

Against Middle Knowledge

PETER VAN INWAGEN

As a result of my teen-age conversion to the Catholic Church . . . I read a work called *Natural Theology* by a nineteenth-century Jesuit . . . and found it all convincing except for two things. One was the doctrine of *scientia media*, according to which God knew what anybody would have done if, e.g., he hadn't died when he did. . . . I found I could not believe this doctrine: it appeared to me that there was not, quite generally, any such thing as what would have happened if what did happen had not happened, and that in particular there was no such thing, generally speaking, as what someone would have done if . . . and certainly that there was no such thing as how someone would have spent his life if he had not died a child. I did not know at the time that the matter was one of dispute between the Jesuits and the Dominicans, who took rather my own line about it. So when I was being instructed a couple of years later by a Dominican at Oxford . . . and he asked me if I had any difficulties, I told him that I couldn't see how that stuff could be true. He was obviously amused and told me that I certainly didn't have to believe it, though I only learned the historical fact I have mentioned rather later.

—G. E. M. Anscombe, *Metaphysics and the Philosophy of Mind* (Introduction)

According to the sixteenth-century Spanish Jesuit Luis de Molina, God's knowledge—the *content* of His knowledge, the things that He knows—may be divided into three parts.¹ There is, first, His *natural* knowledge: the set of true propositions knowledge of whose truth belongs to His *nature*. This is just exactly the set of all metaphysically necessary truths, for knowing that *p* is an essential property of God if and only if *p* is a metaphysically necessary truth. There is, secondly, His *free* knowledge: those propositions that are true but are such that

it is or once was within His power so to arrange (or so once to have arranged) matters that they be false—that there are exactly 80,966 hairs on Tom’s head, for example, or that there is a physical universe. There is, thirdly and finally, God’s *scientia media* or middle knowledge, which comprises those metaphysically contingent truths such that it is *not* (and never has been) within God’s power so to arrange or so to have arranged matters that they be false. Many philosophers and theologians would insist that this third class of propositions is empty. Molina, however, insisted that it contains propositions of the type nowadays called “true counterfactuals of freedom.”

I shall assume that there is such a thing as middle knowledge only if there are true counterfactuals of freedom—or, as I shall sometimes say, only if there are such things as “middle facts.”² A counterfactual of freedom is a “would” conditional that says of a certain agent that that agent would freely have performed a certain act in certain circumstances that do not in fact obtain. For example:

If Curley had been offered a bribe of \$20,000 at noon last Friday to permit the destruction of Old North Church, he would freely have accepted it

is a counterfactual of freedom, provided that Curley was not in fact offered a bribe of the kind specified at noon last Friday. And if this proposition is *true*, it is a middle fact. (If, that is, facts are true propositions. Those who wish to distinguish between facts and true propositions may substitute whatever phrase they think is appropriate for ‘is’.)

My purpose in this essay is to argue that there is no middle knowledge because there are no middle facts. (If there are no middle facts, then, obviously, the proposition that God does not have middle knowledge does not entail that God is not omniscient. If something is not there to be known, then even an omniscient being will not know it. If all counterfactuals of freedom are false—or lack truth-values—then an omniscient being will know *that*, but knowing of each counterfactual of freedom that it is false or truth-valueless does not constitute middle knowledge.)

I will argue that all counterfactuals of freedom are false.³ My argument proceeds from two premises that I believe to be true but will not defend. The first is that David Lewis’s semantics for counterfactual conditionals is the correct semantics for counterfactual conditionals.⁴ The second is incompatibilism, the thesis that free will is incompatible with determinism. Anyone who rejects either of these premises may regard my argument as having a conditional conclusion: *If* Lewis’s semantics is correct and *if* free will is incompatible with determinism, then all counterfactuals of freedom are false. (As regards the second of these premises, it should be pointed out that the defenders of middle knowledge have always assumed that free will is incompatible with determinism—or that, as they sometimes put it, “a libertarian account of free will is correct.”⁵ It would, I suppose, be a part of what most people who use the term mean by ‘middle knowledge’ that the existence of middle knowledge entails that free will is incompatible with determinism. I much prefer to say that if free will and determinism are compatible—and if determinism is true and if people do sometimes act freely—

then there will, of necessity, be true counterfactuals of freedom and an omniscient being would, of necessity, have middle knowledge.⁶)

