THEOLOGICAL DETERMINISM

New Perspectives

EDITED BY

PETER FURLONG

Valencia College

LEIGH VICENS

Augustana University

CAMBRIDGE UNIVERSITY PRESS

CHAPTER 14

Toward a Reactive Attitudes Theodicy

Garrett Pendergraft

14.1 The Problem of Evil for the Theological Determinist

The *problem of evil* typically refers to the problem (or set of problems) that theists encounter when trying to reconcile the existence of God, as traditionally conceived (omnipotent, omniscient, and morally perfect), with the existence of evil. I will understand the problem of evil in a more specialized sense, as the problem of figuring out how to respond to the *argument from evil* against theism. As is often the case with important and influential arguments such as the argument from evil, the label refers to a family of related arguments rather than a single canonical argument formulation. I will focus on a relatively well-known member of the family, namely the argument from gratuitous evil.[1] The argument from gratuitous evil contains two premises. The first premise (call it the *theological premise*) is a claim about the type of world that God (as traditionally conceived) would create; the second premise (call it the *empirical premise*) is a claim about the type of world that actually exists; and the conclusion is that God doesn't exist.

(1) If God were to exist, then gratuitous evil wouldn't exist.
(2) But gratuitous evil does exist.
(3) Therefore, God doesn't exist.

Responding to the argument from gratuitous evil will involve drawing on the resources found in one's doctrinal or theoretical commitments and availing oneself of new resources (and thus taking on board new commitments) as needed. And any new commitments, of course,

[1] This argument is a simplified version of Rowe's argument in Rowe (1979). This particular formulation is presented and discussed in, among other places, Judisch (2012), DePoe (2014), and Mooney (2019).

https://doi.org/10.1017/9781009249362.015 Published online by Cambridge University Press

must be consistent with (or at least require minimal revision of) existing commitments.

One implication of this way of looking at things is that we can evaluate different views of divine providence with respect to the resources they have (or have available to them) when encountering the problem of evil. By these lights, theological determinism is often seen as especially problematic: the determinist is seen as having an impoverished set of resources to draw from in her attempts to respond to the argument from gratuitous evil.[2]

The doctrine of theological determinism is relatively easy to understand at an intuitive level, but difficult to specify in a comprehensive way. We can, however, identify and try to articulate various consequences of a more comprehensive specification. For example, White (2019) has offered a helpful characterization of theological determinism as involving a maximally robust explanatory relation between facts about God's will and contingent facts about the world and its creatures. (White assumes, as do I, that at least some of the facts about God's will are contingent. Thus, facts about God's will cannot explain *every* contingent fact, but at most every *other* contingent fact.) More precisely (and following White), I will take theological determinism to imply an entailment claim and an explanation claim:

(4) Facts about God's will entail every other contingent fact.
(5) Facts about God's will are explanatorily prior to every other contingent fact.[3]

In short, theological determinism implies that *God's will completely explains every other contingent fact*, with the understanding that a complete explanation will be such that facts in the explanans entail the facts in the explanandum.[4]

Our question, then, is whether a view about providence that affirms (4) and (5) suffers from a lack of available resources when engaging with the problem of evil. As we will see below, there is a particularly important

[2] In what follows, I will be using "determinism" and "determinist," by themselves, to refer to theological determinism and the theological determinist rather than causal determinism and the causal determinist.

[3] These propositions are taken verbatim from White (2019: 6), but I have renumbered them.

[4] This understanding of theological determinism is prefigured in McCann (2001: 112), who holds "that we and all of our actions depend upon God for their existence, not vice versa, so that the perfect good for which he creates the world counts as the full and final explanation for what we do."

resource – an appeal to free will – that many have claimed is unavailable to the determinist. I will be arguing that the determinist does not in fact suffer from this particular resource deficit. I will do so by examining some recent work on determinist responses to the problem of evil, analyzing some of the ways in which they fall short, and proposing a new response that is informed by those shortcomings. As we will see, this new response is built on an appeal to the reactive attitudes.

I should acknowledge, however, that my focus will be relatively narrow. Because I will be focusing on deterministic responses to the problem of evil, I won't be addressing general objections to theological determinism as a model of divine providence. (Some of these issues are admittedly closely related to the issues a determinist faces when responding to the problem of evil.) For example, a common (and important) objection to theological determinism is that it makes God in some sense responsible for, or at least approving of, human wrongdoing; in common parlance, it makes God the *author of sin*. If this objection goes through, then theological determinism is fatally flawed and the question of whether the determinist can refute the argument from gratuitous evil loses its interest. (Given my interest in this question, one can infer that I'm not yet convinced that the objection goes through.) For this reason, I think it's better to view a response to the authoring sin objection (and related objections) as part of the broader project of defending theological determinism as a whole, rather than the narrower project – our project here – of examining deterministic responses to the problem of evil.[5]

14.2 Responding to the Argument from Gratuitous Evil

Recall that the argument from gratuitous evil contains (1) a theological premise and (2) an empirical premise:

(1) If God were to exist, then gratuitous evil wouldn't exist.
(2) Gratuitous evil exists.

