

**DOES LIBERTARIAN FREEDOM REQUIRE
ALTERNATE POSSIBILITIES?**

Linda Zagzebski
University of Oklahoma

I

According to causal determinism the chain of causes of an act leads backwards in time to something completely outside the agent. What bothers libertarians is that it seems to follow that something outside the agent is the ultimate cause of the act, what makes it happen. And if what ultimately makes an act happen is what ultimately is responsible for it, then causal determinism seems to imply that no one is ultimately responsible for his acts. Of course, it is a major philosophical challenge to explain what is involved in making something happen and it is sometimes tempting to shift the discussion to something easier to grasp. But whatever digression we take, we should not forget that once we have identified who or what makes an act happen, we have identified who or what is responsible for it, if anything is.

One way to get criteria that insure that an act is not causally determined is to insist that there be no conditions existing just prior to the time of the act that are sufficient for the agent to do the act.¹ That is, the entire history of the world up to but not including the time T at which an act A occurs is compatible both with A's occurring and with A's not occurring. In fact, that condition is often taken to be constitutive of the denial that acts are determined. Call this the Temporal Contingency Principle (TCP):

(TCP) An act A at T is non-determined (temporally contingent) if and only if there is a possible world W that has exactly the same history up to T as the actual world and in which A does not occur at T.

If the condition specified on the right side of TCP is met, A is not determined. A could not be made to happen by anything prior to it if both its happening and its not happening are compatible with everything prior to it. But TCP goes well beyond what is required for A to be causally contingent. Not every event in the

temporal history of A is in its causal history. In fact, it is doubtful that any event is such that the entire past history of the world is part of its causal history. Consider what we ask for when we want to know whether two events are related as cause to effect. When someone inquires, “If the money supply increased, would the interest rates go down?”, he is not asking whether the rates would go down if *everything* that happened in the past remained the same and the money supply increased since most of what happened in the past has nothing to do with interest rate increases, for instance, the invasion of Poland by the Tatars in 1287.

TCP is too strong. The principle we need instead is something roughly like the following Principle of Causal Contingency (PCC):

(PCC) An act A is non-determined (causally contingent) if and only if there is a possible world W in which all the events in the causal history of A in the actual world occur and in which A does not occur.

Perhaps it will be objected that if an event is causally contingent it has no causal history. Isn't a causally contingent event a break in the causal order? Indeed, it is, but a causally contingent event is not a random event. It is not an event that has no causally explanatory tie to the events preceding it. If the causal history of an event includes all events relevant to explaining what makes it happen, then all events have a causal history even if some events are such that their causal history is not sufficient to fully explain what makes them happen. The scope of the causal history of an event is a difficult question, of course, but the difficulty of the question should not lead us to conclude that the causal history of an event either includes the entire past or nothing at all.

But TCP is tempting for another reason. Unfortunately, we are ignorant of a great deal of the causal history of most events. It is even possible that U.S. interest rate increases in 2000 are causally related to a distant military event in 1287. So even though the attempt to identify causes is an attempt to identify that part of the past that is its causal history, since, as far as we know, *anything* in the past might be in its causal history, we can only be sure that we have included everything in its causal history if we include the entire past. The same point applies to the criterion for causal contingency. TCP rules out the possibility that apparently irrelevant events in the past causally necessitate an act A by requiring that nothing in the past is such that its occurrence rules out the non-occurrence of A. Since the causal history of an event is a subset of its temporal history, we can therefore be sure that we have considered everything causally relevant to the occurrence of a given event only by relating it to the entire past. If an event satisfies TCP it also satisfies PCC. Therefore, TCP is a principle of safety even though it is stronger than is required. Temporal contingency is sufficient for causal contingency, but it is not necessary.

The stronger TCP is tempting because of our ignorance. But perhaps it will be argued that TCP is not too strong since if any event A failed the condition of temporal contingency we would automatically judge that it is causally deter-

mined. Suppose that if Poland had not been invaded by the Tatars in 1287 but everything else prior to a particular rise in interest rates remained the same, interest rates would not have risen. Doesn't it follow that the invasion is in the causal history of the movement of interest rates in 2000? No, it does not. To think that it does follow is to think that all counterfactual relations between events imply a relation of temporal modality and that all temporal modalities are causal, and we are in no position to think that. To take an historically important example of a temporal modality that is not causal, consider the fact that it is often thought that divine foreknowledge takes away human free will because of the accidental necessity of past infallible foreknowledge. Accidental necessity is the alleged necessity of the past; presumably, once an event occurs it acquires a form of necessity because there is no longer anything anybody can do about it.² If there is past infallible knowledge that A will occur in the future, there is no longer anything anybody can do about either the fact that the foreknowledge occurred or its infallibility. Hence, A cannot but occur. But notice that this argument does not imply that A is causally determined. Therefore, if temporal necessity is a problem for freedom and responsibility it is not because it entails causal necessity. Causal necessity entails temporal necessity but not conversely. Temporal contingency entails causal contingency but not conversely.