Here, briefly, is Lewis's semantics for counterfactual conditionals (I use his well-known symbol for the counterfactual-conditional connective):

$p \Box \rightarrow q$ is true if and only if either p is impossible or some world in which p and q are both true is closer to actuality than any world in which p is true and q is false.⁷

(The second disjunct of the definiens is sometimes informally stated like this: q is true in all the closest p -worlds.) In the sequel, I shall ignore the first disjunct of the definiens; I shall always assume that the antecedent of whatever counterfactual we may be considering is possible. The above counterfactual of freedom is then true in just the following case:

Some world in which *Curley was offered a bribe of \$20,000 at noon last Friday to permit the destruction of Old North Church and Curley freely accepted the bribe* are both true is closer to actuality than any world in which *Curley was offered a bribe of \$20,000 at noon last Friday to permit the destruction of Old North Church* is true and *Curley freely accepted the bribe* is false.

Is there a world having this property? In attempting to answer this question, we must remind ourselves what incompatibilism actually implies about free actions. We must remind ourselves what the proposition that a certain action is free actually implies if free will is incompatible with determinism. Let us examine the case of the thief and the poor-box that I discussed in *An Essay on Free Will*:

[A hardened thief] is in the act of lifting the lid of the poor-box in a little country church. He sneers and curses when he sees what a pathetically small sum it contains. Still, business is business: he reaches for the money. Suddenly there flashes before his mind's eye a picture of the face of his dying mother and he remembers the promise he made to her by her death-bed always to be honest and upright. This is not the first occasion on which he has had such a vision while performing some mean act of theft, but he has always disregarded it. This time, however, he does *not* disregard it. Instead, he thinks the matter over carefully and decides not to take the money. Acting on this decision, he leaves the church empty-handed. (Van Inwagen, pp. 127–28)

We assume that this was a free act and that free will is incompatible with determinism. What does this imply? It implies that if God had caused the world to revert to exactly the state in which it was just before the thief made the decision not to steal the money (and had then allowed the world to “go forward again”), the thief might well have stolen the money “this time.” If, moreover, there were a large number of occasions on which God caused the world to revert to exactly the state it was in just before the thief made the decision not to steal the money, a suitably placed observer of all these “replayings” would find that sometimes the thief stole the money and sometimes he did not.⁸ In other words, if t is the

moment at which the thief made his decision, there are physically possible worlds (worlds in which the laws of nature are the actual laws) of two kinds: Worlds whose histories are identical with the history of the actual world up to t in which the thief stole the money (no doubt there are infinitely many such worlds if there are any at all); and worlds whose histories are identical with the history of the actual world up to t in which the thief did not steal the money (the actual world, at least, is such a world, and no doubt there are infinitely many others). If there are no worlds of the former kind, the thief's act (leaving the church empty-handed) was determined by the past and the laws of nature, and—if free will is incompatible with determinism—that act was therefore not free.

Now suppose that at the moment the thief was about to decide whether to steal the money, he was alarmed by some noise and fled the church without making a decision. What (free) decision *would* he have made if the noise had not occurred (and if nothing else had prevented his making the decision)? In the actual world—proper name ' α '—the thief freely decided not to steal the money; our question concerns some very close world w in which a sudden noise frightened him and sent him running. We want to know whether the counterfactual conditional *If things had been at t exactly as they were in α at t , the thief would have freely decided not to steal the money* is true or false in the world w —or we might ask the same question about *If things had been at t exactly as they were in α at t , the thief would have freely decided to steal the money*. (These two propositions are contraries: in any given world, either one is true and one false, or they're both false.) Does the thief freely decide not to steal [freely decide to steal] the money "in all the worlds closest to w " in which things were at t exactly as they were in α at t ? If he freely decides not to steal in every member of this class of worlds, then the former proposition is true in w ; if he freely decides to steal in every member of this class of worlds, the latter proposition is true in w ; if he freely decides not to steal in some of them and freely decides to steal in others, then both propositions are false in w .⁹

It would seem that α is one of the worlds in this class—one of the worlds closest to w in which things were just as they were in α at t . (We could certainly so choose w that α is one of the worlds closest to w in which things were just as they were in α at t . Let us suppose that we have done so.) It follows that *If things had been at t exactly as they were in α at t , the thief would have freely decided to steal the money* is false in w . But what about its contrary, the proposition that he would have freely decided not to steal the money? Could that proposition be true in w ? That proposition will be true in w only if *every* world whose history is identical with the history of α up to t and in which the thief freely decides to steal the money is more distant from w than α is.

