Faith and Philosophy: Journal of the Society of Christian Philosophers

Volume 30 | Issue 1

Article 4

1-1-2013

Prophecy, Foreknowledge, and Middle Knowledge

Joseph Corabi

Rebbeca Germino

Follow this and additional works at: <https://place.asburyseminary.edu/faithandphilosophy>

Recommended Citation

Corabi, Joseph and Germino, Rebbeca (2013) "Prophecy, Foreknowledge, and Middle Knowledge," *Faith and Philosophy: Journal of the Society of Christian Philosophers*: Vol. 30 : Iss. 1 , Article 4.

DOI: 10.5840/faithphil20133014

Available at: <https://place.asburyseminary.edu/faithandphilosophy/vol30/iss1/4>

This Article is brought to you for free and open access by the Journals at ePLACE: preserving, learning, and creative exchange. It has been accepted for inclusion in Faith and Philosophy: Journal of the Society of Christian Philosophers by an authorized editor of ePLACE: preserving, learning, and creative exchange.

PROPHECY, FOREKNOWLEDGE, AND MIDDLE KNOWLEDGE

Joseph Corabi and Rebecca Germino

Largely following on the heels of Thomas Flint's book-length defense of Molinism a number of years ago, a debate has emerged about the ability of Molinism to explain God's purported ability to successfully prophesy the occurrence of human free choices, as well as about the merits of other theories of divine providence and foreknowledge in this respect. After introducing the relevant issues, we criticize Alexander Pruss's recent attempt to show that non-Molinist views which countenance only simple foreknowledge fare as well as Molinism in explaining prophecy. We locate two serious problems with Pruss's proposal, and in the process clarify the theoretical costs and benefits of an adequate Molinist account in this sphere.

Introduction

Critiques of Molinist theories of divine providence are not by any means a rarity, of course. A quick survey of the literature shows that the vast majority of criticisms of Molinism are external critiques: attacks on the general philosophical plausibility of the epistemological and (especially) metaphysical tenets that form the central basis for the view. Much more unusual are internal critiques of Molinism: critiques that aim to show that Molinism fails to deliver the goods it promises, even if it is correct in its basic outlines. If sound, though, such criticisms could erode support for Molinism, since most of the reasons for preferring it to the leading alternatives boil down to its purportedly superior ability to accommodate a variety of tenets of orthodox Christianity and other theological desiderata.

Recently, Alexander Pruss has offered such an internal critique—a reason to think that Thomas Flint's paradigmatic Molinist theory fails to account for God's ability to prophesy the occurrence of some free human actions with absolute certainty, an ability that Pruss and others (including Flint) think God has actually exercised on a number of occasions.¹ Pruss's critique, however, is not a mere provincial attack on Flint. (Though Flint is certainly no insignificant player in the debate, and so a successful provincial attack on him is hardly unimportant.) Pruss believes that the

¹See Alexander Pruss, "Prophecy Without Middle Knowledge," *Faith and Philosophy* 24 (2007), 433–457, responding to Thomas Flint, *Divine Providence: The Molinist Account* (Ithaca and London: Cornell University Press, 1998), especially 197–211.

Molinist can remedy the situation, but in order to do so she must help herself to some crucial assumptions which are equally open to those who deny the existence of Middle Knowledge and believe in a God with only foreknowledge. Hence, according to him, Molinism enjoys no theological advantage over these alternative theories, at least where prophecy of free human actions is concerned.

Pruss is mistaken about the general prospects for Molinism and foreknowledge-only views where prophecy is concerned, however, and this paper will show why. We will begin by outlining some basic assumptions and summarizing Pruss's objection to Flint's account. (We will assume throughout that he has identified a problem for Flint's own view.) Next, we will examine Pruss's argument for the possibility of prophecy on foreknowledge-only views. We then argue that Pruss's picture of prophecy on foreknowledge-only views is probably incoherent, and even if it weren't incoherent, we would be completely unjustified in accepting the key relevant similarity principles the view would have to employ. Following this, we examine the parallel amendments to Molinist views that are required to handle the relevant prophecy cases. We then show that the revised Molinism is able to give a coherent account of prophecy, and the prospects for justifying its parallel relevant similarity principles are much better (so long as the basic outlines of Molinism are correct). As a result, Pruss scores a shallow victory against Flint; although he may show that the letter of Flint's theory is incorrect, a view in the same spirit can be constructed that preserves the advertised theological advantages of Molinism in the prophecy sphere. The upshot of this exploration is that we gain a better appreciation for both the theological advantages and theoretical costs of an adequate Molinism, vis-à-vis a theory that accepts only foreknowledge.

Basic Assumptions

As we alluded to above, when we speak of prophecy in this paper, we will be concerned exclusively with prophecies of free human actions *qua* free actions. Moreover, the paradigmatic cases of prophecy we will be addressing are ones where the individual whose free action is being prophesied is informed that the prophecy has been made prior to the time of the prophesied action. (Flint and Pruss focus on the famous case of Jesus predicting Peter's denial of him as their main example, and we will follow them in this convention.²)

We also suppose that God must be 100 percent certain that the prophecy will come to pass in order to issue it (or inspire someone else to issue it). God cannot merely have 99 percent confidence that (e.g.) Peter will deny Christ in order for God (in the form of Christ) to issue the prophecy,

²Flint seems to focus primarily on the version in Luke 22, while Pruss focuses on the version from Matthew 26, but these differences are of no significance. In both cases similar prophecies of free actions are made. Incidentally, we also follow Pruss in simplifying matters by assuming that when Jesus prophesies, God prophesies.

with the understanding that, if Peter does not freely deny Christ, God will simply compel him to deny Christ. For one thing, this would solve the problem in a way that is too simplistic to be interesting—by, in essence, claiming that God doesn't actually issue pure prophecies of free actions *qua* free actions. (If this is the truth of the matter, then the literature on the topic is moot anyway.) And perhaps more troublesome, forcing Peter into denial may be contrary to God's goodness. In addition, it could turn out to be impossible on its own terms, since the idea of an unfree denial may be an incoherent notion.³

Following Flint and Pruss, we understand Molinism to be the thesis that God non-trivially knows the truth values of all contingent subjunctive conditionals (what we will call "F-conditionals") of the form:

- (1) x is in circumstances $C \Box \rightarrow x$ freely chooses to do A in C

Here x is a created agent, C are circumstances which are described in maximal detail (with the caveat that nothing in the specification of C entail either the consequent or the negation of the consequent—otherwise the knowledge requires only a priori inference), and A is some action.⁴ We will take ' $\Box \rightarrow$ ' to be the symbol for subjunctive or counterfactual implication, and the non-triviality qualification is meant to ensure that cases where God knows the truth values of these conditionals only because every such conditional is false (because it is up in the air what a free creature would do if placed in C , as C is non-actual) do not count as a form of Molinism. Following standard practice, we will refer to any knowledge God has of the truth values of such conditionals as "Middle Knowledge."⁵

We will also assume a standard semantics for assessing the truth values of subjunctive conditionals. While the details of the story are notoriously controversial, all that will matter for our purposes is the big picture, and opinion on that is fairly settled. Any subjunctive conditional of the form $M \Box \rightarrow N$ will be true if and only if there exists a world where M and N are both true that is closer to the actual world (i.e., roughly speaking more similar to the actual world, up to the time that M occurs) than any world where M is true and N is false.⁶

Finally, we will not entertain any external critiques of Molinism—critiques designed to call into question the fundamental metaphysical or

³For similar points, see Pruss, "Prophecy Without Middle Knowledge," 436, and Flint, *Divine Providence*, 201. Flint makes only the point about denial being essentially free.