The most popular responses to this argument attempt to cast doubt on the empirical premise by claiming that the existence of some set of sufficiently valuable goods entails, and thus justifies, the existence of some

[5] For recent examples of the broader project (although in both cases the commitment to theological determinism is tentative), see Furlong (2019) and White (2019). For a recent example of the narrower project, see Byerly (2017).

set of evils.[6] (More precisely, some proposition involving the goods will entail some proposition about the evils.) These *greater-good strategies* have been implemented in various ways, but they typically involve arguing that some apparently gratuitous evils are not in fact gratuitous because they are constituents, prerequisites, or necessary consequences of some goods that are valuable enough to justify permitting those evils.[7]

Speak (2013: 206) offers a helpful way of analyzing greater-good strategies, pointing out that they will typically endorse both an *impossibility proposition* and a *value proposition.* The impossibility proposition purports to identify the goods that are impossible to bring about without also bringing about (or at least allowing) some evils, and the value proposition will claim that a state of affairs including both the goods and the evils is more valuable than a state of affairs including neither the goods nor the evils. (If this type of value proposition is true for some goods, then we will say those goods are *valuable enough* to justify the evils.) If these two propositions are true, for some set of goods and evils, then it follows that God has a sufficient reason to create a world that includes those evils, which is to say that they are not gratuitous.

A complete execution of the greater-good strategy would require developing a comprehensive theodicy in which every instance (or at least every type) of apparently gratuitous evil is shown to be entailed by some set of sufficiently valuable goods. Most theodicies, however, focus on a limited set of evils and attempt to show how those evils are justified by some set of goods.

In addition to (or in conjunction with) these greater-good strategies, the theist attempting to respond to the argument from gratuitous evil could cast doubt on the empirical premise in a more literal way by adopting some version of skeptical theism.[8] According to skeptical theism, we should be reluctant to affirm the existence of gratuitous evils because we should be reluctant to think that we have a handle on all of the possible reasons that God might have for bringing about evil. Confidently asserting that a particular evil or set of evils is gratuitous requires confidence

[6] There are some responses that focus instead on the theological premise. For example, Hasker (1992) denies the theological premise on the basis of open theist convictions, and Mooney (2019) denies it on the basis of deontological considerations.

[7] Pittard (2018: 91) makes a further distinction among greater-good strategies, namely between those requiring that every actual evil contribute somehow to the greater good (what he calls "defeat theodicies") and those merely requiring the *possibility* of evil in order to secure the possibility of the greater good.

[8] For a helpful introduction to skeptical theism, see McBrayer (2010).

in the claim that God could have no reason for permitting those evils, and such confidence is simply not warranted for limited beings such as ourselves. Evaluating the merits of skeptical theism is beyond the scope of this essay, but it's worth noting as a potential resource for the indeterminist and determinist alike.[9]

14.3 Characterizing the Project

Given this way of structuring the dialectic (including a focus on greater-good strategies), the question the theological determinist faces is whether she has sufficient resources for mounting an attack on the empirical claim that gratuitous evil exists. An exhaustive treatment of this question would involve an examination of every available response to the argument from evil, to see whether it generates sufficient doubt about the existence of gratuitous evil and also whether it can be adopted by the determinist. This exhaustive treatment is a task that we obviously won't be able to undertake here. But in lieu of an exhaustive treatment, we can make some progress by focusing on representative instances of the most popular responses to the argument from evil. As it turns out, the most popular responses involve some sort of appeal to free will. And since free will is one of the goods that the determinist arguably has the most difficult time incorporating into her model of divine providence, focusing on freedom should provide a good test of the determinist's prospects for responding to the argument from gratuitous evil.[10] For these reasons, we will focus our efforts on developing a determinist-friendly free will theodicy.[11]

Before we proceed with these efforts, however, it might be helpful to say a little bit more to clarify how I am understanding the project. To offer a theodicy, as we saw above, is to claim that there exist some goods that would be impossible to have without the existence of some evils, and that are valuable enough to justify the evils. But the proponent of a free will

[9] Pereboom (2005: 88ff.) argues that the determinist, relative to the libertarian, is going to have an easier time dealing with some of the main challenges to skeptical theism.

[10] If the appeal to freedom is unsuccessful, there are other goods that the determinist could appeal to. Even a critic of determinism as staunch as O'Connor acknowledges (2016: 135) that it is "plausible to assert that … there are highly valuable goods not involving or entailing creaturely freedom that a perfect Creator would rightly not wish to forego and that are not attainable without permitting the existence of suffering."