Temporal contingency entails that the agent had alternate possibilities at the time of the act. If it is compatible with the entire past history of the world that the agent does either A or not A, then when the agent does A she could have done otherwise. Causal contingency as defined by PCC does not entail alternate possibilities. Again, the foreknowledge example illustrates the point. God's foreknowledge of my act might make it the case that I cannot do otherwise even though God's foreknowledge does not cause my act, and if my act has no other determining causes, it is causally contingent.³

In short, temporal contingency entails both causal contingency and alternate possibilities. Causal contingency does not entail alternate possibilities. It should not be surprising, then, if it turns out that a property requiring causal contingency does not require alternate possibilities. Responsibility may be such a property.

II

One way to formulate the libertarian thesis on responsibility given in the first two sentences of this paper is the following:

(LTR) An agent is morally responsible for her act only if the act is causally contingent.

The thesis that the ability to do otherwise is a requirement for responsibility is given in the following well-known principle:

(PAP) An agent is morally responsible for her act only if she could have done otherwise.

The reasons for holding PAP are no doubt different for different people, but I want to argue that the desire to defend libertarian freedom, the kind of freedom that requires causal contingency, ought not to be one of them. LTR may be true even if PAP is false. Since causal contingency does not entail alternate possibilities it is possible for an act to be causally contingent even when the agent lacks alternate possibilities. I have already mentioned that divine foreknowledge is one way in which that possibility might be actualized. If alternate possibilities are required for responsibility, then, it cannot be because alternate possibilities are a requirement of causal contingency.

Three decades ago Harry Frankfurt presented a famous argument that was intended to drive a wedge between responsibility and alternate possibilities, and Frankfurt thought he could thereby drive another wedge between responsibility and libertarian freedom.⁴ In my judgment he succeeded admirably in breaking apart the conceptual connection between responsibility and alternate possibilities, but he went no distance at all towards supporting determinism.⁵ The connection between responsibility and libertarian freedom remains unbroken. On the contrary, I think that a careful look at Frankfurt cases actually supports non-determinism. But let us first review a standard Frankfurt-style case:

Black, an evil neurosurgeon, wishes to see White dead but is unwilling to do the deed himself. Knowing that Mary Jones also despises White and will have a single good opportunity to kill him, Black inserts a mechanism into Jones's brain that enables Black to monitor and to control Jones's neurological activity. If the activity in Jones's brain suggests that she is on the verge of deciding not to kill White when the opportunity arises, Black's mechanism will intervene and cause Jones to decide to commit the murder. On the other hand, if Jones decides to murder White on her own, the mechanism will not intervene. It will merely monitor but will not affect her neurological function. Now suppose that when the occasion arises, Jones decides to kill White without any "help" from Black's mechanism. In the judgment of Frankfurt and most others, Jones is morally responsible for her act. Nonetheless, it appears that she is unable to do otherwise since if she had attempted to do so, she would have been thwarted by Black's device.⁶

Frankfurt believes that his cases support determinism by falsifying PAP and that they do not presuppose a deterministic universe.⁷ Robert Kane and David Widerker have independently argued that F cases do presuppose a deterministic universe and hence cannot be used to support determinism.⁸ My position is that F cases do not presuppose determinism but neither do they support it. In this section I will argue that F cases do not presuppose determinism. In section III I will argue that F cases probably succeed in falsifying morally significant versions of PAP, and in section V I will argue that F cases give no support to determinism and arguably support non-determinism.

Kane and Widerker agree with Frankfurt that the denial of PAP is closely connected with determinism. Their disagreement is over the issue of whether

PAP can be falsified independently of falsifying determinism. Since, as I have argued, non-determinism does not entail PAP, it should be possible to falsify PAP without either entailing or presupposing determinism. Frankfurt cases probably succeed in doing that, but even if they do not, the argument of section I should lead us to expect that it is possible to construct cases that do. In this section I want to look at the Kane/Widerker argument that attempts to show that no Frankfurt-style case can falsify PAP without presupposing determinism.

The problem as Kane and Widerker pose it can be put briefly as follows: The Frankfurt machine cannot operate unless it is possible to be a perfect predictor in a non-deterministic universe. It is not possible to be a perfect predictor in a non-deterministic universe. Hence, a Frankfurt machine is impossible. More precisely, consider whatever it is that the agent is responsible for according to the libertarian—an act of will, some other agent-caused event—whatever is the original undetermined cause of her act. The machine cannot intervene before that event in a non-deterministic world because no matter what the prior signs, the event might or might not occur up until the moment it does occur. So the machine cannot tell for sure whether or not to intervene. But the machine cannot intervene after the event since then it is too late to falsify PAP; the agent is already responsible. The most the machine can do is to make her change her mind, and that's not an F case.

