

Published in *Moral Responsibility and Alternative Possibilities: Essays on the Importance of Alternative Possibilities*. Eds. David Widerker and Michael McKenna. Burlington, VT: Ashgate Publishing Co., 2003, pp. 159-183.

FREEDOM, FOREKNOWLEDGE, AND FRANKFURT

David P. Hunt

The following essay is an exercise in "applied Frankfurtianism." Interest in this piece can therefore be expected to vary with interest in the application; and on this subject I have both good news and bad news. The bad news, for what I presume to be the majority of this volume's readers, is that the application is to theology, and that the only problem to which I claim to offer a definitive resolution is an ancient puzzle involving the theistic concept of God. The good news--which I hope is enough to keep religious skeptics reading past the first paragraph--is that belief in the existence of the theistic God is no more essential to the argument that follows than is belief in the existence of mind-readers and -manipulators essential to Frankfurt's original argument. There might even be some positive advantage to trying out a thought-experiment involving God, inasmuch as thinking about free will in light of a theological threat to alternative possibilities can offer a fresh perspective on familiar questions more often considered in the context of the scientific threat posed by universal causal determinism. (And for readers still unpersuaded, I have included a section on "blockage" counterexamples which is not strictly necessary to my theological thought-experiment.)

Frankfurt and Theological Fatalism

One of the classic puzzles raised by the theistic conception of God is whether infallible foreknowledge is compatible with libertarian freedom. Call this the problem of "theological fatalism." With 'X' designating an agent, 'A' an action, and 't' a time, the problem can be set forth as an extended hypothetical syllogism with three premises:

- (1) Necessarily, if the theistic God exists and X does A at t, then there always exists an infallible belief that X does A at t.
- (2) Necessarily, if there always exists an infallible belief that X does A at t, then it is never possible for X to refrain from A-ing at t.
- (3) Necessarily, if it is never possible for X to refrain from A-ing at t, then X's A-ing at t is not a case of libertarian free agency.

It then follows from these three premises that

- (4) Necessarily, if the theistic God exists and X does A at t, then X's A-ing at t is not a case of libertarian free agency.

Since the argument appeals to nothing that would distinguish X, A and t from any other agent, action and time, it may further be concluded that, necessarily, if the theistic God exists, there are *no* cases of libertarian free agency.

The predominant response of the argument's critics, at least among those who have weighed in on the matter since Nelson Pike's famous restatement of the problem in 1965,¹ has been to question either (1) or (2). The reason (3) has not come in for comparable scrutiny, I believe, is that it is generally thought to constitute part of the puzzle-conditions which have to be respected in canvassing solutions to the problem.² It is so regarded because Pike's argument is intended only to point to a conflict between a certain conception of God and a certain conception of free agency. The libertarian conception of freedom can certainly be rejected (as would be done by, e.g., soft determinists); so can the theistic conception of God (deists would certainly reject it). But Pike's question is whether, *given* these conceptions, a contradiction obtains.

Premise (3) has been treated as sacrosanct, then, because it appears to do nothing more than draw out an obvious implication from one of the conceptions whose consistency is in question. Premise (1), of course, performs a similar function in the argument; but some critics (going back as far as Boethius) have held it to embody a *mistaken* inference from the theistic conception of God, in contrast to the correct inference that premise (3) appears to draw from the libertarian conception of freedom.³ It is not surprising, then, that (3) has seemed the one premise strictly unavailable for revision.

The unremitting focus on the argument's first two premises has not, however, led to the kind of breakthrough around which a consensus might develop. This suggests that premise (3) might merit a fresh look, despite the countervailing considerations briefly touched on in the preceding paragraph. Harry Frankfurt's critique of the "principle of alternate possibilities" (or 'PAP'),⁴ which incidentally appeared just four years after the publication of Pike's article, provides a rather obvious basis for re-examining this premise. Yet the discussions stirred up by Pike's and Frankfurt's articles proceeded initially in complete isolation from each other, and potential connections went unnoticed (at least in the published literature). Only recently has the situation begun to change and attention been directed toward the possibility that Frankfurt's conclusions might be extended to the case where alternate possibilities are excluded by divine foreknowledge. The result is the emergence of a minority report identifying (3) as the problematic premise, alongside the majority report pinning the blame on premise (1) or (2).⁵ Because the inspiration for this report is traceable to Frankfurt's attack on PAP (and not because Frankfurt himself has had any direct hand in producing it), let us call this the 'Frankfurt Objection' to theological fatalism.

Unfortunately, the Frankfurt Objection, as it has come to be formulated, suffers from a dual weakness: it's not clear that the counterexamples on which Frankfurt's argument rests really *do* refute PAP; and even if they did, it's not clear that the critic of theological fatalism should derive much aid and comfort from this refutation. Since I think the minority report is essentially correct in assigning blame to (3), I am concerned to remedy this double flaw. To this end, I present the Frankfurt Objection against (3) and outline what seem to me to be its main weaknesses. I then show how far these weaknesses can be overcome: first, by improving on Frankfurt's counterexamples; second, by arguing that the rejection of (3) is in any case independent of the success or failure of the counterexamples (whether original or improved).

The Frankfurt Objection

At the heart of Frankfurt's argument lies a peculiar kind of counterexample. The best known of these counterexamples (call it 'BSJ' after its three principals, Black, Smith, and Jones) goes something like this. Black wishes to see Smith dead, knows that Jones shares this wish, and thinks it likely that Jones will act on this wish. The optimal scenario from Black's standpoint is for Jones to do away with Smith on his own; but Black also has a backup plan. It takes the form of a science fiction device, capable of monitoring and/or controlling a person's thoughts, which Black has set to monitor Jones's intentions with respect to killing Smith. If the information collected by the device indicates that Jones is *not* going to murder Smith on his own, the device exercises its control function to induce in Jones mental states sufficient for his killing Smith; but if the information indicates that Jones *is* going to murder Smith on his own, the device simply monitors the course of events without interference. As it turns out, Jones goes ahead and murders Smith on his own, and nothing happens to trigger the device's intervention.⁶

Two judgments about BSJ are integral to Frankfurt's argument. One is that Jones is unable to refrain from killing Smith. Conditions are such that either Jones will kill Smith on his own or the device will make him do it; there is no third possibility (at least none that Jones can actualize by his own efforts). So Jones has no alternative to killing Smith. The other judgment

is that Jones may nevertheless remain morally responsible for killing Smith. Though the case could certainly be fleshed out in such a way that Jones's culpability is undermined (e.g., by adding that Jones is really an automaton, or is responding to post-hypnotic suggestion, or has just witnessed Smith rape his daughter), nothing in the case as described warrants any mitigation of Jones's responsibility. In particular, the fact that Jones's action is inevitable fails to shake our sense that he is morally responsible for Smith's death, so long as the other conditions for moral responsibility (whatever they might be) are satisfied. The reason inevitability does not count against responsibility in this case is that the feature of BSJ which makes the killing inevitable does not figure in the sequence of events leading up to the murder. The control function of the device was never activated, and consequently played no role at all in bringing about the murder; Jones would have killed Smith just as he actually did even if the device had not existed. Given our strong inclination to hold Jones morally responsible despite the inevitability of his killing Smith, BSJ appears to provide a compelling counterexample to the principle that

PAP1 A person is morally responsible for what he has done only if he could have done otherwise.

This is the principle that Frankfurt referred to as 'PAP'.⁷

What the Frankfurt Objection needs, if it is to connect in the right way with premise (3) of the argument for theological fatalism, is a convincing critique of a somewhat different albeit related principle, namely,

PAP2 A person is (libertarianly) free in what he has done only if he could have done otherwise.

But BSJ is as good a counterexample to PAP2 as it is to PAP1, assuming that 'free' is being used here in such a way that an agent is free just insofar as he is morally responsible for his actions, or is a fit subject for the "reactive attitudes" generally.⁸ If the device made absolutely no difference to what Jones did, so that he would have performed the very same action even if his alternatives had *not* been blocked, the device's presence should not affect Jones's free agency (so understood) any more than it affects his moral responsibility. Moreover, this result is one that any theory of free will and moral responsibility (including the libertarian) must accommodate.