⁴For a more detailed discussion of the appropriate way to formulate these claims, see Dean Zimmerman, "Gale and the Free Will Defense," *Philo* 6 (2003), 78–113.

⁵For a similar setup, see "Prophecy Without Middle Knowledge." Incidentally, the need for the non-triviality qualification was first pointed out by Robert Adams, "Middle Knowledge and the Problem of Evil," *American Philosophical Quarterly* 14 (1977), 109–117.

⁶See Flint, *Divine Providence*, 7. Although of less central importance, we will also assume an analogous approach for representing the truth of F-conditionals from the perspective of God's creative decision making. This will be discussed in greater detail below.

epistemological underpinnings of the Molinist program.⁷ Although these criticisms may have considerable merit, the purpose of this paper is to assess whether Molinism has a theological advantage over foreknowledge-only alternatives where prophecy is concerned, not to assess the global prospects for Molinism. Consequently, we will at no point challenge the claim that God has Middle Knowledge or any other central tenets of Molinism.⁸

Pruss's Argument Against Flint: A Summary

In a nutshell, Pruss's issue with Flint's account is that Flint fails to explain adequately how God could use Middle Knowledge to secure absolutely certain knowledge of what Peter would freely do under the relevant circumstances, of the sort required for God to issue the prophecy.

With Middle Knowledge, God could easily know the truth values of F-conditionals like the following, where C is a fully specified situation where God does not issue a prophecy about whether Peter will deny Christ:

(2) Peter is in C $\square \rightarrow$ Peter denies Christ

God can even unproblematically know the truth values of such F-conditionals when C involves Peter's *belief* that God has prophesied his denial, so long as the belief is not actually true. But of course what God really needs is substantive knowledge about how Peter would act if he believed *truly* that God had issued a prophecy. Why do things become problematic if God has in fact issued the prophecy? Because then a counterfactual of (2)'s form is true only because:

(3) Peter is in C $\Rightarrow$ Peter denies Christ⁹

(3) holds because C includes God's prophecy that Peter will deny Christ, it is part of the concept of God that he is essentially omniscient and truthful, and so it is a conceptual truth that if God prophesies something, it will come to pass. But why, we might wonder, is this such a problem? Why can't God just use his knowledge of this conceptual truth in order to issue the prophecy in question? Certainly, this knowledge is not Middle Knowledge, but it is knowledge that everyone agrees God would have.

The trouble is that (3) is tautologous, and therefore provides no information that is useful to God in deciding whether to issue a prophecy. After all, God also knows:

⁷Prominent examples of such critiques include a number of arguments in William Hasker, *God, Time, and Knowledge* (Ithaca: Cornell University Press, 1989), and the central attack in Robert M. Adams's "An Anti-Molinist Argument," in *Philosophical Perspectives* 5 (1991), 343–353.

⁸As Flint points out (*Divine Providence*, 42–43), Middle Knowledge is more inclusive than just knowledge of F-conditionals. Nothing we say here should be taken to imply otherwise; in any case, none of these issues should have any bearing in this context.

⁹The ' $\Rightarrow$ ' here is the symbol for entailment.

- (4) Peter is in C^* (where C^* involves God prophesying that Peter will not deny Christ) $\Rightarrow$ Peter does not deny Christ.

(3) and (4) together make it clear that God needs to know something more substantive in order to responsibly decide whether to prophesy Peter's denial. So far, so good—Flint himself is aware of this problem.¹⁰ The trouble, as Pruss sees it, is in how Flint attempts to deal with the issue.

Flint proposes a solution.¹¹ In essence, he wants us to consider all F-conditionals of the form:

- (5) Peter is in C^{**} and Peter believes at t that God prophesies that Peter will deny Christ because of reason R $\square \rightarrow$ Peter denies Christ.

Here C^{**} are general background conditions that describe the circumstances of Peter's decision to deny Christ in fully determinate detail, minus a specification of whether Peter believes at t (the time of the Last Supper, several hours before the denial decisions) that God has prophesied his denial and minus anything that would entail such a specification.¹² R is some reason (aside from God's actually prophesying the denial) why Peter believes God has so prophesied; it could be a complicated hallucination, Peter mishearing Jesus, etc.

If all such F-conditionals have the same truth value, Flint refers to them as "harmonious." If they are harmonious because they are all true, then Flint claims that God can conjoin them with:

- (6) Peter is in C^{**} and Peter believes at t that God prophesies that Peter will deny Christ because God really does prophesy that Peter will deny Christ $\Rightarrow$ Peter denies Christ.

Then, by what Flint takes to be a plausible principle of counterfactual inference, God can conclude that no matter how Peter comes to believe that God prophesies his denial, Peter will deny Christ.¹³ So, God can go ahead and issue the prophecy whenever he is lucky enough that the F-conditionals are harmonious in this way. (Flint thinks this sort of harmony would always be a contingent matter.)¹⁴

¹⁰See Flint, *Divine Providence*, 200.

¹¹We follow Pruss in ignoring what Flint calls his "First Solution" (which is just that the prophecy is not a prophecy of a free action *qua* free). Pruss splits Flint's "Second Solution" into two parts, but the first part leads directly into the second, which ultimately captures the view Flint endorses.

¹²In both the Matthew and Luke accounts (as well as the account in Mark 14), Peter is represented as having forgotten the prediction at the time of the denial decisions, remembering it only after the crowing of the rooster has concluded. This is potentially significant, since it suggests that Jesus's prediction was not within Peter's awareness at the time of his decisions.

¹³The basic principle is that, for any true counterfactuals of the form $G \square \rightarrow I$ and $H \square \rightarrow I$, $(G \vee H) \square \rightarrow I$ will also be true. This can be iterated as many times as needed to accommodate all the possible antecedents, recalling that when an antecedent entails a consequent, it also counterfactually implies that consequent.

¹⁴Flint, *Divine Providence*, 207–208.