[11] Some examples of deterministic theodicies that don't appeal directly to freedom-relevant goods are discussed in Trakakis (2006), Green (2016), Pereboom (2016), and White (2019: 253–259). Trakakis and Pereboom are revising and developing extant theodicies, some of which are more closely tied to freedom goods than others.

theodicy doesn't (or at least shouldn't) claim that freedom *in and of itself* is valuable enough to justify the existence of evil.[12] This is easy to see if we understand free will as entailing the ability to do otherwise – or, more precisely, the ability at some time to either perform an action or refrain from performing that action.

Greater-good strategies require both an impossibility proposition and a value proposition, and the mere ability to do otherwise can support neither. It is possible for this ability to exist without evil, for example, in a world in which the only types of free actions performed involve morally neutral options. Moreover, this ability, considered in isolation, doesn't seem valuable enough to justify the existence of evil. What makes free will so valuable is not the mere ability to choose between options, but to choose between good and bad options; that is, the ability to make choices that have moral significance, and for which one is morally responsible.[13] For this reason, free will theodicies typically appeal not just to freedom, but to *significant freedom*, defined as the ability to freely choose between a variety of options ranging from very good to very bad (cf. Mooney 2019: 447). Thus, the impossibility and value propositions that partially constitute a free will theodicy will not appeal to freedom in isolation, but to something (such as significant freedom) that requires but goes beyond a mere ability to do otherwise.

So a free will theodicy, despite its name, appeals not merely to free will but to something for which free will makes one eligible.[14] Significant freedom is one of the things for which free will makes one eligible, and that is what most free will theodicies appeal to.

The second clarifying note involves the familiar distinction between two different types of response to the problem of evil: a *theodicy* and a *defense*. According to what seems to be the official understanding of this distinction (cf. van Inwagen 2006: 7), the difference between a theodicy and a defense is not a difference in content but a difference in degree of endorsement: the proponent of a theodicy is actually endorsing the content of the

[12] Ekstrom (2016: 77) makes this point in a particularly compelling way, and in fact she wonders whether even morally significant freedom would be enough to justify all the evil that exists.

[13] The claim that a free will valuable enough to justify evil needs to be exercised in a context of moral significance is endorsed, implicitly or explicitly, by both Plantinga (1974a: 30) and Swinburne (1998: 90). For additional discussion, see Speak (2013) and Pereboom (2016).

[14] Bishop (1993: 118) appeals to loving personal relationships as the relevant good for which free will makes one eligible: "For, arguably, loving personal relationships could not possess their highest value, unless those who form and maintain them did so against the background of a robust capacity to do themselves and each other really serious harm."

theodicy, whereas the proponent of a defense is merely claiming that (the truth of) the content of the defense is epistemically possible. On this official understanding of the distinction, someone's theodicy could be identical in content to someone else's defense; the only difference would be the extent to which they are prepared to endorse that content.

Actual usage, however, doesn't always follow this official prescription. For example, some of the philosophers who have done recent work on developing so-called deterministic theodicies (e.g., Trakakis [2006], Byerly [2017]) have done so without explicitly endorsing determinism. Perhaps those who use "theodicy" this way are simply doing it wrong; but on the other hand maybe these instances of unofficial usage should prompt us to reconsider the official prescription. According to the official prescription, there's no way to take a set of propositions in the abstract and accurately describe it as either a defense or a theodicy, because first we have to know which subject is entertaining the propositions and to what extent they endorse those propositions. But it seems that we ought to be able to consider a response to the problem of evil and refer to it as a theodicy (or a defense) even without knowing who is entertaining or advocating for that response.

So I think that the standard defense/theodicy distinction is unhelpful. I will use the term "theodicy" in what follows, but in an unofficial way that doesn't imply full endorsement. Degree of endorsement is important, of course; but I propose that we talk about endorsement directly, rather than trying to represent degree of endorsement using a distinction between theodicies and defenses.

One way of talking about endorsement directly is to situate a given response to the problem of evil on a continuum, running from minimal support to maximal support. Speak (2017) has proposed that we think of responses situated toward the minimal end of the endorsement continuum as *divested*, whereas responses situated toward the maximal end are *invested*. A divested response to the problem of evil aspires only to epistemic possibility and carries very little endorsement, but an invested response comes with a significant degree of endorsement. (There will, of course, be borderline cases in which it's not clear whether the degree of endorsement counts as divestment or investment.) Other things being equal, it seems preferable that someone's response to the problem of evil be invested rather than divested.

So we are interested in the degree to which someone is invested in a theodicy; but we are also interested in whether or not we can *persuade* someone to invest in a theodicy. There's no easy way to measure the persuasive

power of a particular theodicy, but I think we can make some headway by construing different responses to the problem of evil according to their relative *demandingness*. Suppose that we are examining two competing responses to the problem of evil, which are identical except that one of them includes or entails a controversial theoretical commitment (whereas the other does not include or entail that commitment). We can say that the response including the theoretical commitment is more demanding, relative to the response that does not include that commitment.