The compatibilist can accommodate it with equanimity and even relief, since it frees him from the irksome obligation to concoct analyses of 'could have done otherwise' which are compatible with causal determinism. But the libertarian does not have a similarly vexed relationship with robust notions of power--indeed, the denial of PAP2, by leaving open the possibility of libertarian agency in the absence of alternate possibilities, may look self-contradictory. Nevertheless, there may be just as much scope for a new libertarianism instructed by Frankfurt's attack on PAP as there is for a new compatibilism no longer engaged in producing strained and implausible conditional analyses of 'can' and 'could'. The essential insight conveyed by BSJ is that threats to alternate possibilities, insofar as these are located outside the *actual sequence* leading to Jones's action, are irrelevant to Jones's actual exercise of agency; and there is no reason why a libertarian cannot accept this moral as well. This is because the libertarian, even if forced to acknowledge conditions in which free agency is compatible with inevitability, can continue to distinguish herself from the compatibilist through her insistence that free agency is *incompatible* with causal determinism.⁹

In sum, Frankfurt's argument (if successful) is not only lethal against PAP1; it poses an equal challenge to PAP2. Since it is the latter that is directly relevant to the argument for theological fatalism, it is to this principle that I shall be referring henceforth when I mention PAP *simpliciter*. If PAP (in this version) is defective, then so is premise (3), and the fatalist's argument collapses.

This is in many respects an attractive response to the argument for theological fatalism, especially compared to responses which focus on the argument's other premises. Premise (1), for example, is often rejected on behalf of alternative concepts of God which exclude foreknowledge. This is enough to neutralize the argument for theological fatalism, since the existence *at t and later* of an infallible belief that X A's at t makes it impossible only *at t and later* for X to refrain from A-ing at t, and this has no prejudicial consequences for X's libertarian freedom in A-ing at t (which depends only on how matters stand *going into t*, not on the way things are once t has arrived). But this strategy, whatever its theological merits, leaves one with the sense that the basic issue has been avoided. Treated purely as a thought-experiment, for which the actual truth of premise (1) is unimportant, the argument for theological fatalism raises the same baffling question as thought-experiments about time-travel into the past: How could merely knowing what someone else is going to do have any effect on the foreknown action? How, in particular, could foreknowledge alone (absent other factors not logically or metaphysically entailed by it) be sufficient to transform what would otherwise have been a case of free agency into something that is not? That it could do so is highly counterintuitive, and any argument to the contrary presents us with an important conceptual problem. (The time-traveler, like an omnipresent deity, might find that knowledge of the future interferes with her *own* attempts at agency; but that's a different situation, and a different kind of interference, from the one under consideration.¹⁰) Denying premise (1) simply exempts God from complicity in this problem; it does not make the underlying conceptual puzzle disappear, nor does it do anything to address or illuminate this puzzle.¹¹

Attacks on premise (2), for their part, have little to motivate them beyond a desire to escape theological fatalism while retaining divine foreknowledge. Infallibility entails truth. If an infallible belief that X will A at t is part of the past, there are no possible futures in which that belief is false. (If there were, the belief could not count as infallible.) But then it won't be possible for X to refrain from A-ing at t; so (2) is true. Some have hoped that the "power entailment principle" on which this argument rests might prove the Achilles' heel of theological fatalism, but no one has made a convincing case that our present powerlessness over past events isn't closed under entailment.¹² Likewise the "Ockhamist" claim that God's past beliefs about the future are merely "soft" features of the past, lacking the past's immunity to present exercises of power, has little independent plausibility.¹³ For a proposition P to count as a "soft fact" at a time t, the way the world must be up through t in order for P to be true must be consistent with the world continuing *after* t in such a way that P is false. Thus *Lenin was born 47 years before the Russian Revolution* and *Lenin visited London for the last time in 1903* were both "soft" facts in 1910, while *Lenin was born in 1870* and *Lenin believed that socialism would prevail* were "hard" facts in 1910. By the same token, *God believed in 1900 that Lenin would triumph in 1917* is *prima facie* a hard fact relative to 1910. While accounts have certainly been devised under which God's past beliefs about the future qualify as soft features of the past, such analyses often seem precooked to yield this result (not to mention other respects in which they fall short).¹⁴

The Frankfurt Objection to premise (3), on the other hand, shares none of the defects of these first two approaches. Unlike some critiques of premise (2), the Frankfurt Objection is not an ad hoc move to save libertarian theism, but originates in nontheological discussions of free will and moral responsibility and derives from those discussions considerable independent plausibility, prior to its application to cases involving divine foreknowledge. And unlike critiques of premise (1), the Frankfurt Objection goes right to the heart of the problem raised by the argument for theological fatalism. Doubts that Black's device could really destroy Jones's libertarian freedom (as implied by PAP) seem to spring from the same intuitions about free agency as doubts that divine foreknowledge could really preclude X's libertarian freedom (also

implied by PAP). Frankfurt's response, which honors these intuitions by rejecting PAP, seems as suited to the second (theological) case as to the first. If Frankfurt is correct, alternate possibilities are like a symptom which can be expected to accompany free agency in the normal case but which is not (even partly) constitutive of or essential to the underlying condition of which it is symptomatic. BSJ is just one of the abnormal cases in which this symptom is absent and reliance on PAP gives a "false reading."¹⁵ According to the Frankfurt Objection, divine foreknowledge simply provides another context in which this happens. This claim is not only independently plausible, but it appears to address the deep *aporia* at the heart of the argument for theological fatalism.

Problems for the Frankfurt Objection

Despite its manifold merits, this response to premise (3) is not without difficulties of its own. There are two parts to the Frankfurt Objection. The first part affirms the defeat of PAP by BSJ (or some comparable counterexample), while the second part applies this defeat to the case of divinely foreknown actions. Both parts of the Frankfurt Objection are vulnerable. I begin with the second part.

This looks perfectly straightforward. If BSJ is a genuine counterexample to PAP, then PAP is simply false; but if PAP is false, then so is premise (3); and if (3) is false, the argument based on (1)-(3) is unsound. There does not appear to be any logical space in which one could continue to defend the argument while conceding PAP's defeat. But while this is certainly true as far as it goes, it is not enough to ensure a permanent victory over theological fatalism. The problem is that BSJ may constitute nothing more than a local exception to PAP. If that is so, then BSJ does not mandate a wholesale rejection of PAP but only a minor revision which preserves the alternate possibilities requirement for most cases (including ones involving divine foreknowledge) while relaxing the requirement in cases like BSJ.

Two factors favor continued commitment to a modified PAP. One is PAP's overall plausibility. Frankfurt himself drew attention to this, noting that PAP "has generally seemed so overwhelmingly plausible that some philosophers have even characterized it as an *a priori* truth."¹⁶ In its first version PAP is connected with the hoary principle that "ought implies can." The judgment that a person is morally blameworthy for acting in a particular way seems inseparable from the judgment that the person should have acted in some other way, and this latter judgment appears unreasonable in the extreme if circumstances were such that the person *could not* have acted in any other way. It is equally hard to imagine dispensing with PAP in its second version. When we engage in practical reasoning--the quintessential activity of free agents--it's hard not to regard the future as a domain of open alternatives among which it is possible to choose. Given PAP's intuitive force and the central place it occupies in some of our most important thinking about ourselves, a judicious conservatism in theory-adjustment seems appropriate. Whatever form this adjustment takes, it should (the present argument runs) avoid a blanket rejection of PAP based on a single trumped-up anomaly like BSJ.¹⁷

The other factor favoring a restricted PAP, which continues to govern foreknowledge cases while exempting BSJ-type cases, is the existence of significant differences between the two kinds of cases. For comparison's sake, let us set alongside BSJ the foreknowledge case that results from substituting 'Jones' and 'Jones's murdering Smith' for 'X' and 'A' in the argument for theological fatalism on the first page of this essay. Call this case, in which Black and his device have been replaced by God and His infallible foreknowledge, 'GSJ' (for 'God', 'Smith', and 'Jones'). There are some noteworthy similarities between the two cases. One is that God, like Black's device, is a "monitor-controller," albeit a maximally effective one (owing to His omniscience and omnipotence). Another is that God in GSJ is monitoring the future while refraining from controlling it, just as the device in BSJ is monitoring Jones's cognitive states

while failing to exercise control over them. But there are also important differences between the two cases. For starters, God's monitoring activity is itself sufficient to make Jones's action unavoidable (if we grant the first two premises of the argument for theological fatalism), while the monitoring performed by Black's device does not make Jones's action any less avoidable than it would have been had it remained unmonitored. Since monitoring alone is insufficient in BSJ to render Jones's murder of Smith inevitable, the control function of Black's device must be brought into the picture in some way; but its involvement must fall short of actual intervention if BSJ is to be compatible with Jones's libertarian freedom. Frankfurt's solution is to introduce the device's control function via an unactualized *disposition* to intervene. This points to the most important difference between the two cases, which is that the conditions making for inevitability in BSJ are (at least in part) counterfactual in nature, while the conditions entailing inevitability in GSJ are wholly actual.¹⁸ This does seem to mark a potentially relevant difference between the two cases. It is precisely because Black's device intervenes only counterfactually that the device makes no difference to how events unfold in the actual sequence, and it is this failure to affect the actual sequence that renders the device irrelevant to Jones's freedom and responsibility despite the unavoidability of his actions. Given our reasons for rejecting PAP in the case of BSJ, it is not irrational to think that this counterfactual element may be an essential component in exceptions to PAP. Since divine foreknowledge generates inevitability in GSJ through its actual rather than counterfactual operation, it is therefore reasonable to doubt whether PAP's failure in BSJ extends to GSJ, as the Frankfurt Objection requires.