Pruss counters by contending that the bells and whistles provided by this account do nothing but obscure the force of the original problem; they do nothing to solve it. The original difficulty, of course, was that God lacked the substantive grounds needed to issue the prophecy with 100 percent certainty. Although knowing that in a wide array of other circumstances Peter denies Christ (many of them qualitatively indistinguishable from a scenario where God really does prophesy) may give God considerable *evidence* that Peter would deny Christ if a genuine prophecy is made, it does not give God absolute certainty that issuing a genuine prophecy will not somehow alter Peter's free decision from what it would be in one of the non-genuine cases. Hence, God still lacks the substantive information he needs, and this hole cannot be filled by using the tautologous (and hence insubstantial) (6).¹⁵

*Pruss's Solution, and How It Applies to
Foreknowledge-Only Accounts: A Summary*

Now that we have seen the basic issue for Flint, we are in a position to examine Pruss's solution, and how Pruss applies it to foreknowledge-only accounts. Pruss proposes that a relevant similarity principle may hold, of the following form:

(SP) If $C \ R \ C_1$ and if x does A in C , then x would have still done A in C_1

SP here is but a schema. To be made into a full-fledged principle, the R relation must be specified. Beyond claiming that it would need to be an equivalence relation, though, Pruss takes no specific stand on what R could be. Instead, he considers a range of possibilities—everything between complete subjective indistinguishability (across times) for the agent x on one extreme to rough sameness of available reasons for action (at the time of the decision) on the other extreme.¹⁶ But as long as some R relation makes the principle true, this might be enough for a God with Middle Knowledge to identify situations where He could safely issue prophecies.

¹⁵It is important to Pruss's case that Flint be interpreted as claiming that the F-conditionals for all the instances where Peter believes that God has prophesied his denial (but where God has not in fact so prophesied) are harmonious simply *by accident*, as a mere coincidence. This is what motivates Pruss's upcoming solution involving a relevant similarity principle; it is intended to provide grounds for extending (with certainty) the lessons from the non-prophecy F-conditionals to the genuine prophecy ones. If this is an incorrect reading of Flint, then Flint can simply fall back on such a principle to rescue his view. Fortunately, the interest of neither Pruss's own discussion nor our argument hinges on this interpretive issue, because the central matter is not the success of Flint's specific account (which everyone seems to agree can be tidied up if necessary), but rather the adequacy of Pruss's non-Molinist view. Our claim is that there is good reason to suppose that this view is incoherent, and even if it isn't, it provides no theoretical advantages over Molinism in explaining prophecy; in fact, it actually suffers from serious drawbacks. (We are grateful to Flint for pointing out to us in correspondence the potential issues with Pruss's interpretation here.)

¹⁶"Prophecy Without Middle Knowledge," 450–452. Here "subjective indistinguishability" should also be understood to involve sameness of underlying character as well.

The Molinist could then plausibly conjecture that the Peter prophecy is such a scenario.

To be helpful to the Molinist, Pruss also thinks that any such principle would have to be a necessary truth (or at least more than a mere coincidence, as Flint seems to suggest F-conditional harmony would be), because it would need to be in virtue of some general feature of the circumstances that the principle held. (Pruss acknowledges that putting one's finger on the general feature is difficult since we are dealing with circumstances involving libertarian free will—rather than conceptual or nomic necessitation—but he doesn't find the idea that there could be such general features incoherent.) In a situation where God was contemplating prophesying that x would do A , he could only make use of this principle (to justify his belief that prophesying would not alter x 's doing A) if it were true in virtue of some general feature of circumstances R -related to C , along with his recognition that the circumstances associated with a genuine prophecy would be R -related to C . Otherwise—if the principle were merely accidentally true for C and the C_n not involving a real prophecy—there would be no guarantee that just because it held in a variety of other circumstances, it would continue to hold in the crucial case of genuine prophecy.

Pruss thinks that as long as the Molinist is willing to live with this baggage, then she can hold that God is able to use his Middle Knowledge to issue at least some prophecies of free actions. However, Pruss also thinks that those who reject Molinism and accept only foreknowledge can then enjoy basically the same benefits. If the foreknowledge-only proponent accepts such a relevant similarity principle, he too can make sense of prophecy, thus vitiating a theological advantage of belief in Middle Knowledge. According to Pruss, the explanatory account for the Petrine prophecy using only foreknowledge would go as follows (in explanatory order from first to last):

- (i) God decides that at t_1 Peter will be in C^+ either due to a hallucinatory or real prophecy about denying Christ. (Here C^+ are the circumstances Peter finds himself in at the Last Supper as represented in the relevant scriptural account, minus a specification of why Peter believes God has prophesied his denial.)
- (ii) Peter at t_1 is in C^+ due to receiving a hallucinatory or real prophecy about denying Christ.
- (iii) Peter denies Christ at t_1 .
- (iv) God knows (ii), (iii), and some relevant similarity principle where the circumstances of Peter believing God has prophesied his denial because of a hallucination are R -related to the circumstances of Peter believing this because of a genuine prophecy.

- (v) God brings it about that at t_0 (the time of the exchange at the Last Supper) Peter receives a real prophecy about denying Christ.¹⁷

*Why Pruss's Account of Prophecy on
Foreknowledge-Only Views Is Problematic*

If plausible, Pruss's account of prophecy on foreknowledge-only views would be a blow to the Molinist case, since as we alluded to earlier a major motivation for accepting Molinism is its alleged ability to give a theoretical underpinning to what many take to be central claims of Christianity (such as that God can issue prophecies about free actions) where other views fail to provide that underpinning. Alas, the account suffers from serious problems. At least as it is articulated here, Pruss's foreknowledge-only theory of prophecy requires God to bracket that foreknowledge in a way that would be impossible for an omniscient being (or worse, posits that God remains wholly ignorant of some facts).

The crucial assumption we must make to appreciate this difficulty—an eminently plausible one—is that, necessarily, disjunctions (if true) are always true because one of their disjuncts is true (or both). This claim is entailed by the common-sense view of disjunction, as well as by every semantic theory we are aware of. Hence, whenever God knows that a disjunction is true (like the disjunction in (ii) above), God in addition knows which disjunct makes it true.¹⁸ But now we can give the following argument:

- (A) Either God knows that (ψ) "Peter is in C^+ because he has received a hallucinatory or real prophecy" because God knows that (θ) "Peter is in C^+ because he has received a hallucinatory prophecy" or because God knows that (ϕ) "Peter is in C^+ because he has received a real prophecy."¹⁹
- (B) If God knows (ψ) because God knows (ϕ), then God has made a prophecy at some point previously in the explanatory sequence that Peter will deny Christ. (God's knowledge that a certain state of affairs holds—where that state of affairs involves God's action—is posterior in the order of explanation to God's action.)
- (C) If God has made a prophecy at some point previously in the explanatory sequence that Peter will deny Christ, then it is NOT the

¹⁷"Prophecy Without Middle Knowledge," 448.