Consider, for example, the notion of a *counterfactual of creaturely freedom*, which has to do with what some creature would freely do in a given set of circumstances. (Molinists have famously appealed to such counterfactuals, which are typically represented by something like the following schema: If S were to be in circumstances C, then S would freely perform A.) Many philosophers of religion have seen fit to endorse a commitment to counterfactuals of freedom, but such a commitment nevertheless remains controversial. Other things being equal, a response to the problem of evil that appeals to counterfactuals of freedom is more demanding than a response that doesn't make such an appeal. This, however, doesn't mean that every theodicy that appeals to counterfactuals of freedom is at a disadvantage. A rival theodicy that doesn't appeal to such counterfactuals might instead appeal to some other controversial commitment. In this case there won't be any clear difference in demandingness, and a comparison of the two theodicies will need to take into account more than just demandingness. (In such a case we might want to look at explanatory power, or some other theoretical virtue, in addition to demandingness.) A judgment of relative demandingness, in other words, will rarely be the deciding factor when we are evaluating theodicies. But in the interest of good methodological hygiene it does seem that, other things being equal, we should prefer less demanding responses to the problem of evil.[15]

14.4 In Search of a Free Will Theodicy

We have seen that most free will theodicies will appeal to the existence and importance of significant freedom, that is, the freedom to choose between good and bad. Thus they will have something like the following structure (including an existence proposition in addition to the impossibility proposition and the value proposition):

[15] Several philosophers have made something like this methodological point with respect to theories of free will: see, for example, Graham and Horgan (1994), Vargas (2004), and Nahmias et al. (2006).

(6) Significant freedom exists.
(7) Significant freedom would not be possible without evil.
(8) Significant freedom is valuable enough to justify the existence of evil.

So why can't the theological determinist just appeal directly to a significant freedom theodicy involving (6)–(8)? The answer is that such an appeal would take on the burden of arguing that significant freedom is compatible with theological determinism. This is not necessarily an impossible task, but it does involve a controversial theoretical commitment.[16] So a theodicy that needs to argue for (or worse, presuppose) the compatibility of significant freedom and theological determinism is going to be more demanding than a theodicy that doesn't share that argumentative burden.[17] Since we are looking for a minimally demanding theodicy, we should look elsewhere.

14.4.1 Building a Theodicy on an Independence Condition

Perhaps the most straightforward way to build the kind of theodicy we are looking for is to propose an account of free will that includes some condition, satisfaction of which is uncontroversially compatible with determinism but which also precludes God from determining that free creatures always choose the good.

Turner (2013) has offered one of the more widely discussed versions of such an account. This account is built on an anti-manipulation condition that he refers to as an *independence* condition:

(9) If T's arranging matters in way *y* would result in S's being determined to A, and if T knows this and arranges matters in way *y* in order to get S to A, then S does not freely A.[18]

Given an account of free will that adopts the independence condition, it is relatively easy to see why free will would be impossible without evil. If God were to determine that everyone always chooses the good, then he would be arranging things so as to prevent evil, and actions performed on

[16] Michael Almeida argues that the determinist can appeal to a significant freedom theodicy in Almeida (2017).
[17] For a representative example of theological compatibilism, see Helm (2010); for an example of theological incompatibilism, see Pereboom (2016). For a treatment of deterministic theodicies from the perspective of the theological incompatibilist, see Pereboom (2011).
[18] This condition is taken almost verbatim from Turner (2013: 131), except that I have made some notational changes.

those terms would not be performed freely.[19] Thus, on the independence account being considered here, if God wanted to create a world in which free will existed then he would be precluded from determining that everyone always chooses the good. In other words, (9) generates the right kind of impossibility proposition:

(10) Independence from manipulation would not be possible without evil.

What about the value proposition? The straightforward way of articulating it would be as follows:

(11) Independence from manipulation is valuable enough to justify the existence of evil.

On its face, (11) is not very plausible. Independence from manipulation is valuable, but that value is parasitic on the value of other goods that are connected with freedom in a deeper way. Considered in isolation, mere independence from manipulation just doesn't seem valuable enough to justify the existence of evil.

To abandon a Turner-style response for this reason, however, would be too quick. As we saw above, free will – at least when considered in isolation – isn't valuable enough to justify evil either. What's needed for a free will theodicy is something like *morally significant* freedom. Freedom is useful for theodical purposes only insofar as it is necessary for morally significant freedom. But the same move could be made with respect to independence from manipulation: independence is useful for theodical purposes only insofar as it is necessary for freedom, which is in turn necessary for morally significant freedom. The resulting theodicy would be built on (6), (8), and (10). The problem with this maneuver is that it inherits the controversial commitment that we are trying to avoid, namely the commitment to the compatibility of significant freedom and theological determinism.

It would seem, then, that someone who wants to build a theodicy on an independence condition faces a dilemma: they will either be stuck with an implausible value proposition, or they will be stuck with a theodicy that is just as demanding as a theodicy that appeals directly to significant freedom.[20]

[19] Although I will sometimes use masculine pronouns to refer to God, I am not intending to imply that God has a gender or that masculine terms are more revealing of God's nature than feminine terms.