So much for the principal weakness in the second part of the Frankfurt Objection. The problem with the first part is that, even if the moral Frankfurt draws from BSJ were transferable to GSJ, it's not clear that the moral is the right one. Two sorts of worries regarding Frankfurt's argument against PAP have surfaced in the literature: one is that BSJ (and similar counterexamples) succeed in excluding alternate possibilities only when the agent is unfree; the other is that BSJ leaves at least some alternatives open whenever the agent remains free. Contrary to Frankfurt, then, BSJ does not give us a situation in which an action is both free and unavoidable.

The first worry can be unpacked in the following way. If Jones's murder of Smith is to be inevitable (without Jones's agency being curtailed any more than is necessary to ensure inevitability), BSJ must contain

(α) a device programmed to intervene if (and only if) Jones would not otherwise kill Smith. But if BSJ is to constitute a counterexample to PAP, so that Jones's deed is (libertarianly) free despite its inevitability, then BSJ must also contain

(β) nonintervention by the device.

Now the worry here is that, given (α) and (β), BSJ must also contain

(γ) a condition causally sufficient for Jones's killing Smith.

To see why, assume that BSJ does *not* contain (γ). If this is so, then the device will have to intervene, since in the absence of any condition causally sufficient for Jones's killing Smith it is always possible that Jones will not kill Smith, and the device in (α) is required to exclude any such possibility. But this contradicts (β). Therefore the assumption that BSJ does not contain (γ) must be false. It turns out, then, that BSJ constitutes a coherent scenario only if Jones's act of murder is causally determined. Once this covert feature of BSJ is brought to the surface, the conjuring trick performed by BSJ is exposed. PAP asserts an unbreakable connection between free agency and avoidability; BSJ responds, not by breaking the connection, but by concealing it. BSJ is supposed to offer a case in which an action is both free and unavoidable, but the mechanism by which it ensures unavoidability, since it presupposes causal determinism, is actually incompatible with Jones's free agency--or so the libertarian will surely maintain. But if

BSJ flouts the libertarian's key requirement for free agency, it can hardly serve Frankfurt's purposes as an uncontroversial and definitive counterexample to PAP; and it is even less serviceable for purposes of the Frankfurt Objection to the argument for theological fatalism, since the latter specifically concerns the libertarian conception of freedom.¹⁹

The second worry regarding Frankfurt's argument against PAP is that alternate possibilities cannot be entirely eliminated from BSJ, and these residual alternatives may be enough to satisfy PAP. The alternatives in question are not alternatives to Jones's killing Smith, since there simply are none (given the presence of Black's monitor-controller); they are instead alternatives of other sorts, which may nevertheless be integral to the kind of agency Jones exercises in murdering Smith. For this defense of PAP to be successful, PAP itself must be modified so that these other alternatives can satisfy it. An example, adequate for present purposes, is this:

PAP3 A person is (libertarianly) free in what he has done (=A) only if there is something he did (=B) which is such that (i) he could have done otherwise than B and (ii) it is (at least in part) in virtue of his doing B that he is (libertarianly) free in doing A.

(This is equivalent to PAP2 when A=B.) The new formulation need not be viewed as an ad hoc shift in response to Frankfurt's attack, but can instead be seen as a corrected expression of what the principle of alternate possibilities was supposed to mean all along. Frankfurt's argument must extend to any alternative which can plausibly be thought to ground free agency; for if it does not, he can hardly claim to have defeated the "principle of alternate possibilities" (as opposed to a straw man of his own making).

What alternatives relevant to Jones's free agency are left open in BSJ? Two general scenarios are consistent with the presence of a monitor-controller like the one wielded by Black: a "nonintervention scenario," of which BSJ is an instance; and an "intervention scenario," in which Jones does something different from what he does in BSJ and thereby triggers the device to force his killing of Smith. There are a couple of respects in which an intervention scenario provides an alternative to the way things go in BSJ, and both of these have been exploited in criticism of Frankfurt's argument. One looks to what happens *after* the intervention. There is still a murder, but Jones's involvement in the murder is different than it is under the nonintervention scenario. According to William Rowe, the significant difference is that Jones's murder of Smith in the intervention scenario, because it is coerced by Black's device, is not the result of his agent-causing a decision to murder Smith, whereas in the nonintervention scenario the murder *is* the result of his agent-causing a decision to murder Smith.²⁰ For Peter van Inwagen, the crucial difference is that Jones's murder of Smith in the intervention scenario, inasmuch as it has a different causal ancestry, constitutes a different event-particular than the murder in BSJ; so while there is no alternative to the event-universal *Jones's murdering Smith*, there is an alternative to the event-particular that tokens this type in BSJ, namely, the one that tokens it in the intervention scenario.²¹ The main problem with the Rowe-Van Inwagen approach is that what happens after the intervention reflects, not what Jones can do otherwise, but what the device can do to Jones. It is therefore hard to see how this alternative could be (any part of) what grounds Jones's libertarian freedom under PAP.²²

The other respect in which an intervention scenario provides an alternative to BSJ is to be found in what happens *before* the intervention. The intervention is itself the product of Jones's acting in some way other than he acts in BSJ. This deviation from BSJ is ineliminable from the intervention scenario: even a maximally efficient monitor-controller cannot intervene quickly enough to block the mental event that triggered its intervention. Moreover, this deviation is presumably uncoerced and in other respects an example of free agency, since the monitor-controller is the only specified source of coercion in BSJ and it does not come "on line"

until after the deviation occurs. So nothing in BSJ can preclude an uncoerced deviation from the actual sequence prior to intervention by Black's device. But then this nonactual but possible deviation might serve as the alternative required by PAP. Consider, for example, Jones's *decision* in BSJ to murder Smith. While a contrary decision would presumably trigger the device's intervention and prevent Jones from carrying out his intention not to commit murder, nothing in the device's repertoire would allow it to prevent the deviant decision itself. Because Jones's decision to murder Smith could have gone the other way, it can exemplify free agency under the terms of PAP; and because Jones's murder of Smith follows in the right way from this decision, it too can qualify as an exercise of free agency, in virtue of its connection with this avoidable decision.²³

This latter approach, which finds the alternative required by PAP in an uncoerced deviation prior to the device's intervention, is far more plausible than the one that looks to Jones's coerced history subsequent to the intervention for the relevant alternative. The possibility that, e.g., Jones's decision whether to murder Smith remains within his power, despite the presence of Black's device, poses a serious threat to Frankfurt's critique of PAP. A thorough assessment of its merits would of course require a good deal more than the bare summary offered in the preceding paragraph. The same is true of the other concerns over the Frankfurt Objection canvassed in this section.²⁴ But this summary treatment is adequate for my purpose, which is to draw attention to the essential feature of BSJ that sets it up for each of these criticisms. This feature is the counterfactual nature of the mechanism by which Jones's alternatives are excluded.

To see that this is so, let us briefly review the main sources of difficulty for the Frankfurt Objection. The first part of the Frankfurt Objection is simply Frankfurt's critique of PAP based on BSJ. The problem here is twofold. First, the set-up in BSJ is alleged to presuppose a deterministic environment which begs the question against libertarianism. This is because actual control over Jones's actions is replaced in BSJ by *counterfactual* control, and this requires triggering conditions which are causally determinative of Jones's action (barring outside intervention). Second, the monitor-controller is supposed to be unable to eradicate all alternatives, and so cannot ensure PAP's failure. This is because the device operates *counterfactually*, coercing Jones's actions only after a deviation from the actual sequence has already occurred. This brings us to the second part of the Frankfurt Objection, which applies Frankfurt's argument against PAP to the case of divine foreknowledge. The problem here is that the two cases appear relevantly disanalogous. This is because the *counterfactual* nature of Black's device plays a crucial role in preserving Jones's (libertarian) freedom despite the inevitability of his actions, whereas divine foreknowledge is in actual operation, making it doubtful whether GSJ shares the same immunity from PAP as BSJ. At every juncture, then, the counterfactual element in BSJ is central to what makes the Frankfurt Objection vulnerable.