¹⁸Ordinary persons can know disjunctions without knowing which disjunct makes them true, but the same does not hold for God. Bill, for instance, could know that \$100 has been stolen from his office, and he could know that only Bob and Mike had access to the office. Hence, he could know that either Bob stole the money or Mike stole the money, without knowing which. But plainly the same doesn't go for God.

¹⁹We can suppose that God knows the relevant propositions either timelessly or from the earliest instant of time, in keeping with a picture where it is assumed that God has foreknowledge.

case that part of God's justification for His decision to prophesy that Peter will deny Christ is that God knows (ψ). (This would violate prohibitions on the circularity of explanation; it can't be in this context both that God's prophecy explains his knowledge of a claim that justifies his prophesying and also that his knowledge of the claim that justifies his prophesying explains his prophesying.)

- (D) If God knows (ψ) because God knows (θ), then God does not prophesy that Peter will deny Christ. ((θ) involves no prophecy taking place.)
- (E) If God does not prophesy that Peter will deny Christ, then it is NOT the case that part of God's justification for His decision to prophesy that Peter will deny Christ is that God knows (ψ). (This is obvious, since there is no decision to prophesy in this case, and hence no justification for any decision to prophesy.)

But then on either horn of the dilemma, we get:

- (F) It is NOT the case that part of God's justification for His decision to prophesy that Peter will deny Christ is that God knows (ψ).

Where does this leave us? It makes it clear that the explanatory sequence of Pruss's foreknowledge-only account of prophecy (as stated) can't be correct, since (ψ) is equivalent to (ii) in that sequence, and God's knowledge of (ii) played an explanatory role (by justifying his decision to prophesy), contrary to the conclusion of the argument we have just seen. But without this explanatory sequence, the account falls apart. If, on the one hand, God simply foreknows a scenario where there is no prophecy, then obviously God can't use this information to issue a prophecy! (Even God can't erase reality, whatever exactly that would mean.)²⁰ If, on the other, God simply foreknows a prophecy, then we are back where we started—wondering what justified God in issuing a prophecy in the first place.

It might be objected that our argument has misrepresented the situation in which God finds himself. Basically, we have assumed that the only way God could have knowledge of (ψ) is either for God to have knowledge of (θ) or for God to have knowledge of (φ). But there is another way for

²⁰As Pruss himself admits, in a different context: "It is incoherent to suppose that time is rolled back *and something else happens*." ("Prophecy Without Middle Knowledge," 455, emphasis in original). Pruss presumably has the idea that God can leave the temporally earlier parts of the causal story that results in Peter being in C+ open, then somehow fill them in at a later logical moment. But if He foreknows that Peter is in C+, this implies that the causal story of Peter's being in C+ is already complete. Otherwise, there would be no state of Peter being in C+ to foresee, since the pieces would not be in place to have brought about Peter's being in C+. (There can't be an event that plays a role in the causal story of Peter's being in C+ which has no specific cause at an initial logical moment, and which God then returns to at a later logical moment and decides will go from having no specific cause to having a specific cause—e.g., Peter's experience of seeming to hear Jesus predict his denial cannot have no specific cause at an initial logical moment and then come to be caused by Jesus genuinely predicting his denial at a later logical moment. This amounts to just the kind of rewriting of history that is unacceptable, and which Pruss himself rejects.)

God to know (ψ) according to this objection: God can simply *decide* that it will be true, as (i) states, without deciding initially which disjunct will make it true.²¹

While potentially seductive, this objection fails. Recall that Pruss is not proposing a Molinist story, where God has the benefit of knowing various F-conditionals in a moment explanatorily prior to his creation of Peter. Pruss's non-Molinist God can know that Peter will deny God in C^+ (or in other relevantly similar circumstances) only once Peter has actually denied God in C^+ .²² (To reject this would be to give God Middle Knowledge.) But, if Peter actually denies God in C^+ , then Peter's being in C^+ itself already needs to have an explanation (i.e., its explanation must have come in explanatory moments prior to Peter's being in C^+). But then we are faced with the same sort of dilemma: either Peter is in C^+ because God prophesied his denial or not. (If not, then presumably Peter is in C^+ because he misheard or hallucinated something.) If Peter is in those circumstances because God prophesied his denial, then God prophesied without knowing what Peter would do. If Peter is in the circumstances for some other reason, then God did not prophesy Peter's denial (or at least God's prophecy has nothing to do with Peter's belief that God so prophesied). Neither option is palatable, since the scenario calls both for God to prophesy with certainty that Peter would deny Jesus and for God's prophecy to be part of the explanation of Peter's being in the circumstances he is in.

We resist claiming that Pruss's account is incoherent because Pruss offers the disjunctive approach above only as an example of how prophecy of free actions might be possible given only foreknowledge.²³ However, aside from some sketchy suggestions at the end of his paper, he offers us no other examples or accounts that help us to understand how what he is proposing might be coherent after all. We rest content at having shifted the burden of proof strongly back onto his side and conclude that his view as presently defended is inadequate.

A Further Problem For Pruss's Account

There is a further problem for Pruss's foreknowledge-only account. In fact, it is in some ways a deeper problem, because it prompts us to appreciate not just the inadequacy of a foreknowledge-only account of prophecy (which has already been established), but also the theoretical costs of an adequate and potentially plausible Molinist view.

This additional problem surrounds the justification for supposing that a relevant similarity principle of the form (SP) above is correct if a

²¹We are grateful to an anonymous reviewer from this journal for pressing us to clarify our argument by answering this objection.

²²"Once" here is used in an explanatory sense, not a temporal one. To take the temporal implications too seriously here would be to threaten God's foreknowledge.

²³We thank an anonymous reviewer for pointing out the need for this clarification.

foreknowledge-only account is true. We will argue that such a principle is much easier to justify if Molinism is true than if a foreknowledge-only account is true. Before starting out, though, it will be important to provide a bit of background that will allow us to think more precisely about the relevant issues.

The most important background items to introduce are the related concepts of a “creaturely world-type” and a “galaxy” (of possible worlds).²⁴ Intuitively, a creaturely world-type is a maximal set of F-conditionals. While the exact formal specification of the notion remains somewhat controversial, Flint suggests the following characterization, which is satisfactory for our purposes:

(WT) T is a creaturely world-type if and only if for any F-conditional ($C \Box \rightarrow A$), either ($C \Box \rightarrow A$) is a member of T or there exists a proposition A^* such that ($C \Box \rightarrow A^*$) is an F-conditional and ($C \Box \rightarrow A^*$) is a member of T.²⁵

A galaxy is a set of possible worlds where the same creaturely world-type is true in each world. (In other words, the same set of F-conditionals is true at each of the worlds. So, e.g., all the worlds in a particular galaxy will agree with one another about what Curley would freely do if he were offered a \$5,000 bribe in circumstances C.) But, assuming there are non-trivially true F-conditionals, the true creaturely world-type is only contingently true, since what makes a particular F-conditional true is the libertarian free choice of a (possible or actual) being. Clearly that choice *could have been* different and, had it been, it would have changed which creaturely world-type was true. But nevertheless, the only feasible worlds for God to create will be worlds within the galaxy whose creaturely world-type corresponds to the creaturely world-type which is in fact true.²⁶

At this point, it will be helpful to develop a heuristic for thinking about God’s creative decisions according to Molinism and foreknowledge-only views respectively. This will set the stage for us to lay bare the difficulties for foreknowledge-only views and show the costs of an adequate Molinism.