[20] Furlong (2019: 150–152) offers some additional reasons why someone should be reluctant to invest in a theodicy built on the independence condition.

Even if the independence theodicist can overcome this dilemma, however, she will face another problem, which involves God's providential concern for his creatures. Here's how Turner articulates the problem:

> 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*. (2013: 132)

Turner suggests a way out of this problem, but it requires abandoning deterministic divine decrees in favor of probabilistic ones, and thus will be of no help to the divine determinist. It seems, then, that the determinist who relies upon something like (9) will have a hard time explaining how God can exercise providential care over human beings.[21]

So there are two problems with Turner's proposal. First, it inherits the compatibilist's argumentative burden; second, it does not appear to be the sort of proposal that a determinist would want to invest in. The independence condition is plausible, but it doesn't seem to fit the description of what we're looking for.

14.4.2 Building a Theodicy on Counterfactuals of Freedom

One strategy – a surprising strategy – for avoiding the compatibilist's argumentative burden is to appeal to freedom in a way that libertarians (or at least some libertarians) already endorse. An example of this strategy can be found in Howsepian (2007), who appeals to what he calls *middle freedom.*

Consider a world in which agents act freely and that is not deterministic; call it a *libertarian world.*[22] Howsepian proposes (and endorses) a theory built on the *principle of middle freedom*, which says, roughly speaking, that S performs A freely in a world *w* if and only if S would perform A as an exercise of agent-causal free will in the libertarian world nearest

[21] Thanks to Peter Furlong for help with this point.

[22] The label is slightly misleading, since such worlds could exist even if compatibilism is true. Libertarians think that every world with freedom is what we are calling a libertarian world.

to *w*.[23] On this view, the specification of each person's individual essence includes a set of counterfactuals of freedom, and their freedom in action depends, not on the actions themselves, but on the relation between the actions and the relevant counterfactuals.[24] On Howsepian's view, even if I am determined to perform some action, that action can still be free as long as the circumstances in which I perform the action are similar enough to a possible set of indeterministic circumstances in which I exercise agent-causal freedom in performing that same action.

Keeping in mind our earlier caveats about the importance of morally significant freedom, Howsepian's account generates the following impossibility proposition and value proposition:

(12) Significant middle freedom would not be possible without evil.
(13) Significant middle freedom is valuable enough to justify the existence of evil.

Both of these propositions seem relatively plausible. Suppose that God wants to actualize a deterministic world *w* in which, among other occurrences, S exercises free will in C. If S's essence is constituted in such a way that S would (exercise agent-causal powers and thus) freely perform some bad action B in C′ (where C′ is the set of indeterministic circumstances most similar to C), then *w* will be a world in which evil occurs as a result of S's exercise of middle freedom. More generally, a world in which agents exercise middle freedom will be a world in which evil occurs. So (12) seems true. And (13) also seems true, since significant middle freedom is presumably just as valuable as significant freedom *simpliciter*.

Unfortunately for our purposes, Howsepian's view is not going to be the minimally demanding, determinist-friendly, free will theodicy that we are looking for. First, it is relatively demanding, in virtue of its commitment to the existence of counterfactuals of freedom, both the existence and nature of which are controversial. A related and more important problem for Howsepian's view is that free agents are supposed to be in control of whether an action of theirs is praiseworthy or blameworthy. Whether an agent's action is praiseworthy or blameworthy depends on

[23] Howsepian (2007: 222) articulates the principle of middle freedom as governing determined actions, but he also claims that his theory can be extended to cover indeterministic theories of free action as well.

[24] Howsepian describes the essence of his view as follows: "free human actions do not depend upon those actions' mechanisms of implementation, but … on the (non-causal) relation between one's actions and one's counterfactuals of creaturely freedom" (2007: 223).

the relevant counterfactuals of freedom, which in turn depend on contingent properties of the agent's essence. Thus, it would seem that free agents need to have control over contingent properties of their essence (Howsepian 2007: 225–226). Unfortunately, it's not at all clear how anyone could have control over such properties.[25] Thus, an appeal to middle freedom also takes on the substantial burden of showing how we can have this type of control.

Second, Howsepian's view is not determinist-friendly. Counterfactuals of freedom are not by themselves inconsistent with determinism, since it's possible that the counterfactuals are completely explained by facts about God's will. But Howsepian (2017: 225) explicitly says that the "valence" of an action (i.e., whether it's free or unfree) is not up to God. On his view, God determines what people do but not whether they do it freely; whether they do it freely depends on contingent features of their creaturely essence. It's hard to see how this assertion can be squared with the determinist view that God's will completely explains every contingent fact about the created order.[26]

14.4.3 Taking Stock

Both Turner's proposal and Howsepian's view fall short of what we are looking for, but they do so in complementary (and instructive) ways. Turner's view is built on an admirably austere foundation, but as a consequence picks up an additional argumentative burden in virtue of having to borrow a value proposition. Howsepian's view, on the other hand, is too demanding. (Also, it turns out that neither view is going to be very attractive to the determinist.) In light of this situation, some have concluded that the overall prospects for a freedom-relevant, determinist-friendly theodicy are bleak (e.g., Furlong 2019: 155). I think, however, there is reason for optimism. I think we can make at least a little bit of progress toward identifying a free will theodicy that the determinist can endorse without taking on board too many controversial commitments.