These difficulties for the Frankfurt Objection can be addressed piecemeal. This would mean showing (possibly with the help of some ad hoc modifications in Frankfurt's counterexample) why the basic set-up in BSJ does not presuppose determinism after all, and so on. This approach might even succeed; but it is not the approach I intend to pursue in the next section of this paper. What I would like to explore instead is the possibility of formulating a counterexample to PAP which does not make use of any counterfactual trigger *à la* BSJ. If this is indeed possible, then all three of the difficulties raised against the Frankfurt Objection might be amenable to a single solution.²⁵

Blockage to the Rescue?

Student Driver Jones is taking a lesson from Driving Instructor Black. The car has dual steering wheels, which work in the following way. In "normal" mode the student driver's wheel is the one engaged with the car's steering mechanism, and the driving instructor's wheel turns in

concert with the student's wheel. But if (and only if) the instructor grabs his wheel and moves it in a direction contrary to that of the student's, the student's wheel becomes disengaged from the steering mechanism and control over the car's direction shifts to the instructor's wheel. Now imagine that Black and Jones are on a narrow country lane with an unbroken stretch of stone wall just inches away on either side of the car. At the top of a steep hill the car loses its brakes. One hundred yards ahead the road forks: Smith is walking down the right-hand fork with his back to the car; the left-hand fork leads harmlessly off across the countryside. Given his animus toward Smith and thinking what an excellent excuse for murder the faulty brakes provide, Jones steers the car to the right and kills Smith. Black watches with approval, never touching his own wheel.

There is nothing in this case as described to indicate that Jones is other than (libertarianly) free and morally responsible in what he does. If we assume further that nothing prevents Jones from turning left at the fork, we have the alternate possibilities required by PAP. This is not very interesting, however, since Frankfurt's critique of PAP requires a case in which there are no alternate possibilities. So let's look at three ways the case could be embellished to make the killing unavoidable. The cases are distinguished by the different ways in which alternatives are eliminated. In introducing these cases I will consider only the relationship in which Jones stands to the car's direction. Once the three kinds of alternative-eliminators are clear, we can look at what happens when internal factors, such as Jones's intentions, are brought into consideration.

(1) The first of these embellished cases involves what we might call an "active alternative-eliminator." An example is where Black, inspired by his own animus against Smith, takes advantage of the situation to seize control of the car, directing it down the right-hand fork and into Smith. There are a number of ways Jones's part of the story could be fleshed out at this point, but these are unimportant for present purposes, since Black's ultimate control over the car's direction means that Jones does not freely kill Smith. Indeed, Jones does not kill Smith at all--it's Black who employs the car as the instrument of his agency. So this case is not even in contention as a counterexample to PAP.

(2) The next case, like BSJ, involves a "counterfactual alternative-eliminator." Here Black forms the resolve to override Jones's steering wheel if (and only if) Jones shows any inclination to steer the car to the left at the fork, but Jones shows no such inclination and runs Smith down without Black's intervention. In comparison with the previous case and its active alternative-eliminator, the strategy here is to restrict interference with Jones's agency to nonactual worlds, leaving Jones the autonomous agent of Smith's death in the actual sequence. Jones's killing of Smith can therefore exemplify free agency despite its inevitability; as a counterexample to PAP, however, it is open to the same objections as BSJ.

(3) We can now contrast the BSJ-like counterfactual alternative-eliminator in (2) with a rather different kind of alternative-eliminator, which we might call a "passive alternative-eliminator."²⁶ Suppose the driving instructor can lock his wheel at a certain position to prevent the student driver from steering beyond that range, and Black has placed a "left lock" on his steering wheel to block the possibility that Jones might take the road to the left; Jones, however, bears right at the fork and never encounters the lock.²⁷ The principal difference between this kind of case and the one involving the counterfactual alternative-eliminator is that the passive eliminator is in place in the actual world, though the sequence of events actually productive of Smith's death never intersects with it (hence its "passiveness"). But the moral it conveys appears to be the same. A steering lock is no less effective than is Black's counterfactual resolve in ensuring that the car is going to hit Smith and that there is nothing Jones can do to avoid this outcome. Moreover, there is no less reason in this case to regard Jones as a free agent in killing

Smith. The passive alternative-eliminator does not figure in the actual sequence; in its absence, Jones would have done everything the same. If these reasons support Jones's free agency in the face of a counterfactual alternative-eliminator, they equally support his free agency when a passive alternative-eliminator is at work.²⁸

Neither the counterfactual nor the passive case eliminates all alternatives; so while the left-hand alternative to the right-hand fork is indeed eliminated, there remain other alternatives to Jones's actual behavior (e.g., he could steer the car toward the left until control shifts to Black), and these alternatives might be enough to satisfy PAP (in its PAP3 version). Fine-tuning the examples can certainly narrow the range of alternatives considerably. As we saw in the preceding section, however, there is an upper limit on how far the counterfactual alternative-eliminator can be tightened, since its trigger-structure requires that *some* alternative be accessed before the mechanism comes "on line." There is, on the other hand, no evident upper limit on the restrictions imposed by a passive alternative-eliminator. Any alternative can be passively blocked; and because the alternative is eliminated passively, the actual sequence, along with Jones's free agency, is unaffected.

Suppose that Black has taken hold of his wheel and begun to steer. Everything happens as in the first case, involving the active alternative-eliminator, except that by sheerest coincidence Jones's steering motions happen to correspond precisely with those of Black (and Jones fails to notice Black's attempted intervention).²⁹ Because the wheels never move in opposition to each other, control over the car's steering mechanism never passes to Black. But it is also true that Jones is powerless to send the car in a different direction: given the steering arrangement, so long as the driving instructor is gripping his wheel, the car cannot go in a contrary direction. Jones has the power to move his own wheel differently, but this would make no difference to the car's direction, since effective control over the car's steering mechanism would thereby shift to Black's wheel.

What are we to say about Jones's free agency in this case? The same thing, I should think, that we would say about the other cases involving a passive alternative-eliminator (not to mention the case involving the counterfactual alternative-eliminator). Jones's relationship to the car's direction is that of a free agent, even a libertarian free agent. Jones is the one whose wheel is engaged with the car's steering mechanism. He is not being compelled or forced to move the wheel as he is doing. The passive alternative-eliminator--Black's control over his own wheel--contributes nothing to the sequence of events bringing about Smith's death; remove the alternative-eliminator and the actual sequence would be completely unaffected.

So far we have restricted our attention to the overt acts of wheel-turning by which Jones endeavors to control the car's direction. This ignores mental acts, like planning and deciding, which are also critical to free agency and may help to satisfy PAP if alternatives to them remain open. But there is no reason to think that these alternatives cannot be eliminated in the same way (and with the same consequences for Jones's free agency): the relationship between Jones and the car's direction appears to model the relationship between an agent and any action of that agent, no matter how immediate. Suppose that Jones, in the original BSJ, is a paragon of free agency; and consider the actual sequence *S* of Jones's relevant states, mental and otherwise, leading up to his murder of Smith. *S* will of course be endowed with whatever *actual* virtues a sequence must possess if it is to demonstrate free agency. Now add to the picture a super-duper passive alternative-eliminator which, by operating on Jones's neural connections, has managed to block in advance all the alternatives that would ordinarily be accessible from *S*. (That the device happens to block all and only alternatives to *S* can again be put down to chance.) This is essentially the same case as the earlier ones involving passive alternative-eliminators operating on external behavior. Consider, for example, the point in *S* at which Jones decides to murder

Smith. The passive alternative-eliminator, we are assuming, has blocked all alternatives to Jones's decision to murder Smith, but has not interfered with the actual decision. So there are no grounds for thinking that Jones would have decided differently if the alternative-eliminator had been removed; there are no grounds for thinking that the actual sequence involves any coercion; there are no grounds for thinking that Jones isn't agent-causing his decision to murder Smith; in short, there are no grounds for thinking that Jones's free agency has been affected in any way, save for a prior commitment to PAP.

Call this passive version of BSJ, 'BSJ_P', and the original (counterfactual) version of BSJ, 'BSJ_C'. John Martin Fischer has coined the term "blockage case" to describe my use of passive alternative-eliminators in an earlier article³⁰ and in personal correspondence, and the name has stuck.³¹ The burden of the preceding remarks has been to argue that a blockage case like BSJ_P works at least as well as (if not better than) a standard counterfactual case like BSJ_C in eliminating alternatives while preserving the agent's freedom. Can "blockage" then be used successfully to refute PAP in support of the Frankfurt Objection to theological fatalism? Since it is the counterfactual character of BSJ_C that lays it open to the three difficulties raised against the Frankfurt Objection, we must see how the counterfactual-free BSJ_P fares when confronted with the same three difficulties.

