

Compatibilism and the Free Will Defense

JASON TURNER

Faith and Philosophy 30.2 (2013): 125–137

The *free will defense* is a theistic strategy for rejecting a certain argument for the non-existence of God. The argument, sometimes called the “logical problem of evil,” insists that it is logically impossible for an omnipotent, omniscient, perfectly benevolent God to co-exist with evil. Since evil clearly exists, the argument goes, God does not. The free will defense responds by claiming that, since free will is very good indeed, a God with all the characteristics in question might co-exist with evil—provided that the evil was brought about by (other) free creatures and God could not have prevented the evil without making his creatures unfree. It is commonly believed (although, perhaps, seldom argued) that the free will defense only works if the creatures in question have a sort of freedom incompatible with determinism. I used to believe this. I now think it is wrong. There are at least some versions of compatibilism which can, with the help of certain plausible theses about what God can do, respond to the atheist’s argument with a version of the free will defense.

1 THE PROBLEM OF EVIL AND THE FREE WILL DEFENSE

The free will defense is best seen as a tactic in one of the many skirmishes between the atheist and the theist. The atheist begins the skirmish by saying, “God, if there were one, would have to be omnipotent, omniscient, and perfectly benevolent. But, necessarily, any perfectly benevolent being prevents any evil insofar as it is able and can foresee it. And, necessarily, an omniscient being could foresee all potential evil and an omnipotent one could eliminate all of it. So, necessarily, if God existed, there would be no evil: he would have prevented it. But there is clearly some evil in the world—someone nicked my sandwich from the office fridge this morning, for instance—so, by *modus tollens*, God must not exist.”¹

The free-will defender responds, “Wait a minute. I don’t buy your premise that ‘any perfectly benevolent being prevents any evil insofar as it is able and can foresee it.’ A perfectly benevolent being may allow some avoidable, foreseeable evil so long as a much greater good is produced thereby: a perfectly

¹My atheist is here following J. L. Mackie (“Evil and Omnipotence,” *Mind* 64 (1955)). The theist is modeled largely on Alvin Plantinga (*God, Freedom, and Evil* (Grand Rapids, Mich.: William B. Eerdmans, 1974), pp. 7–64 and “Self Profile,” in J. E. Tomberlin and P. van Inwagen (eds.), *Alvin Plantinga* (Dordrecht: D. Riedel, 1985)), although his style of response has become ubiquitous amongst theists.

benevolent doctor may allow, or even cause, the ‘evil’ pain of a vaccination in order to bring about the much greater good of prolonged life.”

“Ok, I’ll grant you your modified premise,” the atheist says. “Let it be said rather that, necessarily, a perfectly benevolent being prevents all evil within its power to foresee and prevent, unless that evil brings about a greater good. But, necessarily, an all-powerful being can get any of these supposed greater goods without relying on the evil as an intermediary: a doctor may have to rely on the pain of a vaccination in order to prolong life, but surely God can administer the vaccination without the needle.”

“Not so fast,” replies the free-will defender. “The goods God is worried about here are not prolonged lives but creatures with the ability to make morally significant choices. Such creatures are very, very valuable, and if God needs to allow a bit of evil in order to have them, they are worth it. If God were to eliminate all evil, he would have to do it by keeping his creatures from acting evilly, which would require his taking away their free will. But were he to do that, they could not make morally significant choices and a very valuable good would be lost.”

“Certainly God knows in advance how any agent is about to act in a given circumstance,” the atheist responds. “Why can’t he just take away their freedom when they’re about to cause some evil, and leave it intact when they’re about to cause good?”