Let β be such a world—a world whose history is identical with the history of α up to t and in which the thief freely decides to steal the money. Let us recapitulate what we know about ways in which the three worlds are related to one another: their histories are identical up to the moment at which the thief is distracted from his deliberations by the noise; thereafter, α and β diverge from w . A brief moment later, at t , the thief freely decides (in α) not to steal the

money, and freely decides (in β) to steal the money and α and β accordingly diverge.

Does it follow from these things we have stipulated that α is closer to w than β is? Well, what are the relevant considerations? If α is closer to w than β is, this must be due to what goes on in the two worlds after t . How should one go about making judgments about the “distances” between these worlds and evaluating these judgments?

There are various ways to think about this sort of question. Suppose we follow David Lewis and think about the “distances” between worlds in terms of their overall similarities and differences. Suppose we identify our question with the question whether α is more similar to w than β is. It is hard to discern any factor that could be relevant to answering this question other than the “ranges of particular fact” that hold in the three worlds. The laws of nature are the same in all three worlds, these laws are indeterministic, and the case of the thief as we have imagined it implies that no miracle (no “large” miracle, no “small” miracle, no miracle in either the usual sense of the word or in Lewis’s special technical sense) is involved in the divergence of α and β .¹⁰

Is there anything about the way things go in β from which we are obliged to infer that β is more unlike w than the actual world is? It would not seem so, if only because we have not said very much about β . In fact, β is not really a particular world at all but an “arbitrarily chosen object.”¹¹ We have picked out β by means of an indefinite description: ‘a world that is identical with α up to t , and in which the thief freely decides to steal’. There no doubt exists a vast infinity of worlds that satisfy this description and nevertheless differ radically from one another. And what it seems most reasonable to believe about this vast infinity of worlds is this: some of them are less like w than α is, some of them are more like w than α is, and some of them are neither more nor less like w than α is. All the worlds that satisfy our indefinite description are, after all, indeterministic, and anything that is consistent with their common indeterministic laws and their common state at the moment of their collective divergence from the actual world will happen in one or another of them given sufficient time—including “reconverging” with w .¹²

We are not obliged to look at the concept of distance between possible worlds as Lewis does—in terms of overall similarity among worlds. I myself prefer to think of the “ p -worlds that are closest to w ” as those worlds in which p is true that could be “obtained” from w by making the smallest possible “changes” in w . (I assume that we are able to distinguish between a change and its consequences, and between the magnitude of a change and the magnitude of its consequences.) A moment’s reflection will show that the class of p -worlds that can be obtained from w by the smallest possible changes in w and the class of p -worlds that are most similar to w need not be the same—for some of the worlds in the former class might, throughout the entire future, diverge radically from w , and thus not belong to the latter. (The minimal changes in w that secure the truth of p might be miracles, replacements of the “entire past” of w by other “entire pasts,” or indeterministic “swerves” of some of the inhabitants of w . Or some of

them might be of one of these types and some of the other two. However we understand the idea of a “change,” some worlds produced by the smallest changes in w that are sufficient to secure the truth of p will not be among the worlds that are—over all—most similar to w .)

Suppose, then, that we look at distance among possible worlds as a function of “size of change.” Does it follow from any of the stipulations we have made about β and w that β is farther from w than the actual world is? It would not seem so. A certain event—the noise that frightens the thief—occurs in w but not in α and β , and thereafter the latter two worlds (still “coinciding”) diverge from w ; a bit later, at t , the thief makes one decision in α and another in β , and those two worlds diverge. To change w into α , therefore, what is required is the noise plus the undetermined course of the thief’s deliberations taking one “fork in the road”; to change w into β , what is required is the noise plus the undetermined course of the thief’s deliberations taking the other fork. The noise, so to speak, cancels out. Can the fact that the undetermined course of the thief’s deliberations went one way in the actual world and the other in β have the consequence that the actual world could be got from w by a smaller change than any that would suffice to produce β ? Remember that if, on a vast number of occasions, God miraculously caused the universe to revert to the precise state it was in just before t , and then allowed things to proceed without further miracles, on some of these occasions the thief would refrain from stealing the money and on some of them he would steal the money. Given that this is the case, it is very hard indeed to see why anyone would say that a larger change in w is required to produce the state of affairs *the thief decides to steal* than is required to produce the state of affairs *the thief decides not to steal*.