²⁴These notions will be familiar to some from the literature, and from Flint, *Divine Providence*, 47–54.

²⁵Flint, *Divine Providence*, 50. Previously we introduced the consequents of conditionals of this form as actions, and here they are being treated as propositions (albeit propositions that describe actions). We will not be overly fussy about these formulations, as no issues of substance hang on them for our purposes.

²⁶We use the term “feasible” here in the way that Flint uses it in *Divine Providence*. A world is feasible for God to create iff God is able to actualize it. One might object: can’t God actualize any possible world? Isn’t that basically what it means for a world to be (metaphysically) *possible*? The reply, of course, is that some worlds involve creatures freely performing particular actions, and there is no way for God to compel creatures to perform these actions. E.g., there is a possible world where Curley freely accepts a \$5,000 bribe when placed in circumstances C. But if the true creaturely world-type states that Curley *would not* in fact freely accept such a bribe if placed in C, God will be powerless to actualize the world in question where Curley freely accepts the bribe.

We begin by looking at Molinism.²⁷ An important aspect of F-conditionals to appreciate is that, in the explanatory moments when God is contemplating what world to actualize, there is (*ex hypothesi*) no actual world yet. Hence, we cannot represent the truth of an F-conditional like $(C \Box \rightarrow A)$ at this stage by claiming that the closest C world (to the actual world) is an A world rather than a $\sim A$ world. Rather, what we need is what we will refer to as a “complete world diagram” (CWD). A CWD is a set with two layers of structure (with other special features); in other words, it is a set of sets. Each CWD is designed to represent all of the modal information encapsulated in any possible creative choice God could make. It is analogous to a possible world, but represents possible scenarios faced by God in making creative choices.

We construct any given CWD by first taking sets of all the possible worlds grouped into galaxies. (So far, all the sets will appear identical, since the worlds and the galaxy each world belongs to is constant.) At this point, we have a representation of all the possible worlds and their closeness relations, as captured by the galactic relationships.²⁸ Next, we take one of these sets and mark one galaxy as having the true creaturely world-type, and hence as containing the only feasible worlds for God to actualize; one can think of this galaxy as “glowing” if one likes. Then, one takes all the sets like this (each with a different galaxy glowing) and forms a second-order set composed of them. Now, we have a representation not just of the possible worlds and their closeness relations, but also a representation of what possible creaturely world-types could be true. (The contingency of creaturely world-types is captured at this second level, where each set represents a different creaturely world-type as being true—a different galaxy glowing.) Next, we take each first-order set and mark one of the worlds in its glowing galaxy (as well as a single world in each of the non-glowing galaxies). The world we mark in the glowing galaxy represents the world God would actualize if that particular first-order set represented the creative situation that faces God (i.e., accurately represents what creaturely world-type is feasible for God—what galaxy is glowing). The world in each non-glowing galaxy represents what God

²⁷Because of spatial constraints, we will not attempt to discuss God’s creative choices in detail, and how these choices would be explanatorily structured if Molinism were true. This is a well-studied phenomenon in the literature, and the details are not important for our purposes.

²⁸Throughout, we will be assuming that within any given galaxy G , each member of G is closer to every other member of G than it is to any world outside G . In fact, we will go slightly further than this, assuming that for every world w in G and every initial segment in w , up through the end of that initial segment w is closer to every other world in G (up through its parallel initial segment) than it is to any world (up through its parallel initial segment) outside G . Both the principles are plausible, and in fact Flint conjectures that the principle for entire worlds is true (*Divine Providence*, 52–53). These assumptions may be dispensable, but it would make the subsequent arguments considerably more complicated. Henceforth, when we speak of the closeness of worlds, we will be supposing that it is the closeness of their initial segments up through the relevant time that is at issue. (Sometimes we will continue to emphasize this if there is the potential for confusion.)

would have actualized in that creative choice scenario had that particular galaxy been glowing instead.²⁹ Finally, to complete our CWD, we mark one of the first-order sets as accurately representing which galaxy is glowing—which creaturely world-type is in fact true. This gives us a representation of God’s specific creative choice, along with a representation of the possible worlds, their closeness relations, and what creaturely world-types could have been true (as well as which one in fact is).³⁰

To sum up briefly, a construction of a single CWD occurs in several stages:³¹

- (1) Form a structured set of all the PWs grouped into galaxies $g_1 \dots g_m$.
- (2) Mark one of the galaxies as glowing and call the resulting set “ s_1 ”.
- (3) Repeat steps (1) and (2) indefinitely, creating new sets $s_2 \dots s_n$ until each galaxy has been marked as glowing in one of the sets.
- (4) Form a set S of the sets $s_1 \dots s_n$.
- (5) In each of the $s_1 \dots s_n$, mark one of the worlds in each $g_1 \dots g_m$ as the world God would actualize in that g_x if that g_x captured the true creaturely world-type.
- (6) Mark one of the $s_1 \dots s_n$ as correctly representing the true creaturely world-type.

We can now picture God’s creative decision according to Molinism as follows: imagine each CWD as a point in a space, similar to a possible world in the space of possible worlds. At the center of this space is God, at the explanatory moment where he is deciding what world to actualize. If Molinism is true, then God has Middle Knowledge, and hence knows which creaturely world-type is true. Thus, we can envision CWDs that represent this world-type being true occupying a nebulous orbit immediately surrounding God, with CWDs that represent other world-types being true occupying orbits substantially further away from God.³² When

²⁹Flint claims that any creative act by God resulting in the actualization of one of the feasible worlds would also include a conditional decision by God about what worlds he would have actualized had other galaxies had the true creaturely world-type instead (*Divine Providence*, 60). This seems right, but the issues raised won’t be relevant for our purposes.

³⁰One might wonder how CWDs represent an improvement over the apparatus introduced by Flint (*Divine Providence*, 51–59). Although Flint’s apparatus is adequate to his task, it is important for our purposes (as will become clear shortly) that we be able to represent God’s relationship to his creative choice possibilities analogously to the way that we represent the closeness of possible worlds to one another. CWDs allow us to do this in a way that Flint’s original machinery does not.

³¹The following summary assumes there are denumerably many galaxies and worlds, but this simplification could be dispensed with if need be.

³²According to most forms of Molinism, God has no power over which creaturely world-type is true. There are variant forms of Molinism that claim otherwise—most notably so-called “Maverick Molinism”—but these views suffer from serious problems. Consequently, we will ignore them.