“Ahh,” says the defender. “Now you’re assuming *Molinism*, the view that, for every possible creature *S*, possible set of circumstances, *C*, and possible action *A*, either the conditional ‘If *S* were put in *C*, *S* would *A*’ or the conditional ‘If *S* were put in *C*, *S* would not *A*’ is true, and God knows in advance which one it is and is able to use this knowledge in deliberating what he is to do. If Molinism is false, then God can’t see when someone is about to do evil and decide to take away their free will.² But if it is true, then it is at least possible that the true conditionals are such that God simply cannot keep *some* people free while preventing others from doing evil. The conditionals may be arranged so that, if God were to prevent this person from doing evil, then someone else would do some evil, and if God were to prevent them *both* from doing evil, then a third person would do evil, and so on. In this case, it is not a necessary truth that an omnipotent, omniscient being can get the desired result—morally responsible creatures—without the cost of some evil.” The free-will defender then goes on to give a complicated modal argument for

²Our theist is not quite right about this; Molinism could be false because God only knows the conditionals about the actually created people. Call the thesis that God knows Molinist-style conditionals about only created people *Semi-Molinism*. Semi-Molinism fares no better than regular Molinism on the dilemma’s next horn, though: there will still be worlds where the spread of Semi-Molinist conditionals are such that God can take no series of actions that will eliminate every instance of immoral action.

the possibility of a world where God cannot get the goods without allowing some evil and, presumably, wins this round.³

2 COMPATIBILISM AND THE FREE WILL DEFENSE

Let *determinism* be the thesis that the laws of nature, conjoined with any proposition describing the state of the world at any instant, entail any true proposition whatsoever.⁴ Let *compatibilism* be the thesis that agents can be free (and, therefore, morally responsible) even if determinism is true.⁵ It is often assumed that the free-will defense is available only to those who think that compatibilism is false.

Here's why. If free will, and thus moral responsibility, is compatible with determinism (the argument goes), then God could create free creatures in a deterministic universe. But clearly God gets to choose the laws and initial conditions for whatever universe he decides to create, and surely there is enough variation in the possible laws and initial conditions for him to create a universe that would contain creatures that met sufficient compatibilist conditions for moral responsibility and still never did any evil. So if compatibilism were true, God really *could* have gotten the goods—the morally responsible creatures—without allowing any evil, and the free-will defense fails."⁶

Many versions of compatibilism certainly seem to have this result. If, like Harry Frankfurt,⁷ we think that an agent acts freely if and only if their higher-order desires are aligned with their first-order desires in the right way, then God needs merely to create a world with creatures determined to act rightly and to have their higher-order desires cohere with their lower-order ones, and he's got himself a freedom-filled world where everyone always does good.

But not every version of compatibilism has this consequence. In the next few sections, I will piece together a compatibilist position according to which God may find himself unable to create a world with free creatures who al-

³Plantinga's "Self-Profile" has a version of the complicated modal argument which does not presuppose Molinism.

⁴Cf. Peter van Inwagen, *An Essay on Free Will* (Oxford: Clarendon Press, 1983), pp. 58-65.

⁵Some philosophers (e.g., John Martin Fischer and Mark Ravizza, *Responsibility and Control: A Theory of Moral Responsibility* (Cambridge: Cambridge University Press, 1999)) think that moral responsibility is compatible with determinism even if some important kinds of freedom are not. I assume here without argument that free will and moral responsibility stand or fall together; if you think otherwise, when I talk of 'free will' and 'compatibilism' imagine me to be directly talking about 'moral responsibility' and 'the compatibility of moral responsibility with determinism'.

⁶Compare van Inwagen, *The Problem of Evil* (Oxford University Press, 2006), pp. 75-76.

⁷"Freedom of the Will and the Concept of a Person", *The Journal of Philosophy* 68 (1971): 5-20.

ways do the right thing. The picture involves some specific compatibilist theses and some specific theistic theses. I do not claim the theses are uncontroversial; but they aren't obvious non-starters, either. I won't be arguing for the theses — I'm not trying to argue that the view outlined is true. Rather, I'm outlining a (not obviously hopeless) position in logical space to show that an apparent incompatibility — between compatibilism and the free will defense — isn't.