It would therefore seem that whether we understand “ q is true in all the p -worlds closest to w ” as

q is true in all the p -worlds that are, over all, the most similar to w

or as

q is true in all the worlds that can be got from w by the smallest changes that suffice for the truth of p ,

it is not the case that the thief freely decides to refrain from stealing in all the worlds closest to w in which his deliberations about whether to steal are uninterrupted by the noise. On either interpretation, therefore, the counterfactual conditional

If the thief’s deliberations about whether to steal had not been interrupted by the noise, he would freely have decided not to steal

is false in w . (We have already seen that ‘If the thief’s deliberations about whether to steal had not been interrupted by the noise, he would freely have decided to steal’ is false in w .) But the arguments that have led us to this conclusion in no way depend on the particulars of the case we have imagined. If these

arguments are sound, then there are no middle facts (no true counterfactuals of freedom) and hence no middle knowledge.

We should note that this conclusion depends on no features of free acts other than the fact that they are undetermined (a consequence of incompatibilism). Arguments that differ from the above in no important way show that there are no facts about what the outcome of any non-actual course of events would have been if its outcome would have been undetermined. Consider, for example, a contraption that has the following design. It is a little box on the outer surface of which there are a button and two lights, a red light and a green light. If someone presses the button, one or the other (never both) of the lights will light. And suppose that it is genuinely undetermined which light will light on any particular occasion on which the button is pressed. (Perhaps the outcome of pressing the button depends on individual quantum events.) Suppose that no one pressed the button at noon. If someone *had*, which of the lights would have lit? If our arguments concerning the thief are correct, similar arguments will show that we can say only that one or the other would have lit; we cannot say that the green light would have lit, because it is false that the green light would have lit; we cannot say that the red light would have lit, because it is false that the red light would have lit. (Of course it is also false that the green light would *not* have lit, and false that the red light would *not* have lit.) God, being omniscient, knows all these things, knows that it is true that one light or the other would have lit, and knows that all the other propositions I have mentioned are false. But He does not know which light would have lit because no such thing is there to be known.

God might, however, know—He would if it were true—that if the button had been pressed at noon, the green light would *probably* have lit. And such a thing could easily be true. (Suppose that there is a magical, truly indeterministic coin—a fair coin such that if it is tossed it is truly undetermined whether it will fall “heads” or “tails.” God cannot know whether, if you had tossed the magical coin once at noon—you didn’t—it would have fallen “tails,” for, again, no such thing is there to be known. Nor does He know that if you had tossed it five times at noon, it would have fallen “tails” at least once, for in some of the closest worlds in which you tossed it five times at noon, it fell “heads” every time. But He can and does know that if you had tossed it five times at noon, it would *probably* have fallen “tails” at least once. He knows this because it is true. It is true because, in *most* of the closest worlds in which you tossed the coin five times at noon, it fell tails at least once (in 96.875 percent of them, to be exact).¹³ Depending on the details of the design of the red-green device, the probability of the red light’s lighting if the button is pressed may be 0.5 or 0.25 or 0.3333 . . . or any other number between 0 and 1. And whatever the probability of the red light’s lighting may be (conditional on the button’s being pressed), God will know what it is, for this thing, this probability, *is* there to be known.

No matter how probable it may be—short of certainty—that the red light would light if the button were pressed, God does not know whether that light would have lit if the button had been pressed on some occasion when it was not pressed.¹⁴ Even if the probability of “red” is 0.9999999, it will still be true that if

the button had been pressed, the red light *might* not have lit, and this proposition is equivalent to the denial of the proposition that, if the button had been pressed, the red light would have lit. As Lewis has pointed out (1973, p. 2), the “would” and “might” conditional connectives are interdefinable:

$$\begin{aligned} p \Box \rightarrow q &=_{df} \sim(p \Diamond \rightarrow \sim q) \\ p \Diamond \rightarrow q &=_{df} \sim(p \Box \rightarrow \sim q). \end{aligned}$$

The second of these schemata suggests another sort of argument for the nonexistence of true counterfactuals of freedom, a very simple argument indeed. Let us suppose that if Curley had been offered a certain bribe (which a building contractor in Chelsea was seriously considering offering him, but, in the end, did not), he would freely have accepted it. I will argue that if free will is incompatible with determinism, then the falsity of this counterfactual of freedom follows from its truth. Assume

- (1) Free will is incompatible with determinism

and

- (2) Curley was offered the bribe $\Box \rightarrow$ Curley freely accepted the bribe.