God makes his creative choice, we can think of him as selecting one of the CWDs from this inner orbit and placing it "right on top of him," as it were, with full awareness of what is going to result. It is at this explanatory moment that he strongly actualizes states of affairs which, together with free creaturely choices (known by God via Middle Knowledge), result in the weak actualization of an entire world.³³ Put all of that together and one CWD becomes completely correct. At this point, all the other CWDs (both in the inner orbit of feasibility and the outer orbits) arrange themselves into fixed positions within their orbital region, based on their closeness to the CWD which has become fully correct. (In this arrangement process, CWDs always remain in the orbital vicinity of the other CWDs that represent the same creaturely world-type being true that they do.)

Things are basically the same for a foreknowledge-only view, with one crucial difference. God does not have Middle Knowledge, and so doesn't know which creaturely world-type is true prior to making his creative decision.³⁴ Consequently, on this approach, when God goes to create, he is blind to which group of CWDs occupies the inner orbit. He is cut-off from this sort of awareness, forced to rest content only in knowing which CWD will wind up correct if any particular creaturely world-type is true (given the states of affairs he plans to strongly actualize). It is only after he strongly actualizes the various strongly actualizable states of affairs that he learns the orbital arrangements of the various CWDs, including which one wound up being correct (and thus which world wound up weakly actualized).

Now, we are in position to fully appreciate the problem for the proponent of a foreknowledge-only view. Recall our relevant similarity principle schema above:

(SP) If $C \ R \ C_1$ and if x does A in C , then x would have still done A in C_1

Recall that the foreknowledge-only proponent hopes that such a principle will justify God in issuing a prophecy (like the prophecy of Peter's denial) by allowing God to extrapolate (with certainty) from a person's

³³Here we ignore issues about the analogues of F-conditionals for indeterministic but unfree events. Incidentally, the terminology of "strong actualization" and "weak actualization" first appeared in the now famous reintroduction of Molinism in Alvin Plantinga, *The Nature of Necessity* (Oxford: Oxford University Press, 1974).

³⁴We assume that even if a foreknowledge-only view is true, there are still non-trivially true F-conditionals of the sort the Molinist believes in. (The central difference in the non-Molinist case, though, is that God lacks (at least) pre-volitional knowledge of these F-conditionals.) While this may be a controversial assumption, Flint points out (*Divine Providence*, 40) that Molina himself basically took for granted that his libertarian opponents accepted it. Moreover (and perhaps more to the point), in order for Pruss's foreknowledge-only view to have any semblance of plausibility, he must assume that there are at least some such F-conditionals, and he openly acknowledges that he is committed to there being some ("Prophecy Without Middle Knowledge," 449). This is really all that is required for our arguments to work, but giving a formulation that takes full account of the technical variations that could arise would make the presentation prohibitively long.

pattern of free behavior in cases where there is not a divine prophecy to a conclusion about how the person would freely behave if there were a divine prophecy.

The scope and nature of the R relation will not matter for our purposes (at least at the moment), so we can simply assume R amounts to whatever the foreknowledge-only proponent thinks is required to make sense of all the important prophecy cases. One thing that does matter for our purposes, though, is whether this principle is supposed to hold purely by accident (i.e., by coincidence), necessarily, or as a result of some general feature of x (which preserves the certainty of x freely doing A in circumstances R-related to C in the actual world's galaxy, but not in all other galaxies). (Pruss, at least, favors the second option, but he is willing to take the third seriously as well.)³⁵ Unfortunately, there are very serious problems for each of these approaches.

The issues are easiest to see for the "by accident" fork. As Pruss himself notes, this answer won't help.³⁶ The trouble is that, if the principle holds by accident for C and C_1-C_n where there is no genuine prophecy, how is God going to use this result to gain confidence (let alone certainty) that, when he genuinely prophesies that Peter will deny Christ, Peter will in fact deny Christ? There is no way it can help him in that case, because there is no guarantee (indeed, no reason to suppose) that the fortuitous coincidence will continue in the event of a real prophecy. So plainly the foreknowledge-only proponent will want to look for other options.³⁷

What about if the principle is true necessarily? This looks quite a bit more promising, at least initially. If it is true necessarily, then presumably God will grasp its necessity and see that he can project its results onto the situation of genuine prophecy. (In fact, he will need to grasp its necessity, because he will need to know that it is true to prophesy with the requisite certainty, and grasping its necessity is presumably the only way someone like God could know a necessary truth of this sort.) But even if we set aside the pesky issue of genuine prophecies, there's trouble brewing here. If such a principle is true necessarily, then let us return to our earlier heuristic to understand better what the situation will look like, focusing on our Petrine paradigm.

In this case, from the perspective of any world where Peter denies Christ in C, it will follow according to the principle that at that world it is

³⁵"Prophecy Without Middle Knowledge," 446–447.

³⁶Ibid., 446.

³⁷It might be tempting to think that God can simply know that circumstances of genuine prophecy would be R-related to C, and hence be covered by the principle. But how could this be so if the principle is true by accident? The suggestion here fails to account for the lessons that came out of contemplating the original problem. God needs to have a special justification for thinking that he can issue the prophecy in the first place, because genuine prophecy is a special kind of case. Lessons about it cannot be easily inferred from non-prophecy situations, and simple *ad hoc* declarations that it is possible are bound to be unsatisfying.

also true that "Peter is in $C_1 \square \rightarrow$ Peter denies Christ."³⁸ Using our preferred semantics for counterfactuals, then, this will imply that the nearest world (to any Peter in C world) where Peter is in C_1 is always a world where Peter denies Christ.³⁹ Similarly (since 'A' is just a variable in the (SP), to be filled in by any action), from the perspective of any world where Peter doesn't deny Christ in C , it will be true that "Peter is in $C_1 \square \rightarrow$ Peter doesn't deny Christ," implying that the nearest world where Peter is in C_1 will be a world where he doesn't deny Jesus. The result will be a partitioning of all the galaxies into two classes—one a class of galaxies whose characteristic world-type claims that if Peter were in C or any of C_1 – C_n , he would deny Christ, and the other a class where the characteristic world-type claims that in any of these circumstances Peter would not deny Christ.