3 HISTORICAL COMPATIBILISM

3.1 *The Zygote Argument Against Compatibilism*

The compatibilist position is motivated, in part, by the need to respond to an anti-compatibilist argument. That argument begins from a thought experiment concerning two possible worlds, w_a and w_b . Both w_a and w_b have the same deterministic laws, but for one caveat: w_b has an omnipotent, disembodied demi-goddess, Diana, able to do whatever she likes, laws of nature notwithstanding.⁸ In w_a an individual, Anna, is conceived at a time t_1 , robs a bank at a time t_2 , and dies at a time t_3 . In w_b , Betty is conceived at t_1 , robs a bank at t_2 , and dies at t_3 . Anna and Betty are intrinsic duplicates of each other, and in fact the portions of w_a and w_b between t_1 and t_2 are intrinsic duplicates of each other, too.

But w_b differs from w_a in one important way. In w_b , Diana intervenes shortly before t_1 to make subtle changes in the environment ensuring that the just-about-to-come-into-being Betty-zygote is an intrinsic duplicate of Anna's zygote, and that the rest of Betty's environment will be an intrinsic duplicate of Anna's. And Diana does this because she wants Betty to rob a bank at t_2 and — being omniscient — knows that if she fiddles with the zygote and the state of the world in this way, and then stays out of the picture after Betty's zygote comes into being, Betty will be causally determined to rob the bank at t_2 .

That's the story. Now the argument: Individuals that are intrinsic duplicates throughout their lives, living in worlds that are intrinsic duplicates throughout those lives, do not differ in any way relevant to free will. But Anna and Betty are such individuals, so they don't have any free-will-relevant differences. Clearly, though, Betty didn't freely rob the bank — so neither did Anna.

More precisely:

⁸More precisely: the Diana-free state of w_b at a time t , conjoined with the laws of nature and the proposition that Diana does not intervene between t and a later time t' , entails the Diana-free state of the world at t' .

- (1) If worlds w_1 and w_2 are intrinsic duplicates between t and t' , and if S_1 and S_2 are intrinsic duplicates in w_1 and w_2 whose lives are wholly contained within t_1 and t_2 , then for any action they each perform at a time t , either S_1 and S_2 both do it freely or they both do it unfreely.
- (2) So (given the story) Anna and Betty either both rob the bank freely or both rob the bank unfreely.
- (3) Betty didn't freely rob the bank.
- (4) So Anna didn't freely rob the bank, either.

But Anna was just an agent living in a deterministic world. Furthermore, for any agent S and action A that S performs at t in a deterministic world w , there is a Diana-like possible world that's an intrinsic duplicate of w throughout S 's life and where Diana fiddles with things before S is born to ensure that S A 's at t . Repeating the argument for arbitrary A and S gets us that no agent in a deterministic world ever acts freely.⁹

3.2 *Historical Compatibilism*

The Zygote Argument is valid, and there seems no reason to doubt the alleged possibility of either w_a or w_b . (More generally, there's no reason to think there's anything amiss with the possibility of a Diana-like world for each arbitrary agent S and action A .) So if the compatibilist wants to resist the argument, they'll need to deny one of the two premises: (1) or (3).

The (3)-denier is the *hard-line compatibilist*, who insists that, Diana's meddling notwithstanding, Betty doesn't act freely.¹⁰ The (1)-denier, by contrast, denies instead that an agent's free-will-relevant properties supervene on only how the world is during her lifetime.

The most plausible denial of (1) is a form of *historical compatibilism*. Historical compatibilism holds that the property of acting freely is, like the property of being a Rembrandt or a genuine one-dollar-bill, a *historical* property. A counterfeit dollar or Rembrandt may be an intrinsic duplicate of a genuine one; what makes it a counterfit is that it has the wrong history. A dollar bill is genuine only if it was produced in the right way and under the authority of the US Mint; a Rembrandt is genuine only if it was Rembrandt who put its pigment on canvas. According to historical compatibilism, whether an action is free or not depends on the history that led up to the action.

⁹Cf. Alfred R. Mele, *Free Will and Luck* (Oxford University Press, 2006), pp. 184–195. My presentation differs from Mele's in several respects, but the essentials are all here.