Premise (2) entails

- (3) Curley was offered the bribe $\Box \rightarrow$ Curley was able not to accept the bribe.

Premise (1) and proposition (3) together entail

- (4) Curley was offered the bribe $\Diamond \rightarrow \sim$ Curley accepted the bribe.

By the second $\Box \rightarrow / \Diamond \rightarrow$ schema above, proposition (4) entails

- (5) \sim (Curley was offered the bribe $\Box \rightarrow$ Curley accepted the bribe),

and (5) entails

- \sim (Curley was offered the bribe $\Box \rightarrow$ Curley freely accepted the bribe),

which is the denial of (2), the counterfactual of freedom with which we started. It is obvious that a parallel argument could be constructed with respect to any counterfactual of freedom.

I wish finally to consider two minor but interesting points about counterfactuals of freedom.

1. Alvin Plantinga has presented a very pretty little argument for the conclusion that there are true counterfactuals of freedom (Plantinga, p. 177). Suppose (he argues) that Curley was offered a bribe of \$20,000 and accepted it. Then, surely, it is evident that if he had been offered a bribe of \$35,000 on the same occasion to secure the same service, he would have taken it? But then there is at least one true counterfactual of freedom.

One might wonder, however, how we know that Curley would have *freely* accepted the larger bribe. For all anyone but God knows, Curley was able to re-

fuse the smaller bribe but would have been unable to refuse the larger. Well, let us simply *assume* that it is true that he would have been able to refuse the larger bribe. But then—given incompatibilism—we know that if he had been offered the larger bribe, he might have refused it. And how could we know, how could even God know, that what might have happened would not have happened? We certainly could *not* know this (for it would not be true) if ‘it might have happened’ is logically equivalent to ‘it is false that it would not have happened’, as it is on Lewis’s semantics. (At *t*, I am asked to toss the magical coin once; I immediately do so, and it lands “heads.” Do I know that if, at *t*, I had been asked to toss it twice and had immediately done so, it would have landed “heads” at least once? I do not, for that is not true.)

2. If my arguments are correct, they show not only that there are no true counterfactuals of freedom, but that there *could* not be any (for my premises, if they are true at all, are necessary truths). Does it follow that counterfactuals of freedom are necessary falsehoods? That they are is not a formal consequence of my conclusion, since *being a counterfactual of freedom* is an accidental property of those propositions that have it, and it might be that some counterfactuals of freedom are only contingently false because they are true in some worlds in which they are not counterfactuals of freedom. Counterfactuals of freedom are only accidentally counterfactuals of freedom because they are not counterfactuals of freedom (as I have defined the term) in worlds in which their antecedents are true. Proposition (2) above, for example, is not a counterfactual of freedom in those worlds in which Curley was offered the bribe. Might (2) be true in some of these worlds? It not only might be true in some of these worlds but *is* true in some of them on one version of Lewis’s semantics, the version that incorporates the assumption of “unweakened centering”—that is, the assumption that no other world is as close to any world as that world is to itself.¹⁵ (And that assumption derives a certain degree of plausibility from the normal associations of words like ‘closer’ and ‘distance’. In their normal, nonmetaphorical sense, these words apply to places, and no other place is as close to any place as that place is to itself. If we think of closeness among worlds as Lewis does, as a function of similarity, the unweakened centering assumption can seem undeniable: if w_1 and w_2 are distinct worlds, then things must go differently in them—and in that case, how could w_2 be as similar to w_1 as w_1 is?¹⁶ Perhaps, however, the assumption can seem less obvious if we think of the distance from w_1 to w_2 as being given by the size of the change in w_1 that would be required to “turn it into” w_2 ; if w_1 is indeterministic, then it may be that, in one sense, *no change at all* is needed to turn it into a world with which it coincides up to some point in time and in which the same indeterministic laws hold.) On the version of Lewis’s semantics that incorporates unweakened centering, ‘ $p \Box \rightarrow q$ ’ is a logical consequence of ‘ $p \& q$ ’, and (2) will therefore be true in any world in which Curley is offered the bribe and takes it. If, however, we reject the unweakened centering assumption—if we accept “weakened centering,” if we allow the formal possibility of other worlds that are as close to a world as that world is to itself—there is no formal or logical reason to suppose that there is any counterfactual of freedom that is only contin-