The first difficulty is that Frankfurt's critique of PAP might covertly presuppose causal determinism, since the triggers for counterfactual intervention are apparently required to provide causally sufficient guarantees for Jones's behavior. Since there are no counterfactual triggers in BSJ_P, this particular difficulty is no longer an issue.

The second difficulty is that Frankfurt's critique of PAP employs an alternative-eliminator which cannot eliminate all alternatives--in particular, it cannot eliminate those alternatives that serve as triggers for the eliminator. This is not a problem for BSJ_P, since it involves no triggers. A passive alternative-eliminator avoids entanglement in the actual sequence, not by confining its operations to nonactual scenarios (accessible only via counterfactual triggers), but by blocking untraveled side roads in the actual world.

In eschewing counterfactual triggers, however, BSJ_P may generate a new problem, and this should be acknowledged before moving on to the third difficulty for the Frankfurt Objection. By causally eliminating *all* alternatives to *S*, BSJ_P brings it about that only *S* is causally possible. But then every state or event of possible relevance to Jones's exercise of agency is causally determined by the device, leaving no opening for the causal indeterminism critical to libertarianism.³²

It is clear how the critic of PAP who has adopted BSJ_P as her prize counterexample will want to respond to this worry. The key Frankfurtian insight is that what happens in the actual sequence is all that's relevant to judgments of free agency and moral responsibility. If Jones's murder of Smith, along with such crucial preliminaries as Jones's decision to murder Smith, are determined by a causal chain operating within the actual sequence itself, then the libertarian must deny that Jones is functioning as a free agent; but if Jones's agentially relevant states are determined by causal factors operating outside the actual sequence, the libertarian who has taken Frankfurt's critique of PAP to heart might well deny that such causal determinism counts against Jones's freedom. This strategy obviously requires a pretty solid grip on the notion of an actual sequence, and clear criteria for when something would (or would not) belong to such a sequence. This is a large and important subject which cannot be pursued adequately here. Intuitively, one wants to say that the actual sequence in BSJ_P consists of those conditions, sufficient for Jones's action, which supply an apt substitution for the ellipsis in "Jones murdered Smith because . . ." (When more than one substitution satisfies this requirement, there is overdetermination in the

actual sequence.) Now there does seem to be a sense in which, given BSJ_P , S provides an apt substitution but the passive alternative-eliminator does not, even if the latter causally determines Jones's action. Jones murders Smith because of the deliberations, decisions, intentions, etc., to be found in S ; he does not murder Smith because of the alternative-eliminator. So the claim that the actual sequence in BSJ_P can be distinguished from the conditions causally determining Jones's murder of Smith, in just the way required by Frankfurt's critique of PAP, has considerable prima facie plausibility. In the absence of a satisfactory analysis of "actual sequence" and the "because"-relation, however, this discussion of the new objection must be judged inconclusive. (Providing such an analysis may be the most important job facing post-PAP metaphysics.)

Finally, the third difficulty is the disanalogies between BSJ_C and GSJ which make the former's success against PAP (such as it is) nontransferable to the latter, particularly given the general presumption in PAP's favor. The most important of these disanalogies is that Black's device operates counterfactually, whereas divine foreknowledge is in actual operation; and this seems important in turn because the counterfactual set-up in BSJ_C helps to explain how there can be unavailability without any interference in Jones's actual exercise of agency, whereas no comparable explanation is available for an active mechanism like foreknowledge. This difficulty is less troublesome when BSJ_C is replaced by BSJ_P , since it shows that counterfactual mechanisms are not the only way to avoid entanglement in the actual sequence. Still, it is not clear that GSJ is more analogous to BSJ_P than it is to BSJ_C ; so this difficulty remains a real one.

In sum, the first two difficulties lapse completely, though the second is replaced by a new difficulty of doubtful force; and the third difficulty survives, albeit somewhat mitigated. This certainly marks an advance over BSJ_C , though lingering worries mean that the Frankfurt Objection is not yet home free. Bringing it home requires addressing the third difficulty, and an adequate response to this difficulty (as it turns out) provides a sufficient response to the residual second difficulty as well.

The Self-Refutation of Theological Fatalism

Why, again, are disanalogies between GSJ and BSJ (whether counterfactual or passive) supposed to undermine the Frankfurt Objection? The idea, apparently, is that the decisive counterexample to PAP must come from BSJ or the general vicinity; once a successful BSJ -type counterexample is in place, other claims for exemption from PAP must then be based on relevant similarities to this *Ur*-counterexample. What this analysis ignores, however, is the possibility that GSJ is itself a perfectly suitable counterexample to PAP, requiring no support from BSJ and its ilk.

Divine foreknowledge of Jones's action, in virtue of its infallibility, is just as effective as any BSJ -type device in precluding alternatives; and in virtue of its omniscience, it is at least as effective as a perfect passive alternative-eliminator in rendering the entire actual sequence unavoidable, leaving no alternatives that could support Jones's free agency under any reasonable version of PAP. At the same time, the claim that PAP gives the right result in GSJ is every bit as counterintuitive as it is in BSJ . Recall what it is about Frankfurt's original counterexample that is supposed to leave Jones (libertarianly) free with respect to Smith's murder despite its inevitability. Frankfurt himself characterized the relevant features this way: Black's device "played no role at all in leading [Jones] to act as he did;" indeed, "everything happened just as it would have happened without Black's presence in the situation and without his readiness to intrude into it;" for this reason the counterfactual intervener is "irrelevant to the problem of accounting for a person's action" and "does not help in any way to understand either what made [Jones] act as he did or what, in other circumstances, he might have done."³³ But to whatever

extent this is true of the alternative-eliminator in Frankfurt's original example (not to mention the one in BSJ_P), it is equally true of divine foreknowledge in GSJ. God's infallible beliefs about the future rule out alternate possibilities without their making any difference to, or playing any role in, or helping in any way to explain the future. If these grounds for rejecting PAP are cogent in BSJ, they should be cogent in GSJ as well; and if they lack cogency in the latter case, it's hard to see how they could have it in the former.

If GSJ replaces BSJ as the critical counterexample to PAP, the third difficulty for the Frankfurt Objection becomes moot (GSJ can hardly be disanalogous to itself!). And it obviously escapes the first two difficulties, since (like BSJ_P and unlike BSJ_C) it does not employ any counterfactual triggers. Finally, the additional worry generated by BSJ_P does not even arise for GSJ. The problem for BSJ_P is that a perfect passive alternative-eliminator might leave the actual sequence causally determined, contrary to the requirements of libertarian freedom. GSJ's answer to this potential problem is to deploy as its alternative-eliminator a divine cognition which is not related to its object as cause to effect. This is not to say that foreknowledge is *never* causally responsible for what it foreknows--perhaps it is in some contexts. God might send Wanda a vision of herself winning the lottery, and make the vision so striking and self-certifying that Wanda can have no doubt of its veracity; this foreknowledge then leads her to a liquor store where she buys the winning ticket. Or if one prefers a miracle-free example, Stephen might foreknow that he will teach his ethics class today at 9:00 AM (it's the class's scheduled meeting time and he is conscientious about his duties), but the complex action-guiding interplay between belief and intention is such that his foreknowledge ends up playing a causal role in getting him to class.³⁴ In such cases, however, foreknowledge is causative of what it knows only because some special condition, beyond the knowledge itself, has been added (e.g., that the foreknower is making practical use of her foreknowledge). But no such condition is included in GSJ, nor is any presupposed in the argument for theological fatalism. Here it is only raw, unadorned, antecedent belief, of the infallible and omniscient variety, that is entailing inevitability. In a case like this, God's foreknowledge that Jones will murder Smith precludes Jones's alternatives without causally determining his action.³⁵

GSJ's success on this score can also be brought out by considering the distinctive way in which GSJ isolates its alternative-eliminator from the actual sequence. BSJ_P does it with a device that operates on untaken pathways; BSJ_C does it with a mechanism that operates in unactualized worlds; but GSJ does it by introducing its alternative-eliminator only *after* the actual sequence has been completed. Temporally speaking, of course, foreknowledge does not come after but *before* its object. What is relevant here, however, is the logical or explanatory order, not the temporal order. Knowledge of contingents normally *follows* its objects in the order of explanation, no matter where it is located in the order of time: the *fact* that *p* may help explain or account for someone's *knowledge* that *p*, but the knowledge does not explain or account for the fact. Recall our intuitive understanding of an actual sequence as those conditions, sufficient for Jones's action, which complete the phrase "Jones murdered Smith because . . ." "God foreknew that Jones would murder Smith" satisfies the sufficiency requirement but does not correctly fill the ellipsis; so God's foreknowledge remains unimplicated in the actual sequence, despite preceding it in time.³⁶ As a counterexample to PAP, GSJ is therefore unrivaled in its capacity for eliminating all of Jones's alternatives while leaving the actual sequence unaffected (again barring the presence of extra features which make the future explanatorily dependent on God's knowledge of it).

