You maintain your distance. You do not intervene. Such eventualities are—after all—a foreseen but unintended consequence of the laissez-faire system that you've adopted. This child's suffering is not itself a means to a greater good that compensates for it. You've not used this child as a means to anything, but you have allowed the child to suffer when you could have stopped it.

Did you have the right to allow this to happen? Well, I suggest that the answer to this question depends on a number of things. One of these is how much suffering the child has actually undergone. If the game you allowed the children autonomy to develop had been a William-Golding-esque one involving the piggy in the middle being killed, then obviously you should have intervened; you'd have done something wrong in allowing the children in your charge to have *that much* freedom and power over one another. If, on the other hand, the suffering was of a relatively minor sort—a sort that would all be forgotten about within five minutes or so of the start of the next lesson—then, it strikes me, the answer is that you wouldn't have done anything wrong in taking this laissez-faire attitude, in allowing this child to suffer to the extent that he or she did. The child could have had a better playtime, but he or she has no cause to complain to you as a result of this.

So, our question must be, 'What determines how much evil you have the right to allow creatures in your system to suffer?'

I suggest that there are three relevant factors.

First, it depends on how good the overall effect of your system is. If it's really very good indeed that these children have the amount of freedom your system provides, then that will make it more morally justifiable for you to have allowed those who suffered in your system to suffer when you could have intervened. If, on the other hand, it's not that important whether or not those in your system have the level of freedom it provides, then that will make it less morally justifiable for you to have remained distant when one of them suffered as a result of your having given others the level of freedom you had given them. You could have intervened; stopped the suffering; and not thereby deprived anybody else of anything that valuable. So that's one factor.

Secondly, it depends on your capacity and intention to provide compensation for those who suffer in your system. If you know that after playtime is over you

will compensate those who've suffered, that will make it more morally justifiable for you to allow them to suffer to the extent that you have. Conversely, if you know that you will not compensate sufferers for the suffering they've undergone, that will make you less morally justified in allowing them to suffer to this extent.

Thirdly, it depends on whether or not the people in question have refused to participate in your system. If, knowing of the sort of laissez-faire attitude you were going to adopt, the children had all agreed to be participants in the system, that would make you more morally justifiable in subjecting them to it. Conversely, if, hearing of the sort of system you were going to adopt, a child had asked you if he or she could stay inside this playtime, that would make you less morally justified in throwing that child out into the playground anyway.



Let me leave those three factors on the table for a moment or two and turn to consider things from the child's perspective.

Imagine now, then, that, rather than being the teacher, you're a child. On arriving at school one day, you're greeted by the headmaster. 'Today,' the headmaster tells you, 'is a special day. You have a choice of which playground to play in. There are a number of playgrounds. In each playground the supervising teacher will adopt a certain level of this laissez-faire approach. In playground one it's zero. Each child is completely controlled in their every movement by teaching assistants, who guide the children's limbs inside the cotton-wool suits that every child wears. No child ever has freedom to be less than perfect in his or her relations with others, but then again, of course, there's nobody who suffers to any extent as a result of the choices of others. Playground one guarantees those children who reside in it that they won't suffer as a result of the system in the sense relevant here to any extent whatsoever. In playground two, there's a little bit of freedom. Every ten minutes, each child is taken out of their cotton-wool suits and allowed ten seconds in which they can act as they wish; thus, occasionally, one of these children uses their autonomy to punch another. Playground three has a bit more autonomy and thus offers a bit more danger of suffering than does playground two. And so on.'

'There's another feature of the meta-system we're running today,' says the headmaster. 'Each person who suffers will be compensated for any suffering after playtime is over. So, those who've been in playground one won't need any compensation. Some of those who've been in playground two by contrast will have suffered as a result of the system there and thus they will need some compensation, but on average not as much as those who've suffered in playground three and so on. I want to stress, though, that no child—whichever playground they've been in—will leave school at the end of the day having suffered in a way that he or she will think has not been adequately compensated for.'

You thank the headmaster for apprising you of his meta-system and consider which playground you'll sign yourself down for.

Is the only playground it would be rational for you to choose to play in playground one?

I think the answer to that is 'no'. I'll justify that negative answer later. Before I do so, I want to speak to what I imagine are a large number of you who are growing rather impatient with my analogy for another reason.



There's a crucial difference between my headmaster analogy and our case. God didn't ask any of us whether we'd mind being put into the universe he'd created. The headmaster—as it were—didn't ask us to choose a playground at the start of our school day; he just threw us into one, this one.