Now we reach a crucial juncture. It is important to recognize that the differences in closeness between worlds here are what we might call "primitive" differences in closeness; in other words, they do not supervene on qualitative similarity or similarity of laws in the way that is familiar from the evaluation of standard nomic counterfactuals like "If the crystal vase had been dropped from the top of the Eiffel Tower, then it would have shattered." To see why, consider a world (w) where Peter is in C and does deny Christ. Now, consider two worlds where Peter is in C_1 ; in one of these worlds, Peter denies Christ (w_1) and, in the other (w_2), he doesn't. (We know that both worlds exist because Peter's decision is free, and so it is possible for him to choose either option.) *Ex hypothesi*, because Peter's decision is free, we can choose these two worlds so that their initial segments (up through the time of C_1) are qualitatively identical to one another; the same qualitative circumstances hold in each, the same laws of nature, etc. And moreover, we can choose them so that they represent the closest C_1 worlds where Peter denies Christ, on the one hand, and doesn't deny him, on the other. If the relevant similarity principle holds, then w will be closer to w_1 than to w_2 , but plainly, since the relevant initial segments of w_1 and w_2 are qualitatively identical to one another, whatever is responsible for the closeness difference between each of them and w could not be a matter of qualitative similarity. And since there is really no other candidate to serve as a basis for the closeness relation, the differences have to be *sui generis*.⁴⁰

³⁸There is another way to interpret the claim that the principle is necessary. If one reads "x does A in C " as "x *actually* does A in C ," and one interprets the term "actual" to rigidly designate a particular world, then the principle will be trivially necessary if true at all, since (on plausible assumptions about modal logic) the principle will be necessarily actually true if it is actually true *simpliciter*. In this case, the necessity option will simply collapse back into one of the contingent options for all intents and purposes. Consequently, we ignore it.

³⁹Note that the sort of F-conditional we have here is a regular counterfactual that can be evaluated using a regular semantics in terms of closeness of worlds. This stands in contrast to the F-conditionals that apply when God is preparing to make his initial creative decision.

⁴⁰Strictly speaking, of course, there is another candidate; we could treat F-conditionals at worlds as having primitive truth values, and then count similarity of counterfactuals in

Thus, on this foreknowledge-only picture, after God engages in his creative decision and strongly actualizes some states of affairs (resulting in the weak actualization of a world), he will be in a position to know primitive closeness relations between worlds, and hence to know the non-trivial truth values of F-conditionals at this world. But if God is able to know these primitive closeness relationships, why deny God knowledge of primitive closeness relationships between his initial starting point and the various orbits of CWDs—in particular, the knowledge of which set of CWDs represent the true creaturely world-type (i.e., the set of CWDs that represent the correct galaxy as glowing)? To do so is totally unprincipled. The main reason to deny this sort of knowledge to God in his initial state of creative deliberation is suspicion about God's ability to know primitive closeness relations (or their very existence)—suspicion about whether it makes sense to speak of God knowing which F-conditionals are true, and hence God being aware of primitive closeness relations of one sort or another. But once we have set aside our objection to primitive closeness (and indeed, to God's knowledge of primitive closeness relationships), the bulwark against full-on Molinism seems to have been removed. For there appears to be no principled reason in that case to maintain that God *can* know primitive closeness relationships between worlds, but *can't* know primitive closeness relationships between CWDs and his creative starting point.

A possible objection here—suggested by some of Pruss's remarks—is that Molinism requires God to know F-conditionals that have no actual truth-makers, while the F-conditionals Pruss thinks God would know do have an actual truth-maker (Peter's actual decision about whether to deny Christ).⁴¹ According to this line of thinking, it is only knowledge of ungrounded F-conditionals that is problematic.

Although initially plausible perhaps, the objection fails. While it is seductive to think of Peter's denial in the actual world as somehow relevant to the truth of an F-conditional about what he would have done in different circumstances, no proposed semantics for any sort of subjunctive conditional allows for this. It is crucial to keep in mind that only similarity up through the end of the initial segment is relevant, and the decision comes after the initial segment. (The antecedent of an F-conditional is defined so that it does not entail the consequent or its negation, or otherwise lay bare on conceptual or nomic grounds whether the consequent holds.)

different worlds as important factors in fixing their closeness to one another. (This appears to be Plantinga's view in *The Nature of Necessity*, 178 and also in other writings.) For our purposes, though, going this route winds up being equivalent to the primitive closeness picture; the only difference between the approaches is where we choose to locate our fundamental primitiveness. The same points will apply to this primitive truth value picture *mutatis mutandis*.

⁴¹See "Prophecy Without Middle Knowledge," 449. Incidentally, the famous anti-Molinist Robert Adams also seems to endorse the truth of these sorts of F-conditionals (and God's knowledge of them). See his "Middle Knowledge and the Problem of Evil," *American Philosophical Quarterly* 14 (1977), 109–117.

The seductiveness here is probably due to our tendency to smuggle assumptions that are appropriate for ordinary nomic counterfactuals into a context where they don't belong. If an unknown object is actually dropped from a height of five feet onto a marble floor and shatters, it is probably safe to assume that if it had been dropped from a height of ten feet onto that floor (or from a height of five feet onto a concrete floor), it still would have shattered. But there is no mystery in this—no primitive closeness relations between worlds. The bottom line is that the initial segment of the nearest 10-foot drop world where the object shatters is going to be closer to the actual world than the initial segment of the nearest 10-foot drop world where it doesn't shatter because the segment in the shattering world is more qualitatively similar to that of the actual world. (The laws in the non-shattering world are likely to be quite a bit different from those of the actual world, for instance, or else the floor is likely to be qualitatively different—padded perhaps.) This fact has nothing whatsoever to do with the outcome—the shattering—in the actual world, except insofar as the outcome in the actual world gives us a clue as to the similarity between nomic processes there and in the other world.

Things are completely disanalogous with F-conditionals, of course, because nomic processes are not going to drive the transition from antecedent to consequent. Any C_1 world where Peter denies Jesus is going to have a parallel C_1 world where Peter doesn't deny Jesus that is perfectly qualitatively identical to it (and identical in its laws) until the moment the decision is made.

So much, then, for believing that the relevant similarity principle holds necessarily (in the context of a foreknowledge-only view). This position tries to find a middle ground between denying that God has knowledge of any primitive closeness relations and believing that God has expansive knowledge of primitive closeness relations, and it fails to do so in a principled fashion. As you'll recall, though, there is one option remaining—claiming that the relevant similarity principle holds contingently, but via a general enough fact about agent x that God can extrapolate with certainty as to x 's behavior if x were given a genuine prophecy.

According to this approach, galaxies will be considerably less "orderly" than they were on the necessity view, because only in a select few galaxies will (e.g.) Peter perform the same action in C , C_1 , etc. Rather than cleanly dividing the galaxies into two classes, a contingent relevant similarity principle would leave a large variety in the sorts of world-type profiles one would find in the respective galaxies. To be sure, thanks to the truth of the principle, the galaxy with the true world-type will still be quite orderly, though. In that galaxy, Peter denies Christ in every circumstance R -related to C , but in other galaxies, things will be quite a bit more chaotic. (What it is for the principle to hold contingently is for the true world-type to conform to the principle, but for other world-types—world-types which could have been true—to fail to conform.)

While there is no obvious incoherence in the idea of this sort of arrangement, alas, there are once again serious problems (again, in the context of a foreknowledge-only view). The main problem is essentially the same one we saw above for the necessity fork—the unprincipled attempt to walk the fine line between completely denying God’s knowledge of primitive similarity relationships and embracing widespread divine knowledge of these things. But the problem is made even more stark with a contingent principle. A contingent principle would have the consequence that there will be distinct qualitatively identical worlds (qualitatively identical at all times) that differ from one another only in their primitive closeness relationships with other isomorphic worlds (and hence, in what F-conditionals are true at them).