¹⁰Cf. Michael McKenna, "A Hard-Line Reply to Pereboom's Four-Case Argument," *Philosophy and Phenomenological Research* 77 (2008): 142–159.

Historical compatibilism comes in many forms. Some are consistent with (1).¹¹ But others aren't. Historical compatibilists may insist that the causal history of an action relevant to its freedom extends very far back in time, even predating the existence of the agent, in which case the freedom-relevant properties of an action won't supervene on the intrinsic state of the world just during the agent's lifetime.

Why would they say this? Well, first, why do we judge Betty unfree in the Zygote scenario? Those sympathetic to incompatibilism will insist that it's because her action was *determined*. Compatibilists will resist such a diagnosis. But they don't have to look far for another. The events in w_b don't unfold 'on their own', governed only by impersonal law and culminating in Betty's actions. Rather, another agent intentionally engineers the whole thing. Diana 'makes Betty do it', deciding what she wants Betty to do and engineering both Betty's intrinsic nature and her environment from before her birth. Whatever Betty decides, Diana has already made the decision for her. Betty's choice, in addition to being determined, is *dependent* on Diana's will, and the compatibilist may plausibly think that free actions must not be dependent in this way. Call this view — that free actions can be determined but must not be dependent on another's will — *independent compatibilism*.

Of course, even determined agents can be influenced by others to *some* degree, so independent compatibilists will need to think hard about what sorts of fiddling make actions 'dependent' (and hence unfree) and what sorts don't. But, in response to the Zygote Argument, independent compatibilists will at a minimum think that if one agent engineers things so as to get another to *A*, and if the first agent knows that her tinkering will result in the other's being causally determined to *A*, the resulting *A*-ing isn't free. That is, he will endorse

Independence: If *S*'s arranging matters in way *w* would result in *T*'s being causally determined to *A*, and if *S* knows this and arranges matters in way *w* in order to get *T* to *A*, then *T* does not freely *A*.

This is completely consistent with thinking that in a wide variety of other cases determined agents act freely.

¹¹Historical compatibilism has been developed and defended most prominently by Mele (*Autonomous Agents*, Oxford University Press, 1995), who holds that if an agent *S* freely *As*, then (*very* roughly) *A* didn't come about in a way that *S* wouldn't have wanted it to (see *Autonomous Agents*, p. 193 for the official proposal). Mele's official proposal is consistent with (1) — and he argues in *Free Will and Luck* for a hard-line response to the argument — but compatibilists may be historical while disagreeing with him on these details.

4 A COMPATIBILIST FREE WILL DEFENSE

Independent compatibilists can use the free-will defense. Here, in broad brushstrokes, is the idea. Suppose God decides to create a deterministic world with free creatures. Since the world is going to be deterministic, he can create a world simply by actualizing some laws and initial conditions. We can imagine God considering all the laws and initial conditions to see what kind of world he would get if he chose *that* laws-and-initial-condition pairing, and choosing a pairing on the basis of what he sees.

Suppose God looks at laws L and initial conditions I , and notices that, were he to actualize L and I , he would get a world where everyone acted rightly all the time. Still, given independent compatibilism, if he wants morally responsible creatures he can't appeal to the fact that an L -and- I world is one where everyone acts rightly all the time as a reason to actualize L and I . If he did, he would be arranging matters that would result in each individual being causally determined to act rightly, and he would knowingly be doing in this order to get them to act rightly, and so they would not freely act rightly.

This isn't to say that God couldn't create L and I . Perhaps he can. But if he can, it will have to be for some reason *other* than that, if he did, people would act rightly all the time. For all we know, though, God has no such reason, in which case if God wants free creatures he'll have to allow some evil to slip in through the cracks.

That, at any rate, is the basic idea. It will be refined a bit in answer to a complaint.