Despite the brief just offered on its behalf, I suspect that GSJ will be viewed by many with lingering suspicion, and regarded as far less convincing than a comparable counterexample

drawn from the neighborhood of BSJ. I want now to consider three reasons why this might be so, only the last of which raises a serious difficulty for GSJ.

The first is the sense that it is ad hoc or question-begging to cite GSJ as a counterexample to PAP for purposes of defending the Frankfurt Objection. The Frankfurt Objection, after all, is essentially the claim that PAP breaks down in cases like GSJ, and answering an attack on this objection with the rejoinder that PAP breaks down in cases like GSJ may seem circular. Note that the objector is not claiming that GSJ, taken simply as a counterexample to PAP, is specially defective. (She might even agree that GSJ is *better* than BSJ when it comes to counterinstancing PAP.) What the objector *is* claiming, however, is that GSJ begs the question against PAP *in a particular context*, namely, when PAP is being marshaled on behalf of theological fatalism. As a counterexample to PAP, GSJ must be such that

(G1) If PAP is true, then Jones is not a free agent in GSJ.

(G2) But Jones *is* a free agent in GSJ.

(G3) Therefore PAP is not true.

But the argument for theological fatalism presupposes the soundness of

(P1) If PAP is true, then Jones is not a free agent in GSJ.

(P2) But PAP *is* true.

(P3) Therefore Jones is not a free agent in GSJ.

What concerns the objector is that anyone who rejects (P2) by citing GSJ is simply counterposing the *modus tollens* of (G1)-(G3) to the *modus ponens* of (P1)-(P3); and that is to beg the question against (P2).

If this is how the objection goes, it proves too much; for the objection, if sound, would show that *every* refutation by counterexample is circular. Such refutations take the following form: if principle P is true, then case X is F; but X is not F; so P is not true. Against the critic's *modus tollens*, the defender of P can *always* raise the corresponding *modus ponens*. But it surely does not follow from this that all refutations by counterexample commit the fallacy of *petitio principii*. Who (if anyone) is begging the question depends on the independent cases to be made for the arguments' differing second premises. Likewise, the dialectical situation that obtains when the Frankfurt Objection appeals to GSJ to support its denial of PAP depends entirely on how strong a case can be made out on behalf of (G2). Appreciating the case for (G2) involves, among other things: coming to GSJ with a (relatively) open mind; examining it with due care and deliberation; noting its characteristic features (e.g., that divine foreknowledge does not force or compel Jones to act as he does, that his action is logically prior to God's knowledge of it, that he would have done everything the same if God's foreknowledge were subtracted from the picture, etcetera); and bringing to one's meditation on GSJ the same pretheoretical intuitions about free agency that are supposed to be reflected in PAP. Such considerations won't convince everyone that PAP goes subtly wrong in capturing those intuitions; but insofar as anyone *is* persuaded by these considerations, that person has a good (noncircular) reason for rejecting arguments like (P1)-(P3) that are based on PAP.

A second source of doubts about GSJ as a counterexample to PAP is the simple fact that it is less psychologically compelling than BSJ--perhaps because thought-experiments whose materials are drawn from science fiction are easier to credit, imagine, or think about than ones constructed from theological materials. (The reverse would undoubtedly have been true 1000 years ago.) The short answer here is that an impoverished imagination does not entitle anyone to veto-power over an argument. But this flippant response does not give the present objection its due. The fact is that, over the last thirty years, the percentage of those considering BSJ who have found it to pose a *prima facie* challenge to PAP has been fairly large, while the percentage of

those thinking about divine foreknowledge scenarios like GSJ who have reached a similar conclusion has been minuscule. So BSJ evidently enjoys some sort of practical advantage over GSJ as a counterexample to PAP.³⁷ But this advantage is in the end compatible with the claims I have made on behalf of GSJ. There are at least two senses in which one counterexample to P may be "better" than another: (i) the first may be more effective than the second in getting some person(s) to see that P is false; or (ii) the first may come closer than the second to constituting conditions actually falsifying P. My defense of GSJ has focused on the second sense. This does not exclude the possibility that BSJ_C or BSJ_P is better in the first sense, providing the catalyst for doubts about PAP which GSJ then brings to fruition.³⁸

Finally, GSJ may look dubious as a counterexample to PAP owing to fears that divine foreknowledge might be logically incoherent. This possibility cannot be entirely dismissed. In his original article, Pike noted the conceptual peculiarity of connecting truth analytically with divine belief. Divine knowledge differs from human knowledge, Pike wrote, in that "it is *not* an additional, contingent fact that the belief held by God is true. 'God believes X' entails 'X is true.'³⁹ Pike did not suggest that this is anything more than a *difference* in the logic of divine 'belief'; but conceptual peculiarity is a step in the direction of conceptual incoherence. If a divine forebelief is analytically connected with the foreknown event, it's hard to see how the two could retain sufficient independence from each other to count as anything like a recognizable belief or event.⁴⁰ (This crux in the concept of infallible prescience is probably responsible for the competing intuitions over whether divine forebeliefs are "hard" or "soft" facts about the past.) None of this shows that infallible foreknowledge *is* logically incoherent; but there is enough to worry about here that its coherence cannot simply be assumed.

If we *could* assume it, the best counterexample to PAP would be GSJ itself. But given reasonable doubts about the logical coherence of infallible foreknowledge, this conclusion must be conditional in form, to wit:

If infallible foreknowledge is logically coherent, PAP is false.

If we add to this conditional the one that seems supported by BSJ_P, namely,

If a passive alternative-eliminator can eliminate all relevant alternatives without causally determining or otherwise affecting the actual sequence, PAP is false,

and the one supported by BSJ_C, namely,

If a counterfactual alternative-eliminator can eliminate all relevant alternatives without presupposing a deterministic environment inimical to libertarianism, PAP is false,

we have the makings of a cumulative case argument against PAP. The argument is not probative, however, so long as it remains possible to deny all three antecedents.

When we turn from PAP to the problem of theological fatalism, however, it is possible to be more definite. We have seen that, if infallible foreknowledge is logically coherent, then GSJ is available as a decisive counterexample to PAP; and if it is, the argument for theological fatalism fails, since premise (3) of the argument depends on PAP. But of course the argument also fails if infallible foreknowledge is *not* logically coherent, since this concept plays an essential role in premises (1) and (2). This allows an unconditional verdict on the argument to be reached via constructive dilemma. The argument for theological fatalism fails to prove its conclusion, and the fundamental reason it fails may be brought out as follows:

- (a) If infallible foreknowledge is logically coherent, PAP is false (because GSJ is then available as a counterexample to PAP).
- (b) If PAP is false, the argument for theological fatalism fails (because there is then no reason to think that premise (3) is true).
- (c) So if infallible foreknowledge is logically coherent, the argument for theological fatalism fails.

- (d) But if infallible foreknowledge is *not* logically coherent, the argument for theological fatalism also fails (because the argument's alternative-eliminator cannot then exist).
- (e) Either infallible foreknowledge is logically coherent or infallible foreknowledge is not logically coherent.
- (f) Therefore the argument for theological fatalism fails.

The Frankfurt Objection to theological fatalism is therefore at least semi-correct. Assuming that the constituent concepts of the argument for theological fatalism are coherent, premise (3) of the argument is the one that must be rejected. (And if its constituent concepts are not coherent, the argument doesn't even get out of the starting gate.)

The Frankfurt Objection does misrepresent the situation in one respect, however. Its reliance on Frankfurt's critique of PAP, for purposes of defeating the argument for theological fatalism, looks initially like a case in which developments in the metaphysics of freedom provide succor to a crisis in philosophical theology. But given the key role played by premise (a) in our counterargument, the final version of the Frankfurt Objection looks more like a classic puzzle in philosophical theology shedding light on a contemporary controversy in the metaphysics of freedom.⁴¹

NOTES

¹Nelson Pike, "Divine Omniscience and Voluntary Action," *Philosophical Review* 74 (January 1965), pp. 27-46.