[25] Furlong (2019: 141–146) explains in detail why the prospects for an account of this type of control are bleak.

[26] It would also seem that Howsepian's view is inconsistent with soft determinism (the view, roughly speaking, that free will is not only compatible with determinism but *requires* determinism). According to the principle of middle freedom, an action is free only if it's freely performed in some indeterministic world. According to soft determinism, there are no such worlds.

14.5 Toward a Reactive Attitudes Theodicy

As noted above, the main reason why we care about free will (at least in this context) is that it is required for significant freedom – the freedom to choose between good and bad – which in turn seems valuable enough to justify the existence of evil. It's worth noting, however, that the freedom to choose between good and bad is not the whole story about valuable human agency. If significant freedom is just the ability to make a free choice in a context of moral significance, then it's possible that someone could exercise significant freedom and yet not be *morally responsible* for the choice they made. (For example, if someone makes a free choice that has moral significance but isn't aware of that moral significance, then they will typically not be responsible for such a choice.)

This is not a serious problem for the notion of significant freedom, because the theodicist could simply stipulate that significant freedom also requires an awareness of the moral significance of the relevant choice. But this observation does highlight the relevance and importance of moral responsibility. Theodicies that focus on significant freedom tend to focus on the moral status of the options that are available to free agents, rather than on the moral responsibility typically possessed by such agents. I propose that we focus instead on moral responsibility itself, and that this shift in focus will open up a path toward a more promising (determinist-friendly, freedom-relevant) theodicy.

Agents who exercise significant freedom will often be morally responsible for the free choices they make. The determinist who is responding to the problem of evil can't simply appeal to moral responsibility directly, though. This is because the relevant propositions (existence, impossibility, and value), while plausible, are just as controversial (give or take a challenge or two) as the corresponding propositions involving significant freedom. A moral responsibility theodicy of this type would not represent progress because it would face the same argumentative burden as a significant freedom theodicy.[27] We need something that's uncontroversially (or at least *less* controversially) compatible with theological determinism but that can also generate a plausible impossibility proposition and a plausible value proposition. Perhaps we can make some progress by looking a little more closely at the concept of moral responsibility.

[27] As P. F. Strawson notes, incompatibilists would argue that "*just* punishment and *moral* condemnation imply moral guilt and guilt implies moral responsibility and moral responsibility implies freedom and freedom implies the falsity of determinism" (1962: 2–3).

One influential way of thinking about moral responsibility traces back to P. F. Strawson's (1962) discussion of, among other things, the *reactive attitudes*: attitudes that represent our reaction to displays of good will or ill will toward ourselves or others. Moral praise and moral blame are obvious examples, as are resentment, indignation, and guilt. (Expressions of ill will toward us tend to give rise to resentment; others' expressions of ill will toward each other tend to give rise to indignation; and our own expressions of ill will toward others tend to give rise to guilt.) Our tendency to have these reactive attitudes, and our practices of expressing them, represent (and help regulate) a system of interpersonal demands and expectations, the content of which involves the quality of will expressed in the actions of ourselves and others. There is widespread agreement about the importance of the reactive attitudes (notwithstanding significant amounts of controversy over the exact relationship between the reactive attitudes and moral responsibility), and numerous theorists have incorporated an appeal to the reactive attitudes as a core ingredient of their theory of moral responsibility.[28] We can't take on board a complete theory of moral responsibility, lest we violate our directive to find a relatively minimally demanding theodicy, but perhaps we can build on some sort of appeal to the reactive attitudes.

Unfortunately, we can't make a *direct* appeal to the reactive attitudes either, since there is plenty of controversy about whether the reactive attitudes could be justified given the truth of determinism. (For example, the incompatibilist will want to say that moral praise and blame would not be justified if determinism were true.)[29] But, I would like to argue, if we zoom in just a little bit more then we will see some resources that the determinist-friendly theodicist can make use of.

I propose that we borrow some of the elements of Fischer and Tognazzini's (2011) detailed and helpful analysis of the "physiognomy" of responsibility.[30] They start by pointing out that a judgment of moral responsibility is not a unified or monolithic activity; instead, responsibility judgments (and the practices built around those judgments) can be divided into numerous analytical or conceptual stages. The most important stage, at least for present purposes, involves what they call "openness

[28] For a helpful summary of Strawson's reactive attitudes approach and some prominent criticisms of it, see Talbert (2019, §2.2).