4.1 *How to Choose*

The complaint: Given independent compatibilism, if God decides to create a deterministic world with free creatures he will have no way to provide for their well-being. After all, once God decides to create (determined) free creatures, he can't use the moral valence of these creatures' actions to decide what kind of world to make. But then, given the sorts of considerations he *can* appeal to, for all we've said the world God would have most reason to create would be one in which people torture babies for sport all the time. Balancing God's providential control with human freedom is notoriously difficult, but surely we want a God who can do more than *this*.

The complaint overstates the case, though. To see why, I'll sketch an Independence-respecting model of divine providence according to which God exercises (probabilistic) providential control while creating determined free creatures. I make no claims as to the model's accuracy. Since we're in

the business of providing a compatibilist defense — of assuming (independent) compatibilism and showing the logical compatibility of an omniscient, omnipotent, omni-benevolent God with evil — the model's intelligibility is all we need.

First, some background. Notice that God may not need to pick a *particular* set of laws and initial conditions in order to create a world. Peter van Inwagen has argued that God could create by disjunctive decree: he could say 'Let there be a cat or mouse', and if he did, then either a cat would come into existence or a mouse would come into existence — although there's no fact about which one it would be.¹² Likewise, God could decide to create a deterministic world with certain features, and decree that there be such, leaving the other chips to fall where they may.

If God can do this, then he can presumably also give essentially *probabilistic* decrees: he can say 'Let there be a cat with likelihood .75 or a mouse with likelihood .25,' at which point either a cat or a mouse will pop into existence, and the objective chance of the decree's giving rise to one or the other is .75 and .25, respectively. If this is right, then God could make general 'chancy' decrees, and could do so because these decrees would *increase the likelihood* of people acting rightly — and, more generally, would increase the expected overall goodness of the world to be created — without doing anything that he knows would bring about everyone acting rightly all the time.

Now for the model. According to the Bohmian interpretation of quantum mechanics, the physical world consists ultimately of two separate components: (i) a wavefunction, and (ii) positions of point-sized particles. The wavefunction can be represented by a function from possible configurations of particles to complex numbers. It evolves deterministically, and what it's like at later times is determined entirely by what it's like at earlier times. The positions of the particles do not affect the wavefunction or its evolution. But the wavefunction affects the positions of the particles: a particle's position at a given time, plus how the wavefunction is at that time, tells the particles how to move. The overall system is deterministic: the state of the wavefunction at a given time, combined with the precise configuration of particles at that time, entails (via the laws) the state of the wavefunction and precise configuration of particles at any later time.

Any wavefunction is compatible with any configuration of particles.¹³ Knowledge of the exact location of the particles plus the state of the wave-

¹²"The Place of Chance in a World Sustained by God," in T. V. Morris (ed.), *Divine and Human Action* (Ithaca: Cornell University Press, 1988), pp. 211–235.

¹³Well, more or less, keeping in mind that the wavefunction's dimensionality must mesh with the number of particles. Note also that the theory is non-relativistic, so we're here imagining God creating non-relativistic worlds. See Tim Maudlin, *Quantum Non-Localities and Relativity* (Wiley-Blackwell, 2011, third edition) for discussion of the relevant issues.

function at a time would allow a Laplacian demon to calculate the exact location of the particles and state of the wavefunction at any later time.

(We can't know both the exact location of the particles and the state of the wavefunction at a given time. But, on the Bohmian picture, that's not thanks to any deep metaphysical spookiness but just because any attempt to measure the exact location of particles corresponds to a shift in the wavefunction. An analysis of what we'd have to do to measure particle positions and wavefunction-states shows us that, when we find out where the particles are we lose information about how the wavefunction is, and when we find out how the wavefunction is we lose information about where the particles are. We can't simultaneously track *both* the positions of particles and the state of the wavefunction. That's no problem for God or a Laplacian demon, who can know these locations 'directly', independent of measurement.)

Suppose a Laplacian demon knows the state of the wavefunction at a time t but, rather than knowing the precise configurations of the particles at t , only knows the probability of each possible configuration. This demon won't be able to figure out which configurations the particles will have at each later time, but he *will* be able to figure out the probabilities of each configuration at each later time.