gently false. If, moreover, it is indeed possible for some world to be as close to some other world as the latter is to itself, there would seem to be no reason to doubt the following thesis: For every world w , and for every “would” conditional of freedom p whose antecedent and consequent are both true in w , there is (assuming incompatibilism) a world as close to w as w is to itself in which the antecedent of p is true and its consequent is false. And if that is the case, then every “would” conditional of freedom is a necessary falsehood.

NOTES

1. The brief account of Molina’s concept of middle knowledge that is contained in this paragraph is based on Flint (1988). Flint’s article may be consulted for citations of Molina and for an overview of the literature on middle knowledge and related topics. Many useful references can also be found in Adams (1977). (I may say that I find little to disagree with in Adams’s classic essay, and that there is perhaps little in the present essay that cannot be found, at least in embryo, in Adams. This essay is a rather more technical version of what I have been saying in undergraduate lectures in courses on the philosophy of religion for many years. It is my hope that the reader who finds the general outline of the arguments presented here “old news” will at least find some of the details interesting.)

2. My business in this essay is with counterfactuals of freedom. I will not address the question whether, if there are true counterfactuals of freedom, at least some of them would have to be items of God’s middle knowledge in Molina’s sense. Nor will I address the question whether my assumption that there is middle knowledge only if there are true counterfactuals of freedom is in fact correct. It should be noted, however, that if God does indeed have middle knowledge of some truths, it cannot be that all of them are true counterfactuals of freedom. It should also be noted that, if true counterfactuals of freedom are possible, it would seem to be possible for at least some true counterfactuals of freedom to belong to God’s free knowledge. As to the first point: If a true counterfactual of freedom belongs to God’s middle knowledge, then so does any contingent truth that it entails; for example, the material conditional with the same antecedent and consequent. (And if true counterfactuals of freedom are items of God’s middle knowledge, then no doubt so are certain “true counterfactuals of indeterminism”—such as the counterfactual about the “red-green device” that will be considered later in the body of the essay—that are not themselves counterfactuals of freedom.) As to the latter point: Suppose that in all the closest possible worlds in which Adam is placed in a certain determinate set of circumstances, he freely disobeys God. Then it is true that if he were placed in that set of circumstances, he would freely disobey God. But, consistently with this, we can go on to suppose that there is a world that it is within God’s power to actualize in which Adam is not placed in those circumstances and in which it is false that if he were placed in those circumstances he would freely disobey God.

3. My argument for this conclusion will depend on the assumption that every proposition is either true or false. Without this assumption, my argument would show only that no counterfactual of freedom is true. But this conclusion is sufficient to establish the further conclusion that there is no middle knowledge. My argument for the nonexistence of middle knowledge does not therefore depend on my assumption that every proposition is true or false.

4. My official term for the type of conditional of which Lewis’s semantics treats is “would” conditional. Officially, I apply the term ‘counterfactual conditional’ only to “would” conditionals with false antecedents. But I shall not usually speak officially. When I am speaking generally about “would” conditionals I shall mostly follow established usage and call them ‘counterfactual conditionals’ or ‘counterfactuals’. But I shall in every case apply the special term ‘counterfactual of freedom’ to a conditional only if it has a false antecedent.

5. Or, worse, that in the phrase 'counterfactual of freedom', the word 'freedom' is to be understood "in its libertarian sense." I keep telling people not to use the word 'libertarian', but no one listens. There are no "libertarian accounts of freedom," or, more exactly, all incompatibilist accounts of freedom that anyone has actually offered are entirely unsatisfactory. And the word 'freedom' does not have a "libertarian sense," because the compatibilist and the incompatibilist use the word in the same sense.