[29] For a criticism of Strawson's position that touches upon these incompatibilist concerns, see Pereboom (2008).

[30] Fischer and Tognazzini are analyzing the *physiognomy* of responsibility in the sense that they are examining some of the specific features of the two faces of responsibility that Watson identifies in Watson (1996).

to reactive appraisal," or *reactive attributability* (2011: 386). To ask whether someone has met the conditions of reactive attributability is to ask whether they have exercised whatever capacities render them a sensible target of the reactive attitudes.

It is important to distinguish, however, between a mere openness to reactive appraisal and an instance of reactive appraisal itself. To engage in reactive appraisal is to actually target someone with the reactive attitudes, which is one way of holding them accountable for their action(s). To see how these two conceptual stages come apart, consider someone who is a sensible target of reactive appraisal (and thus meets the conditions of reactive attributability) and performs an action that is morally wrong, but has an excuse for the wrongdoing that renders reactive appraisal unjustified. Fischer and Tognazzini provide the following example:

> Suppose that a mother is faced with an awful choice: either she can save her own child from drowning, or she can save five other children from drowning. Perhaps the children are swimming in a lake and the mother is driving a boat but cannot get to all six of the children in time. She must choose. Suppose, further, that she chooses to save her own child instead of the five other children. Such a decision is surely understandable, but even more than that, it is at least arguable that it would be unjustified, even in principle, for anyone to target this mother with any negative reactive attitudes. This will presumably be because of the enormous difficulty involved in doing the right thing. And yet – here's the important point – it's plausible to suppose saving her own child is a morally wrong action that is attributable to her in a reactive sense. (2011: 388–389)

The upshot (2011: 391) is that there are important differences between the question of whether it even makes sense to target someone with the reactive attitudes and the question of whether targeting someone with the reactive attitudes would be justified in a particular set of circumstances. One of the important differences has to do with how these two different features of moral responsibility interact with determinism. According to the incompatibilist, determinism makes it so that there are no circumstances in which engaging in reactive appraisal would be justified. But the incompatibilist doesn't (or at least shouldn't) think that determinism undermines reactive attributability. Whether someone is a sensible target of the reactive attitudes has to do with (a subset of) their intrinsic properties, and the abilities that they have in virtue of those properties[31] – not with the circumstances

[31] Vihvelin (2013: 27) distinguishes between *narrow* abilities and *wide* abilities, where our narrow abilities are the ones we have in virtue of our intrinsic properties.

in which those abilities are exercised (or thwarted). To meet the conditions of reactive attributability is to be eligible for the reactive attitudes, and the truth or falsity of determinism does nothing to affect that eligibility. (Similarly, someone might be eligible for early access to a vaccine in virtue of being a healthcare worker. They won't have access to the vaccine before it's been developed, but that lack of access is not an eligibility problem.)

Reactive attributability therefore shows initial promise as a minimally demanding potential resource for the determinist who is mounting a response to the problem of evil. The next question is whether the corresponding theodical propositions are plausible.

As we work our way toward the impossibility proposition, let's focus on resentment. Resentment is a fundamental – if not *the* fundamental – reactive attitude, which suggests that it plays an essential role in representing and regulating the system of interpersonal demands and expectations that are characteristic of moral responsibility. (Since we are using "reactive attributability" as a label for openness to reactive appraisal, let's use "resentment attributability" for openness to resentment.) The centrality of resentment, in other words, suggests the following preliminary impossibility proposition:

(14) Reactive attributability would not be possible without resentment attributability.

Someone might object to (14) by pointing to a world that contains a system of moral expectations similar to our own, but doesn't contain resentment; in such a world, the relevant expectations are never violated (or if they are, there's always a known excuse). I'm not confident that such a world is possible;[32] but even if it were possible, it seems to me that resentment attributability would still be a feature of that world. Suppose that someone satisfies a moral expectation, and in virtue of that satisfaction it would not be appropriate for someone else to target them with resentment. This judgment that resentment would be unjustified presupposes that the actor has some relevant capacities such that if they *hadn't* satisfied the expectation, then resentment *would have* been justified – which is to say that the actor is at least in principle a sensible target of resentment. It would seem, then, that (14) is plausible enough to take as a working hypothesis.

[32] Satisfying moral expectations in a world where those expectations are always satisfied would be something like behaving in accordance with a law of nature, and "expectation" doesn't seem to be exactly the right word for our attitude with respect to behavior governed by the laws of nature. So perhaps such a world isn't possible.

The next step involves asking about the relationship between resentment attributability and evil. As we noted above, even in a world in which resentment is never justified (assuming that such a world is possible), free creatures would still be sensible targets of resentment. In other words, it would make sense to feel resentment toward them, were they to violate some relevant expectation. But it can't make sense for an agent, S, to target another agent, T, with resentment if S doesn't even know what it's like for moral expectations to be violated. (And violation of a moral expectation is a type of evil.) A world in which resentment is never justified would have to be a world in which moral expectations are sometimes violated, but there always exists a morally relevant excuse for such violations. This suggests another intermediate impossibility proposition:

(15) Resentment attributability would not be possible without evil.