What goes for the Laplacian demon goes for God. If God knows what the wavefunction is initially like, and if God knows how likely each configuration of particles is, he will know how likely each subsequent configuration of particles is, too. And so if God decrees 'Let there be a Bohmian world with thus-and-so initial wavefunction, and with probabilities of initial particle distributions being like so,' he will know how likely each possible future will be given that decree.

(What probability distribution should he use? In order to use the laws to figure out the later probabilities, the initial probabilities need to follow the 'Born rule': for a given initial wavefunction Ψ , the probability of each configuration c should be $|\Psi(c)|^2$. Such a decree would explain why quantum mechanics' probabilistic predictions pan out so well.¹⁴ Let a decree of a wavefunction Ψ and some initial configuration c or other with probability $|\Psi(c)|^2$ be called Ψ 's *characteristic decree*.)

Since individuals' actions supervenes on the configurations of particles,

¹⁴See e.g. Craig Callendar, "The Emergence and Interpretation of Probability in Bohmian Mechanics," *Studies in History and Philosophy of Science B* 38 (2007): 351–370 for a discussion of how this works and related issues. Callendar also there compares the Bohmian picture of probabilities with that of statistical mechanics, another deterministic physical theory. Note that we could adapt the present model of divine providence for use with statistical mechanics rather than Bohmian mechanics by replacing wavefunctions with thermodynamic states and the Born-rule based decree with one based on David Albert's past hypothesis (*Time and Chance*, Harvard University Press, 2003).

God can use this trick to exercise some providential control while issuing a decree which leaves open his determined creatures' choices. God can examine each initial wavefunction and see how likely it is that his creatures choose the right were he to issue the characteristic decree for that wavefunction. He can use this information to rank wavefunctions as to how likely it is that he'll get a good world from such a decree, how likely he'll get a bad one, and so on, to determine which wavefunction's characteristic decree would best maximize the expected goodness of his creation. And he can then issue the characteristic decree that maximizes expected goodness. Nowhere in that process does he arrange matters in such a way as to determine that everyone acts rightly in order to get them to do so, so Independence presents no bar to such a process getting him free creatures.

4.2 *The Atheist Objects*

The atheist objected to the original free will defense on the grounds that God could know in advance how people would (freely) act in a given circumstance and take their free will away just when they were about to act wrongly. The atheist has a similar objection to the compatibilist free will defense spelled out here. Sticking with the Bohmian model of divine creation outlined above, the objection comes in two forms.

In the first form, the atheist says, "I'll grant that, if God decrees 'Let A occur with probability p or B occur with probability q ', then either A occurs with probability p or B occurs with probability q . But furthermore, God knows *which* one will occur if he issues such a decree. Now, at least many wavefunctions will have non-zero probability of everyone acting rightly all the time. Why doesn't he pick one, Ψ , where he knows that, *if* he issues Ψ 's characteristic decree, the world he'll get is the one where everyone *in fact* acts rightly all the time?"

We might respond to the atheist as follows: "If God were to do that, then he would be issuing a decree because he knows it will get him people determined to act rightly all the time, and he'll be doing that in order to get people to act rightly all the time, and so by Independence those people won't be morally responsible for acting rightly all the time."

This response seems right to me, but the objection faces a bigger problem. The atheist's objection to the original free will defense presupposed Molinism — the view that for every possible circumstance C , agent S , and action A , exactly one of "If S were put in circumstance C , he would A " and "If S were put in circumstance C , he would not A " is true, God knows which it is, and God can use this knowledge in deciding what to do. His objection now presupposes another Molinist-esque thesis: the view that, for every possible

probabilistic decree D and possible outcome O of D , exactly one of “If God decreed D , O would occur” and “If God decreed D , O would not occur” is true, God knows which it is, and God can use this knowledge in deciding what to do.