6. It does not, however, follow that for any given pair of contrary counterfactuals of freedom, one of them will be true. If the mechanisms underlying human deliberation were, although fully deterministic, sufficiently "chaotic"—if small, indeed humanly imperceptible, variations in the conditions under which these mechanisms operated could change the outcome of a deliberation—it might be that both members of many such pairs of contraries were false. Let us suppose that we are compatibilists and determinists. Still, if we believe that the mechanisms that underlie human deliberation are chaotic, we must admit that there are many worlds in which the proposition *Curley was offered a bribe of \$20,000 at noon last Friday to permit the destruction of Old North Church* is true. And if we believe that human deliberation is a "chaotic" process, we must admit that, for all we know, there are two such worlds, both of them among the closest worlds in which Curley was offered the bribe, in one of which he freely accepted the bribe and in the other of which he freely refused the bribe. Then it is false that if Curley had been offered a bribe of \$20,000 at noon last Friday to permit the destruction of Old North Church, he would freely have accepted it—and false that he would freely have refused it. And we should generalize this concession: for all we know, both members of many such pairs of contraries are false, owing to the extreme sensitivity of the outcome of human deliberation to imperceptible variations in the conditions under which that deliberation is carried out.

7. Lewis (1973).

8. Or, to be pedantic, the probability that this is what the observer would see would approach 1 as the number of replays increased without limit.

9. Here is another way they could both be false in w : suppose that for every world in which the common antecedent of the two counterfactuals is true and he steals, there is a world closer to w in which the antecedent is true and he does not steal; and that for every world in which the common antecedent of the two counterfactuals is true and he does not steal, there is a world closer to w in which the antecedent is true and he steals.

10. This paragraph contains some allusions to the discussion in Lewis (1979) of the factors that are relevant to judgments of overall similarity among worlds. The reader can follow the argument of the present essay without understanding the allusions.

11. As is w , which is specified as follows: a world that coincides with α up to some point after the thief has begun his deliberations, in which a noise frightens him away before he can reach a decision, and which is such that α is among the worlds closest to it in which the thief completes his deliberations.

12. And even if we neglect the fact that these worlds are indeterministic, we should note that there are lots of ways in which the event by which β diverges from α —the thief's deciding to steal—might happen. (The thief might take a shorter or longer time to decide to steal; he might decide to put the money in his left pocket or his right pocket.) If, therefore, the physical world is at all "chaotic" then the worlds in which the thief decides to steal might differ radically from one another, some being more like w than α is. (Cf. n 6.)

13. No doubt there are infinitely many possible worlds. If there are infinitely many worlds, no doubt there are infinitely many worlds in which the coin is tossed five times and falls "tails" at least once and also infinitely many worlds (an infinity of the same cardinality) in which the coin is tossed five times and falls "heads" every time. What we are supposing is that many infinite sets of worlds have a "size," a measure, that is not determined by their cardinality. Sets of possible worlds are thus analogous to sets of points on a plane, many of which have sizes or measures that are not determined by their cardinalities: one set of points may occupy two square feet, and another set of the same cardinality (the power of the continuum) may occupy one square foot—or a square light-year or square micron. In my view, it is impossible to make

sense of probabilities without making this assumption. When I say that the coin falls "tails" at least once in 96.875 percent of the worlds in which it is tossed five times, I mean that the ratio of the measure of the set of worlds in which the coin falls "tails" at least once out of the five times it is tossed to the measure of the whole set of worlds in which it is tossed five times is 96.875.

14. A set of points may easily have the same area as one of its proper subsets. No doubt the same thing may hold, *mutatis mutandis*, for sets of worlds and their measures. There is, therefore, at least a formal possibility that $p \Box \rightarrow q$ might be false even if the probability of q , given p , is 1. I shall not explore the implications of this formal possibility for the problem of God's middle knowledge.

15. See Lewis (1973), pp. 14–15, 26–31.

16. But see Lewis (1973), p. 29.

REFERENCES

- Adams, Robert Merrihew. 1977. "Middle Knowledge and the Problem of Evil." *American Philosophical Quarterly* 14: 109–17.
- Flint, Thomas P. 1988. "Two Accounts of Providence." In Morris.
- Lewis, David. 1973. *Counterfactuals*. Cambridge, Mass.
- . 1979. "Counterfactual Dependence and Time's Arrow." *Noûs* 13: 455–76.
- Morris, Thomas V., ed. 1988. *Divine and Human Action: Essays in the Metaphysics of Theism*. Notre Dame, Ind.
- Plantinga, Alvin. 1974. *The Nature of Necessity*. Oxford.
- Van Inwagen, Peter. 1983. *An Essay on Free Will*. Oxford.