²That it is specifically *libertarian* free agency that is in question was not explicitly recognized during the first phase of the debate over Pike's article; indeed, Pike himself, following St. Augustine's classic formulation of the problem in *On Free Choice of the Will*, spoke simply of *voluntary* agency, without elaborating on how he understood this notion. William P. Alston, in "Divine Foreknowledge and Alternative Conceptions of Human Freedom," *International Journal for Philosophy of Religion* 18 (1985), pp. 19-32, demonstrates that the argument presupposes a libertarian conception of freedom, and shows how lack of clarity on this point has led some contributors to the debate to talk at cross-purposes with each other.

³If 'always' means *at every point in time*--which is how I am using it in premise (1)--then Boethius rejects this premise on the grounds that God and His cognitions are not in time at all. Boethius's modern followers include Eleonore Stump & Norman Kretzmann, "Eternity," *Journal of Philosophy* 78 (1981), pp. 429-58, and Brian Leftow, "Timelessness and Foreknowledge," *Philosophical Studies* 63 (1991). Another possibility for rejecting premise (1), without placing God outside time, is that God's beliefs about future contingencies are not infallible. Theists who have adopted this position include Peter Geach, *Providence and Evil* (Cambridge: Cambridge University Press, 1977), ch. 3; Richard Swinburne, *The Coherence of Theism* (Oxford: Oxford University Press, 1977), ch. 10; and William Hasker, *God, Time, and Knowledge* (Ithaca & London: Cornell University Press, 1989).

⁴"Alternate Possibilities and Moral Responsibility," *Journal of Philosophy* 66 (December 1969), pp. 829-39.

⁵Contributors to this "minority report" include Eleonore Stump, "Intellect, Will, and the Principle of Alternate Possibilities," in Michael Beaty, ed., *Christian Theism and the Problems of Philosophy* (Notre Dame: Notre Dame University Press, 1990), pp. 254-85; Linda Trinkaus Zagzebski, *The Dilemma of Freedom and Foreknowledge* (New York: Oxford U. Press, 1991); and John Martin Fischer, *The Metaphysics of Freedom*, Aristotelian Society Series, v. 14 (Oxford, UK, & Cambridge, Mass.: Blackwell, 1994). I identify St. Augustine as the first "Frankfurtian" in my "Augustine on Theological Fatalism: The Argument of *De Libero Arbitrio* III.1-4," *Medieval Philosophy and Theology* 5 (1996), pp. 1-30, and I argue in favor of this Augustinian/Frankfurtian approach to the problem of theological fatalism in my "On Augustine's Way Out," *Faith and Philosophy* 16 (January 1999), pp. 3-26.

⁶Frankfurt himself is a bit vaguer than this, though BSJ is consistent with the considerations that he marshals against PAP and is a natural development from his suggestion that "Black manipulate[s] the minute processes of Jones's brain and nervous system . . . so that causal forces running in and out of his synapses and along the poor man's nerves determine that he chooses to act and that he does act in the one way and not in any other" (pp. 835-36).

⁷*Ibid.*, p. 829.

⁸This is Peter Strawson's term for such other-directed attitudes as gratitude, resentment, indignation, esteem, contempt, benevolence, forgiveness, and the like. See his "Freedom and Resentment," *Proceedings of the British Academy* 48 (1962), pp. 1-25.

⁹Frankfurt left open the possibility that the actual sequence in BSJ is causally undetermined, precisely in order to accommodate the intuitions of libertarians. John Martin Fischer maps the logical space for this libertarian appropriation of Frankfurt, as well as highlighting the importance of the "actual sequence" in judgments of free agency and moral responsibility, in his excellent "Responsibility and Control," *Journal of Philosophy* 89 (January 1982), pp. 24-40.

¹⁰I discuss this problem in "Omniprescent Agency," *Religious Studies* 28 (September 1992), pp. 351-69, and in "Divine Providence and Simple Foreknowledge," *Faith and Philosophy* 10 (July 1993), pp. 394-414.

¹¹A fuller treatment of this issue may be found in my "What Is the Problem of Theological Fatalism?" *International Philosophical Quarterly* 38 (March 1998), pp. 17-30.

¹²In making this claim, it should be obvious that I am not thereby claiming that every *formulation* of this closure principle is problem-free! Critics like Michael Slote, in "Selective Necessity and the Free-Will Problem," *Journal of Philosophy* 79 (1982), pp. 5-24, have shown that power-necessity is not always well-behaved, and David Widerker ("On an Argument for Incompatibilism," *Analysis* 47 (1987), pp. 37-41) and Thomas J. McKay & David Johnson ("A Reconsideration of an Argument against Compatibilism," *Philosophical Topics* 24 (Fall 1996),

pp. 113-122) have described counterexamples to one of the best known transfer principles, Peter van Inwagen's rule β . But these counterexamples do not show that the idea underlying the flawed formulations is fundamentally wrong--van Inwagen, for example, has recently revised rule β in light of the critique of McKay and Johnson (see his "Free Will Remains a Mystery." in *Philosophical Perspectives* 14: Action and Freedom (2000), pp. 1-19.) For an excellent discussion which distinguishes power-necessity from other modal contexts in which the closure of necessity under entailment is problematic, see John Martin Fischer, "Power Necessity," *Philosophical Topics* 14 (1986), pp. 77-91; cf. also Timothy O'Connor, "On the Transfer of Necessity," *Nous* 27 (1993), pp. 204-218.

¹³This strategy is named after William Ockham, who first developed it. See his *Predestination, God's Foreknowledge, and Future Contingents*, trans. with intro., notes, and appendices by Marilyn McCord Adams & Norman Kretzmann (New York: Appleton-Century-Crofts, 1969). It is probably the most popular strategy among contemporary critics of the argument. Important (and to my mind persuasive) critiques of the "Ockhamist" position may be found in John Martin Fischer, "Freedom and Foreknowledge," *Philosophical Review* 92 (January 1983), pp. 67-79; William Hasker, "Hard Facts and Theological Fatalism," *Nous* 22 (September 1988), pp. 419-36; and David Widerker, "Troubles with Ockhamism," *Journal of Philosophy* 87 (1990), pp. 462-80. See also Nelson Pike's "A Latter-Day Look at the Foreknowledge Problem," *International Journal for Philosophy of Religion* 33 (June 1993), pp. 129-64.

¹⁴One Ockhamistic analysis that escapes this charge may be found in Alvin Plantinga's "On Ockham's Way Out," *Faith and Philosophy* 3 (July 1986), pp. 235-69, and again (in somewhat different guise) in Ted A. Warfield's "Divine Foreknowledge and Human Freedom are Compatible," *Nous* 31 (March 1997), p. 80-86. The Plantinga-Warfield line goes something like this: according to the classical theist, it is a necessary truth that God exists and is omniscient; so *It was the case that X will A at t* both entails and is entailed by *God believed that X will A at t*; since these propositions are broadly logically equivalent, they cannot differ with respect to their hardness/softness; but the former is clearly a soft fact about the past; therefore the latter must be a soft fact as well. Though not "precooked" (inasmuch as it purports to *prove* the "softness" of God's beliefs about the future), this analysis is nevertheless unconvincing, as I show on pages 14-16 of my "On Augustine's Way Out," *op. cit.*

¹⁵In his "Responsibility and Control," *op. cit.*, John Fischer argues (pp. 37-38) "that what is wrong with lack of control [i.e., freedom to do otherwise] is that it usually (but not always) indicates actual-sequence compulsion. Thus the reason why lack of control normally rules out responsibility is that it normally points to actual-sequence compulsion. But when lack of control is not accompanied by actual-sequence compulsion, we need not rule out responsibility."

¹⁶"Alternate Possibilities and Moral Responsibility," p. 829.

¹⁷These are not only reasons for limiting the damage if PAP does prove vulnerable to counterexample; they are also reasons for suspecting that any proposed counterexample must (somehow) be mistaken. Any anti-PAPist (and not just the Frankfurt Objector) is therefore obligated to explain why the two pro-PAP considerations canvassed in this paragraph are deeply misleading despite their surface attractiveness. I have addressed the second of these considerations in a number of places, most fully in "Omniprescient Agency," *op. cit.* The first and tougher of the two--the "ought-implies-can" objection--has been transformed by David Widerker into an independent argument in favor of PAP. Widerker calls it the "what-should-he-have-done-instead" defense, or the "W-defense". See his "Theological Fatalism and Frankfurt Counterexamples to the Principle of Alternative Possibilities," *Faith and Philosophy* 17 (April 2000), pp. 249-254, and his "Frankfurt's Attack on the Principle of Alternative Possibilities: A Further Look," *Philosophical Perspectives* 14 (2000), pp.181-201, where he develops this

defense more fully. My reply to Widerker--"On a Theological Counterexample to the Principle of Alternate Possibilities," *Faith and Philosophy* (forthcoming)--gestures in what I believe to be the right direction, but is admittedly short on details.