If the Molinist-esque presupposition is false, then the atheist’s proposal makes no sense. But even if it is true, we have no reason to suppose that *any* wavefunction’s characteristic decree would result in a world where people act rightly all the time. It may be that, through no fault of God’s, no conditional of the form “If God were to issue Ψ ’s characteristic decree, everyone would act rightly all the time” is true.

So much for the first form of the objection. But the atheist has a second, smaller-scale objection. “Suppose God issues some wavefunction’s characteristic decree,” he says. “Then, as the world unfolds, he sees that Herman is about to perform a heinous crime. Why can’t God intervene and take Herman’s free will away just before the crime is committed? And why can’t God repeat this process, leaving people’s free will intact whenever they’re about to act rightly, but removing it when they’re about to act wrongly?”

Two initial questions before responding to the theist. First, what would it be on the Bohmian model for Herman to be about to commit a heinous crime? Presumably it would be for the particle configuration, as driven by the wavefunction, to be about to evolve into one that constitutes Herman committing the crime.

Second, what would it be for God to intervene and take away Herman’s free will with respect to this act? Presumably, it would be for God to modify the world in so as to lead to Herman’s not committing the crime, while leaving as much else intact. The most natural way to do this would be to leave the wavefunction alone and change the particle configurations just a little bit, so as to shift them from a configuration that the wavefunction is driving towards a crime-committing Herman to one that the wavefunction is driving towards a non-crime-committing Herman. (It might be enough, for instance, to move one particle in the configuration just a little to the left.) If he does this, of course, thanks to Independence Herman will not be free with respect to not committing the crime; but that was in the cards anyway, since the atheist is having God take away Herman’s freedom for the sake of crime-avoidance.

There will generally be multiple ways God could change the particle configurations to keep Herman from committing the crime. God could, in principle, evaluate each such way to determine the one with the best overall outcome, not just for Herman’s behavior but for the subsequent world as a whole. But if he chooses some particular configuration because it would lead to *others* acting rightly down the road, then that choice would take away the freedom of those others. Since he’s trying to surgically remove Herman’s

freedom while leaving others' intact, God's best bet is to again look at the possibilities, make some expected overall goodness calculations, and say 'Let the particle configuration switch to one of *those*, with thus-and-so probabilities.' Call such a decree a *switching decree*.

God can indeed make switching decrees, and so can indeed surgically remove Herman's free will shortly before he would otherwise commit the crime. But in order to undermine the compatibilist's free will defense, the atheist has to think that God can make switching decrees in a way that will lead to everyone acting rightly all the time. And we have no reason to think that God can do that — at least, not without taking away *everyone's* free will.

Suppose that, to keep Herman from his crime, God issues a switching decree *D*. It may be that the result of *D* is to move the particles in such a way that someone else, who was previously on track to act rightly, is now lined up to do wrongly. Of course, God can intervene again; but such an intervention may lead to yet someone else doing wrongly, and so on and so forth. In fact, the Bohmian God may be in a situation precisely analogous to the Molinist God the theist described in the first section: God may simply be dealt conditionals about what happens given each possible probabilistic decree on which there is simply no pattern of decree-issuing that leads to a world with (some) free creatures where everyone acts rightly all the time.

5 CONCLUSION

Despite appearances, merely being a compatibilist does not put the free will defense beyond reach. (For that matter, merely being a fan of the free will defense does not put compatibilism beyond reach, either.) I have outlined one sort of strategy for a compatibilist free-will defense, and for all I know there may be others.

The strategy outlined here is not for everyone. Those independently convinced of incompatibilism, or of certain compatibilist views inconsistent with Independence, or of certain pictures of divine providence, may find little to their liking. Those with other commitments may be more attracted to the defense. And those still developing their commitments can at least do so knowing that a compatibilist free will defense is on the table.¹⁵

¹⁵Thanks to Ross Cameron, Peter French, Robbie Williams, some anonymous referees, audiences at the University of Leeds, Calvin College, and the 2006 Inland Northwest Philosophy Conference, and most especially Dean Zimmerman for helpful comments on and conversations about this paper and various predecessors. Research supported in part by the British Academy, grant SG-53931.