Someone might object, then, that even accepting that the first condition is met (that overall the level of freedom we enjoy really is worth the level of suffering necessary for it) and that the second condition is met (God can and does provide all of us with sufficient compensation for our sufferings in an afterlife), the third condition isn't met, God didn't ask us beforehand if we would be willing participants in the system he was about to create. That's a crucial disanalogy between God and the headmaster case. And this shows that God didn't have the right to put us into this world. It is indeed true that there's this disanalogy, but there's another crucial disanalogy. God couldn't—of necessity—have asked us in advance of our existence whether or not we'd be willing to suffer the evil that our existing in this world would entail for the simple reason that we didn't exist in advance of our existence. Does this let him off the hook, morally speaking, with regard to the third condition?

Can we find an analogy to guide our moral intuitions here? I think we can. The analogy is again the choice of whether or not to have children.

Ours is a world where there is a significant risk that any children we bring into existence will suffer. We can't guarantee that the system is overall worthwhile; or that they're going to get enough compensation if they suffer as a result of it; and we can't ask our children before they're born whether or not they're willing to be born into the world on these terms. Nevertheless, we do not regard ourselves as generally under an obligation *not* to have children. We certainly don't regard ourselves as under an obligation not to have children simply because we can't ask them in advance of having them whether or not they're willing to be born. So, I conclude that God not—of necessity—being able to act as the headmaster does, and ask us in advance of our existence whether or not we're prepared to take the risks that existence will bring, *does* let him off the hook morally speaking with regard to the third condition.

God's let off the hook with regard to the third condition and he can easily satisfy the second; as we've already observed, any finite amount of suffering a creature undergoes in this world must be capable of compensation ultimately in an infinitely extended afterlife, something that we saw in the first half of the book God must—if he exists—extend to all creatures to whom it would be good to

extend it. The only question that remains, then, is whether the level of freedom to be less than perfect that we enjoy is overall worth the suffering it entails. How to answer this question?

Unfortunately, the answer that one gives to this question will depend entirely upon the probability one has previously assigned to theism. If one is asked by one's host at a dinner party whether one would like to try a dish that is certainly different from anything else one might have later and that some people enjoy even though others violently dislike, one's answer might reasonably depend on whether this dish is being offered to one as an option for the hors d'œuvre or for the main course. If one is told that it is an option for the starter and one is assured that the taste (if it is found to be unpleasant) may be washed away very quickly by the drink accompanying the main course, one would no doubt try it. If, on the other hand, one is told that the dish is an option for the main course (and there will be no desert), one would be more reasonable in refusing. It is not that one would be more risk-averse, just that the risk would be greater in relative terms, for what it was relative to would be smaller. Similarly then, if one sees the suffering of this world as a prelude to an infinite afterlife of perfect fulfilment in God's presence, the chance to enjoy a freedom that we will not be able to enjoy when directly exposed to God in Heaven will be judged worth the suffering that accompanies it. However, if one sees this world as all that there is, one's judgement will differ. It seems then that the mere existence of evil cannot be taken as in itself evidence against the existence of the theistic God for it would only be so on the hypothesis that there is no compensating afterlife, a hypothesis that is false on theism precisely because of God's omnipotence and perfect goodness.



Let me tell you a bit about some ground-breaking technology that's been installed in the spine of the book you're holding.⁶ (If you're not actually holding it, pick it up.)

This book is linked up—via the internet—to a computer that is running a programme called 'The Best Life You Could Lead'. If you squeeze the book as hard as you can between both hands, then it will painlessly implant into your hands electrodes that will then send signals to your brain meaning that the ideas you have can be shaped by the programme on the computer. You will be immediately plunged into a virtual reality world. You won't realize this, because your virtual-reality world will start off by being very much like the real world. It will seem to you that you decided *not* to squeeze the book (or perhaps that you did squeeze it but nothing happened) and that you're still sitting in the room reading it, etc. But after a few moments *in the virtual-reality world* (not in the real world), a good friend will rush in through the door to tell you that you've won the lottery. In the real world of course, no such thing will be happening—you'll just be sitting with the book in your hands, a rather fixed expression on your face, oblivious to your surroundings, a source of curiosity for anybody who might wander in.

Back in the virtual-reality world, it won't seem strange to you that you've won the lottery. You won't be able to choose to reflect on whether or not this is all 'too good to be true', because the computer will have implanted a false memory of your having bought a ticket (if such is needed) and it will suppress—or rather redirect—your reasoning. Your freedom to do less than is maximally conducive to your own happiness will be eliminated—painlessly, immediately, and totally—from that moment on because the computer can guide you through the virtual world more effectively, in the sense of more optimally for your happiness, than you'd be able to guide yourself. So it is that in the virtual world it will seem to you that you've chosen to leap up; hand out glasses of champagne to your friends and family (who've all appeared); and—over the next few weeks—use your winnings to make investments that by the end of the month have given you enough political power to unite all governments and bring world peace. You'll believe yourself to have found cures for all diseases and released a number one cover version of *Sitting on Top of the World*. Of course, all this will only be happening in the virtual world. In the real world, what will have happened is that I'll have hooked-up your body to an intravenous drip at one end (feeding you nutritious fluid) and a catheter at the other (removing waste products). For the rest of your life, as in the virtual world you go from strength to strength, in the real world you'll just be getting more and more dusty, and those who use the room will have to change the nutrition and waste bags that lie beside your body every week or so.