Now, given (14) and (15), we can infer the following:

(16) Reactive attributability would not be possible without evil.

This impossibility proposition is not indisputable, but we have seen some reasons to think that it's plausible. And it's important to note that these reasons should be equally acceptable to both compatibilists and incompatibilists alike. So far, then, we haven't picked up any untoward compatibilistic argumentative burdens.

The final remaining task is to argue for the value proposition:

(17) Reactive attributability is valuable enough to justify the existence of evil.

At first blush, this value proposition might seem like a non-starter. We have been focusing on resentment, and surely resentment by itself is not valuable enough to justify evil; you might even think that adding resentment to a world (so to speak) would make that world *worse*. Remember, though, that resentment is a reaction to a failure to satisfy an expectation of good will. This means that a world in which creatures are sensible targets of resentment is a world in which creatures are capable of expressing good will toward each other. And the existence of those capacities seems to be a very valuable thing. I don't have a knockdown argument in favor of this value claim, but I do think it's instructive to note just how much of what P. F. Strawson says about the importance of the reactive attitudes actually focuses on quality of will. For example:

> The central commonplace that I want to insist on is the very great importance that we attach to the attitudes and intentions towards us of other human beings, and the great extent to which our personal feelings and reactions depend upon, or involve, our beliefs about these attitudes and intentions. (1962: 5)

Strawson, in other words, wants to emphasize

> how much we actually mind, how much it matters to us, whether the actions of other people – and particularly of some other people – reflect attitudes towards us of goodwill, affection, or esteem on the one hand or contempt, indifference, or malevolence on the other. (1962: 5–6)

He also enjoins us to consider "in how much of our behaviour the benefit or injury resides mainly or entirely in the manifestation of attitude itself" (1962: 6).

The attitudes and intentions Strawson refers to here are what contribute to one's quality of will. Thus, Strawson is claiming that quality of will is of the utmost importance to us. If this is right (and it certainly seems right), then that would suggest that the capacities in virtue of which we possess and manifest quality of will are in turn extremely valuable – perhaps even valuable enough to justify the existence of evil.

It would seem, then, that a theodicy built on reactive attributability has the characteristics we have been looking for. It is determinist-friendly, it is minimally demanding, and it is freedom-relevant. It may seem like a stretch to refer to this theodicy as a *free will* theodicy, since the goods involved in reactive attributability are not directly related to free will. These goods are, however, at least first cousins of free will: susceptibility to the reactive attitudes, like free will, is an eligibility requirement for moral responsibility. Thus, I think the goods identified in a reactive attitudes theodicy are relevant enough to be included in the extended family of free will theodicies.

14.6 A Final Challenge

One challenge that will face any theodicy, as briefly alluded to earlier in the essay, is the challenge of explaining not just the existence of evil, but the disturbing frequency with which evil occurs and the unfathomable depths of so much of that evil. I have argued for the value of the capacities in virtue of which we manifest quality of will, but that point can only go so far; and it certainly seems to produce diminishing returns beyond a certain magnitude of evil. So this explanatory task remains as important

unfinished work for the reactive attitudes theodicist.[33] It's important to note, however, that in this respect a reactive attitudes theodicy seems to be roughly on a par with other popular theodicies. And given its relatively minimal theoretical commitments, any progress made on this task in the context of other theodicies will likely also be endorsable by the proponent of a reactive attitudes theodicy.

Although I have only sketched the outlines of (some essential ingredients of) a reactive attitudes theodicy, I hope to have shown that an appeal to the reactive attitudes – and to reactive attributability in particular – is a valuable resource that the theological determinist has at her disposal when facing the problem of evil. An appeal to reactive attributability doesn't require taking on board controversial theoretical commitments, since the importance of the reactive attitudes is almost universally affirmed. And that near-universal affirmation makes it clear that a theodicy built on reactive attributability will be the kind of theodicy that invites endorsement. The determinist who relies on a reactive attitudes theodicy still faces various challenges and explanatory burdens, and in the end those burdens may prove too great. But the prospects for mounting a response to the argument from gratuitous evil, at least, seem to be promising.[34]

[33] White (2019: 287–291) argues that the theological determinist can make some progress on this task by adopting either Adams's (1999) theodicy or (a slightly modified version of) Stump's (2010) theodicy.

[34] For helpful comments on earlier drafts of this essay, I am grateful to Justin Coates, Taylor Cyr, Andrew Law, Ben Mitchell-Yellin, Dan Speak, Philip Swenson, Matt Talbert, Patrick Todd, and Neal Tognazzini. For helpful advice and comments throughout the writing and publishing process, I am thankful to the editors of this volume, Leigh Vicens and Peter Furlong.