¹⁸At least divine foreknowledge is traditionally understood as actual or "occurrent" in nature. For an alternative conception, see my "Dispositional Omniscience," *Philosophical Studies* 80 (December 1995), pp. 243-78. I show how far this conception can (and cannot) go in solving the problem of theological fatalism in my "Does Theological Fatalism Rest on an Equivocation?" *American Philosophical Quarterly* 32 (April 1995), pp. 153-65.

¹⁹David Widerker makes a similar argument against certain Frankfurt-style counterexamples (but without the application to theological fatalism) in his "Libertarianism and Frankfurt's Attack on the Principle of Alternative Possibilities," *Philosophical Review* 104 (April 1995), pp. 247-61.

²⁰William L. Rowe, "Two Concepts of Freedom," *Proceedings and Addresses of the American Philosophical Association* 61 (1987), pp. 43-64.

²¹Peter van Inwagen, "Ability and Responsibility," *Philosophical Review* 87 (April 1978), pp. 201-24.

²²In "Responsibility and Control," John Fischer claims (correctly, I believe) that van Inwagen "confuses the ability deliberately to do otherwise with the possibility of something different occurring" (p. 32).

²³David Widerker pursues this defense of PAP in his "Libertarianism and Frankfurt's Attack on the Principle of Alternative Possibilities," *op. cit.* See also his "Libertarian Freedom and the Avoidability of Decisions," *Faith and Philosophy* 12 (January 1995), pp. 113-118.

²⁴I look at some of the discussion of these issues in my "Frankfurt Counterexamples: Some Comments on the Widerker-Fischer Debate," *Faith and Philosophy* 13 (July 1996), pp. 395-401.

²⁵I reach a similar conclusion about the counterfactual element in standard Frankfurt counterexamples in my "Moral Responsibility and Unavoidable Action," *Philosophical Studies* 84 (January (II) 2000), pp. 195-227, and describe three "improved" counterexamples that do not employ counterfactual interveners. The next section of the present essay develops one of the three counterexamples from this earlier article.

²⁶I owe this term to David Widerker.

²⁷In *The Metaphysics of Freedom* (p. 133) John Fischer presents a similar case, in which a "car's steering apparatus *works properly* when I steer the car to the right. But unbeknownst to me, the car's steering apparatus is broken in such a way that, if I were to try to turn it in some other direction, the car would veer off to the right . . . Since I actually do not try to do anything but turn to the right, the apparatus functions normally and the car's movements are precisely as they would have been, if there had been no problem with the steering apparatus."

²⁸Cases which include passive alternative-eliminators are certainly not new. In addition to Fischer's malfunctioning steering mechanism mentioned in the previous footnote, there is van Inwagen's inoperative telephone which would have prevented someone from calling the police if he hadn't decided against it on his own ("Ability and Responsibility," p. 205), not to mention John Locke's famous example of a man who remains willingly in a room which he cannot in fact leave because it has been surreptitiously locked from the outside (*An Essay concerning Human Understanding*, II. 21.10). But these cases tend to be mixed together in the literature with cases involving counterfactual alternative-eliminators. What has not been grasped is that there is some point to distinguishing the two sorts of alternative-eliminators, inasmuch as they fare differently with respect to some of the main objections to Frankfurt's argument.

²⁹If a story is needed to explain the remarkable harmony between the wheels' motions, imagine that Black is equipped with a device that can "see" the future movements of Jones's

steering wheel with 99% accuracy. Data from the device show a sequence of movements culminating in the car's colliding with Smith, but in order to guard against the 1% chance of error Black programs the same movements into his own steering wheel. The information from the device turns out to be accurate, however, and the two steering wheels move in perfect unison. The best story, though, since it does not depend on a retrocausal device, is the one where the synchronicity between the wheels is simply coincidental. The fact that it's a million-to-one--or even a trillion-to-one--shot is irrelevant, since all that's needed is a single counterinstance to PAP.

³⁰David P. Hunt, "Moral Responsibility and Unavoidable Action," *op. cit.* Another strategy with obvious affinities to "blockage" is laid out by Alfred R. Mele and David Robb in their "Rescuing Frankfurt-Style Cases," *Philosophical Review* 107 (January 1998), pp. 97-112.

³¹Fischer uses this term when speaking of "Hunt's approach" on pp. 114-115 and 119-120 of his "Recent Work on Moral Responsibility," *Ethics* 110 (October 1999), pp. 93-139. Others who have picked up the term include Derk Pereboom on pp. 125-28 of his "Alternative Possibilities and Causal Histories," *Philosophical Perspectives* 14: Action and Freedom (2000), pp. 119-137.

³²Alison McIntyre, for example, cites the case of the Princess who stood in her box at the opera for the customary one minute, then continued standing an additional three minutes because she enjoyed the adulation of her subjects, not realizing that a scientist employed by a tabloid had generated a force field around her which would have kept her standing if she had tried to sit down; the force field was then removed, as it happens, just seconds before she sat down on her own (the tabloid's photographers having just then finished snapping their photos). McIntyre claims that this case raises a problem for the incompatibilist, since it seems pretty clear that the Princess is acting as a free agent and is morally responsible for her breach of etiquette, even though it is causally determined that she will remain standing for the extra three minutes. McIntyre's example, however, leaves too many alternatives (especially mental ones) for a savvy libertarian to be flummoxed by it, given the availability of PAP3. See her "Compatibilists Could Have Done Otherwise: Responsibility and Negative Agency," *Philosophical Review* 103 (July 1994), pp. 453-88. McIntyre discusses the Princess case on pp. 474-78.

³³"Alternate Possibilities and Moral Responsibility," pp. 836-7.

³⁴I discuss cases like Wanda's and Stephen's in my "Divine Providence and Simple Foreknowledge," *op. cit.*

³⁵Aquinas is in apparent disagreement with this claim, maintaining that "God's knowledge is the cause of things" (*Summa Theologica* Ia.q14.a8). But insofar as this includes knowledge of future contingencies, particularly exhibitions of free choice, it raises the problem of *predestination*, which is best kept logically separate from the problem of divine foreknowledge.

³⁶This is essentially Augustine's analysis of the relation between divine foreknowledge and human choice, as I argue in my "Augustine on Theological Fatalism," *op. cit.* Key passages for my interpretation of Augustine include "God foreknows all the things of which He Himself is the cause, and yet He is not the cause of all that He foreknows" (*On Free Choice of the Will* III.4) and "A man does not sin *because* God foreknew that he would sin" (*The City of God* V.10). Cf. John Fischer in *The Metaphysics of Free Will*: "On one standard view of the nature of God, God's belief is not a part of the mechanism issuing in (say) my action. On this view, His belief is not what causes my action; rather, my action explains His belief" (p. 180).

³⁷Actually, it's not at all evident to me--I'm making a concession to the objector at this point. I suspect that BSJ's "practical advantage" over GSJ has little to do with the examples themselves and much to do with the fact that BSJ was first *presented* as a challenge to PAP whereas stories like GSJ have always been presented as PAP-based challenges to traditional theology.

³⁸Another possibility is that the Frankfurt Objection would be strengthened if GSJ were joined by some counterexample(s) from the BSJ-family, inasmuch as this would enrich the "data base" by including in the evidence against PAP "ordinary" cases of human-imposed unavoidability alongside the theological case. But this is compatible with GSJ having all the virtues I have ascribed to it.

³⁹"Divine Omniscience and Voluntary Action," p. 43.

⁴⁰This peculiarity of divine omniscience is one of the reasons why William Alston, following Aquinas, is led to deny that God's knowledge is carried by beliefs. See his "Does God Have Beliefs?" *Religious Studies* 22 (September/December 1986), pp. 287-306.

⁴¹A shorter version of this paper was read in 1996 at the Central Division Meeting of the American Philosophical Association, where Linda Zagzebski was the respondent. Longer versions were shared with audiences at Talbot School of Theology, Arizona State University, and the University of Notre Dame, where Al Plantinga, Tom Flint, and the 1999-2000 Fellows at the Center for Philosophy of Religion (Dan Howard-Snyder, Frances Howard-Snyder, Ken Perszyk, Franklin Mason, Zongkun Liu and Omar Mirza) critiqued the paper sentence-by-sentence. I am also indebted to John Fischer, Dale Tuggy and Michael Zimmerman for sending me written comments, and to two anonymous referees for *The Philosophical Review* who had valuable suggestions for rewriting and resubmitting (I did the former but not, in the end, the latter). Finally, my special thanks to John Fischer, whose favorable references to the unpublished manuscript have already introduced some of the paper's arguments into public discussion.