As you decide whether or not to squeeze the book, there's no need to worry that your friends and family might not be happy seeing you hooked up in this way to the virtual-reality machine: I've got books for them too. If you decide to squeeze it, I'll get them to squeeze books of their own so that they'll be in their own virtual-reality worlds, worlds where they'll think that they themselves are world leaders; pop stars; or whatever.

Assuming you believe all this, is it irrational for you not to squeeze the book? If you asked me, I would answer this question negatively; I would say that it's pretty obviously not irrational for me to think that freedom is worth it.⁷ I may hazard that all of you reading this will agree with me in giving the question this negative answer. But we can easily imagine people differing from us in this regard; indeed, we can imagine that if our own lives were going much worse than they are, we would give the question a different answer.

Consider for example being on the torturer's table, about to be subjected to torture for twenty-three and a half hours a day for the rest of your life, with merely a half-hour slot each day in which you might freely pursue your own objectives. If you were then offered the choice of squeezing the book and thus avoiding the twenty-three-and-a-half-hours-a-day torture by renouncing any further freedom, it would seem to me pretty obviously irrational not to squeeze the book, to think that a half-hour-a-day's worth of freedom isn't worth the suffering that it would involve. Now imagine reducing the proportion of the

days ahead that you will spend on the torturer's table until we reach the cut-off point where it becomes a matter of indifference to you whether or not you squeeze the book. It is perhaps strange to posit that there is a cut-off point—rather than that the matter becomes indeterminate—but let's try not to worry about that. Once we've settled on an approximate amount that would more or less balance the corollary freedom, imagine being told that suffering a slight but non-negligible amount of torture more than this (five minutes per day say?) was actually the only way to secure freedom for the rest of humanity. That fact would, I suggest, decisively tip the balance of reasons in favour of your not squeezing the book. The good of humanity as a whole would compensate for this extra suffering befalling you as an individual. Of course, it wouldn't compensate you personally for it; it would compensate the aggregate of people that is 'the system as a whole' for your suffering. This being so, it might well be that the sacrifice was not one it would be reasonable for you to accept (unless you were assured that you personally would be compensated for having made it in an afterlife). But whether or not it's prudent or would be overall reasonable for you to choose to make this sacrifice in the knowledge that you would actually happen to have in the situation we're trying to imagine, it remains the case that your making this sacrifice would, overall, be better than your not doing so. Overall, a system where you were forced to undergo this sacrifice would be worth it, and so I suggest that if the sufferers in a system are indeed going to be compensated for their suffering and if, were they to have been fully informed beforehand, they would have reasonably chosen to participate in it (for they would have seen that overall the system that entails their suffering is worth it) but one is not able to ask them beforehand as they have yet to be created, one is morally justified in creating such a system.

If all this is right, God's perfect goodness then allows him to create universes with all sorts of evil in them. If there are creatures for whom there is a best or joint best of all logically possible worlds, he might yet create such creatures in worlds that aren't the best or one of the joint best of all possible worlds for them. He might allow creatures to suffer in ways that produce no good for them whatsoever and produce no greater good for anyone else either. The only thing his perfect goodness prevents God from doing is creating a world of creatures who suffer to an infinite extent at a given time or a world of creatures such as Tantalus and Sisyphus, who are destined to suffer to some finite extent for ever. An infinite amount of suffering can never be compensated for (even by God) either as regards the individual creature who has suffered from it or as regards the system as a whole. But it is obvious that ours is not a world in which creatures can suffer to an infinite extent at a given time or one in which there are immortal creatures destined to suffer for ever (well, it's perhaps not obvious that it's not the latter, but we've certainly no reason to believe it is—how many immortal sufferers have you ever come across (again, note that a denial of the traditional doctrine of Hell seems necessary for this point to go through)?). On theism, as

we have seen, after our finite lives here an infinite life awaits us hereafter. For every creature who suffers, there will come a day when they say that as individuals their suffering has been more than adequately compensated for and on which they will be able to see how their suffering fitted into a greater whole that was overall worth it. On that day, even those who were broken on the wheels of the machine as they turned will thank God for it.