For example, there will be a world (call it w_3) where Peter mishears Jesus and believes that Jesus prophesied his denial, and where Peter goes on to deny Jesus. There will be another world (call it w_4) where things unfold (before, during, and after Peter’s mishearing) exactly as they do in w_3 . But w_3 will be in a different galaxy from w_4 ; w_3 is in the glowing galaxy, so (e.g.) the F-conditional “Peter hallucinates Jesus prophesying (in the relevant circumstances) $\square \rightarrow$ Peter denies Christ” is true in w_3 , while it so happens that this same F-conditional is false in w_4 . But this entails that the closest hallucination world to w_3 is one where Peter denies Christ, while the closest hallucination world to w_4 is one where Peter doesn’t deny Christ (or perhaps there is a tie). But then suppose Peter does mishear Jesus and believes Jesus has prophesied his denial (and in addition all else goes as w_3 and w_4 state); in this case, how will God even know *which world* has been actualized?⁴² Both w_3 and w_4 are qualitatively alike in all respects. The only possible answer, it seems, will be to appeal once again to God’s mysterious ability to grasp primitive closeness relationships. God will have to sense how this world is related to other worlds—and hence be aware of what counterfactuals are true at it—in order to zero in on which world it is. And, just as we saw before, there is no principled way to take this step and refuse to take the further step of embracing Molinism’s more expansive affirmation of God’s ability to know closeness relationships.

So, to sum up, all of the three forks open to the foreknowledge-only proponent in answering the question of the modal strength of the relevant similarity principle (accidental, necessary, contingent in virtue of something general) lead to dead ends. Consequently, we have yet another

⁴²One might object that God did not, in fact, allow Peter to mishear Jesus in this way, and so the issue doesn’t arise. It is true that God didn’t allow Peter to mishear in this way, although there are presumably many cases where individuals actually performed free actions in particular circumstances. The problem arises in all such cases where the relevant similarity principle is supposed to apply; cases related to prophecy are not required to make the point. So, e.g., if Curley freely accepts a \$5,000 bribe from Fat Tony on a particular street corner on a specific day, then he would have accepted the same bribe from Fat Tony’s identical twin (whom he is unaware is not Fat Tony) had it been offered at that same place on that same day. But in a different galaxy, the same pattern will not play out.

reason to reject the claim that foreknowledge-only cases can handle prophecies using a relevant similarity principle.

Can Molinism Do Any Better?

Molinism, by accepting an extreme (but principled) view of God's ability to know primitive closeness relationships, is able to get around the concerns that plagued foreknowledge-only views in the last section. Since the Molinist is happy to admit that God is aware both of the primitive closeness of sets of CWDs to his creative starting point and primitive closeness relationships between worlds, there is no issue of opportunistic shiftiness on the question of God's abilities in this sphere. And the Molinist can also avoid the earlier problems associated with knowledge of disjunctions that plagued the foreknowledge-only proponent; as Pruss acknowledges, the Molinist God can simply employ knowledge of the relevant similarity principle along with his Middle Knowledge to arrive at absolute confidence about the acceptability of issuing prophecies in at least some particular cases. This will involve no unacceptable circularity in explanation or unpalatable claims about God's knowledge of disjunctive propositions.⁴³

There are a couple of lessons that the Molinist should bear in mind, however, as she attempts to fine-tune this account. One issue is that, if the Molinist attempts to hold that the relevant similarity principle is contingent, there may be difficulties in motivating the claim that it can ground the absolute certainty God needs in order to prophesy responsibly. Recall that, in a sense, the problem Pruss originally took Flint to task over was his failure to explain how a harmonious set of F-conditionals could be extrapolated to a situation of genuine prophecy. While it does not seem that Flint really attempts to distinguish the purely accidental harmony of a set of F-conditionals with harmony grounded in something general but nevertheless contingent, it may be that at the end of the day this distinction is very hard (or even impossible) to draw. What would this sort of general (but nonetheless contingent) feature look like, and how would we know that a genuine prophecy world would be a member of the "glowing galaxy," as it were?⁴⁴

Consequently, the Molinist might have better luck going with the claim that the relevant similarity principle holds necessarily. While this approach might come with its own set of mysteries, those mysteries would probably not be unpalatable, nor would it be unprincipled to live with them.⁴⁵ (An example of such a mystery: how would we represent in

⁴³It will involve no unpalatable circularity in explanation assuming the basic outlines of Molinism are correct. As discussed previously, we are not entertaining what we have called "external" critiques of Molinism here—i.e., challenges to the metaphysical or epistemological plausibility of Molinism as a general position.

⁴⁴Pruss raises a similar issue ("Prophecy Without Middle Knowledge," 447).

⁴⁵Pruss does acknowledge issues for the view that the principle holds necessarily, but points out that many of these are general problems that face us in our understanding of necessary truths ("Prophecy Without Middle Knowledge," 449).

possible worlds terms a relevant similarity principle that held in all possible situations *except* genuine prophecy and one that held in all possible situations *simpliciter*?)

If the Molinist holds that the relevant similarity principle is true necessarily, though, there may be a further price if Molinism is to avoid its own charges of unprincipled conduct. For we might wonder: why should a relevant similarity principle apply only to a particular agent? (Why, in other words, should the principle only attend to circumstances that stand in a particular relation to one another, but where the circumstances always involve the same agent?) Maybe different agents who are qualitatively similar enough to one another would always freely do the same thing if placed in the same circumstances. So, given that Curley accepts a \$5,000 bribe in circumstances C (and so, consequently, Curley would also have freely accepted a \$5,000 bribe in circumstances R-related to C), why suppose—to use an extreme example—that someone qualitatively identical to Curley would not also have freely accepted such a bribe in C? Taking the principle this far might be enough to make the Molinist a little squeamish, but it also might be required if the Molinist is not to wind up in her own subtle (but nonetheless potentially damning) state of hypocritical half-heartedness.

In any case, we have gone far enough for present purposes. It is clear that foreknowledge-only accounts fail on their own terms to account for prophecy of free actions for two separate reasons—one involving an implausible view of God's knowledge of disjunctions and the other involving an unprincipled attempt to strike a balance between two competing extreme views of God's knowledge (as it relates to F-conditionals). And it is also clear that Molinism does quite a bit better on its own terms. While it may not have yet cleared up every little mystery and it may not be without its theoretical costs, so long as its basic outlines are coherent, it does not run afoul of circularity problems (like the ones associated with disjunctive knowledge for the foreknowledge-only proponent) nor does it have difficulty taking a principled and consistent stand on God's knowledge of primitive closeness relations.

St. Joseph's University