As I recall, Herodotus tells a story of the Barbarian despot, Xerxes, talking with a general in his court about his plans to invade Greece. Xerxes asks the general how many men he thinks the Greeks would need to muster before they would dare to oppose him in battle. The general asks Xerxes whether he wants an answer that will please him or the truth. Xerxes asks for the truth. The general tells him that if the Greeks have ten thousand men, then ten thousand will fight him; if they have only a thousand, then a thousand will take to the field; if they have only a hundred, then still those hundred will stand before him. Xerxes cannot believe this, for he plans to invade with the largest army the world has yet seen. If these Greeks were under the iron control of a tyrant, such as himself, he reasons, then perhaps they might go forward, even against impossible odds, from their fear of that tyrant and his lash. But these Greeks, he has heard, are free men and freedom is the end of discipline. The general replies that the Greeks are indeed free, but this is only because they have a master whom they respect more than they could fear any tyrant. This master is their duty. This they listen to and this they obey. And what it commands is ever the same: not to retreat in the face of barbarism, however great the odds; rather, to advance against it; to stay firm in their ranks; and to conquer or die.

A world without evil would be a world where we could turn every sword into a ploughshare; it would be a world where we never needed to fight because it would be a world where there was never anything worth fighting. A world with terrible barbarians is a world where there are people worth fighting; it is a world where we need swords as well as ploughshares; and it is a world where it's open to us to choose either to go forward into battle against the barbarians like free Greeks or meekly acquiesce to them as would the craven slaves of a barbarian despot. We are free to choose to be heroes or villains, to sacrifice ourselves or to save our own skins, to do our duty or to shirk it.⁸

Would a life without any evils at all be better than a life filled with such choices? It would certainly be easier—but then a life in the virtual world I've just described is easier than a life in the real world and I don't think you think it'd be better for you to go into that virtual world. Playground one is not the only reasonable choice. Would a life with *more* terrible evils than there are in our world and thus more such choices be better than a life with less evils, but of course less of this sort of freedom as a result? Is playground number infinity the only reasonable choice? As one goes up in playground numbers one gets more

and more of this sort of freedom, but of course one gets more and more evil as a result. However, if this evil is compensated for across the system by the good of the corollary freedom and each sufferer will individually ultimately be compensated for his or her suffering in an afterlife, then—as this freedom is a good—it would seem that one should say ‘Yes’, playground infinity is the only rational choice. However, in fact here my analogy breaks down again—there is no playground infinity that God might have created. Of necessity, any creatures God could have created would have been creatures with a finite amount of freedom (in virtue of his necessary omnipotence, no creature can be as free as him). So it is that if theism is right, God was faced with a choice to create nothing; to create a world with no such freedom but no evil (Heaven straight away); or to create a world with a finite amount of this freedom and thus evil, a world in which he compensates everyone for their suffering in an afterlife (a world like ours, with Heaven afterwards). That our experience gives us reason to believe that if he exists, he has chosen the latter does not—I suggest—give us any reason to believe that he doesn’t exist.

I conclude then that the argument from the existence of evil to the non-existence of God cannot be rendered as a good deductive argument; nor can it be rendered as a good inductive argument; nor again does evil inductively support the claim that there is no God. The occurrence of evil in the world provides us with no reason whatsoever to think that there’s not a God.



It is sometimes objected that to offer a ‘solution’ to the Problem of Evil in the manner that I have done in this chapter is to blunt our awareness of the evil or at least to blunt our motivation for combating it, either of which would provide a moral reason to object to the very process of undertaking a theodicy such as that sketched in this chapter.⁹ This charge may be sustainable against some theodicies, but it is not sustainable against the one I have outlined. As we have seen, all that theism commits one to saying is that overall the system as a whole is worth it. Reconciling the existence of evil in the world with the existence of the theistic God in the manner sketched need not therefore diminish our awareness of particular instances of evil or remove our motivation for seeking to combat them. The theodicy sketched is compatible with accepting that in the actual world there are lots of evils that are completely gratuitous, that don’t lead to any good end at all; there are lots of evils which are partially gratuitous in that even though they lead to a good end which could not, even in principle, have been achieved without them, lead to a good end which isn’t good enough to compensate for the evils that produce it; and there are lots of evils that are dispensable, which is to say that even though they do lead to some good end which compensates for them, the good end in question could in principle have been achieved without them. We might very well be under an obligation or it might be a supererogatory good for us to remove some or all of these evils. In short, to

justify God in the face of evil is not to justify evil in the face of God and, if we are conceptually clear-headed, neither will justifying God in the face of evil erode our motivation for fulfilling our obligations and performing supererogatory good acts to combat evil. If theism is right, there will come a time when every sword may safely be turned into a ploughshare, but if theism is right, that time is not yet. For now, we are called to act as free Greeks.

As I think it would be needlessly evil for me not to do so, I find myself motivated to tell you where I think all this leaves us. I'll do so in the next chapter, having looked at the nature of faith.

This page intentionally left blank