One way out of the Kane/Widerker objection proposed recently is that of Alfred Mele and David Robb.⁹ Their idea is to postulate two independent causal processes leading up to the act, one deterministic and unconscious, and the other non-deterministic and resulting from conscious deliberation and choice. In their scenario the F-style machine does not simply monitor the agent's neurological activity. It is set to deterministically cause the agent to decide to steal a car at T2. But we are to imagine that he decides on his own to steal the car at T2, the exact moment the deterministic mechanism is set to cause the same decision. Since there are two independent causal sequences in operation in this scheme, there are several possibilities to be considered. The first is the F-style scenario. The deterministic process and the indeterministic process coincide at the moment of decision, T2. Mele and Robb specify that the machine is so designed that in that case the deterministic process ceases operation and the indeterministic process is the one that is effective. The second possibility is that the two processes diverge at T2. Since two contrary choices cannot occur simultaneously, what happens if the agent makes the "wrong" choice at T2? Mele and Robb have thought of this problem and they stipulate that the machine overrides the agent's choice.¹⁰ We need to accept, then, that the machine is designed in such a way that the indeterministic choice overrides the deterministic one when they coincide, but the deterministic choice overrides the indeterministic one when they diverge. This is a stretch, but I would not claim it is impossible.

There are two other possibilities. Since the process of deliberation is also indeterministic, the agent can make his decision prior to T2. Suppose the agent decides on his own before T2 to steal the car. Mele and Robb specify that in

that case the deterministic process in place causes him to decide to steal the car again at T2.¹¹ Suppose instead that the agent decides on his own before T2 not to steal the car. Mele and Robb say that if the device were to issue at T2 in the agent's deciding to steal the car, it would erase any memories that are incompatible with its so issuing.¹² Presumably, then, the machine causes him to change his mind at T2 and to decide to steal the car after all.

These features of the Mele/Robb scenario make it importantly different from more standard F cases. In the latter the agent cannot do otherwise during the entire period of time that the agent is deliberating and making up her mind. The device is set to intervene whenever something in the agent tips off the machine that the "wrong" choice is ahead. That duration is generally assumed to be short, but it is more than an instant. In Mele and Robb's scenario the agent can do otherwise at every moment up to but not including T2. The device is already in operation, leading to a deterministic decision to steal the car at T2 unless the same decision is made indeterministically at T2. Since there is only a single moment at which the agent cannot do otherwise (assuming the description of what happens at T2 is possible), it is not the case that the device insures that a deliberator cannot choose otherwise without making her change her mind. It is only chance that the agent indeterministically makes her decision at T2. She might easily have made it moments earlier, in which case the scenario would not have been an F case. Hence, it is only chance that the scenario turns out to be like Frankfurt's and it is only chance that the Mele/Robb scenario falsifies PAP. That is, it is only chance that in the Mele/Robb scenario the agent is responsible for an act she performed at a moment when she could not do otherwise.

Eleonore Stump has a different approach to answering the Kane/Widerker objection. She argues that if libertarian free choices are correlated with temporally extended neural sequences, an F-style machine can operate in a non-deterministic world by intervening after an initial neural sequence occurs which is a necessary condition for making the choice the machine does not want.¹³ As Stewart Goetz has pointed out, it is not clear exactly which processes/events Stump thinks are causally undetermined and which are not, and indeterminacy in this matter can make her description appear vulnerable to the Kane/Widerker objection.¹⁴ But Stump's point seems to be that intervention does not require infallible prediction since all that is required is that it occur after a causally necessary condition for the choice (or other responsible event) has occurred, but before a causally sufficient condition obtains. For example, the process of making a decision might itself be temporally extended and temporally coincide with a neurological sequence. The machine can intervene after the neurological process begins but before it ends. Alternatively, the neurological sequence precipitating intervention might be a causally necessary condition for the choice which occurs subsequently and is causally contingent.¹⁵

Either option should work. If the process of making the decision has already begun but the outcome is still undetermined, and if the decision cannot be completed without such a beginning, the machine can intervene without need-

ing to perfectly predict the outcome. For the same reason, the machine can intervene after any necessary condition for making the “wrong” choice has occurred. Of course, this means the machine might intervene unnecessarily, but that is not a problem. Stump’s description of F cases in an earlier paper suggests this interpretation.¹⁶ In fact, Frankfurt’s own description of his principal case suggests this interpretation as well:

Suppose someone—Black, let us say—wants Jones₄ to perform a certain action. Black is prepared to go to considerable lengths to get his way, but he prefers to avoid showing his hand unnecessarily. So he waits until Jones₄ is about to make up his mind what to do, and he does nothing unless it is clear to him (Black is an excellent judge of such things) that Jones₄ is going to decide to do something *other* than what he wants him to do.¹⁷

Presumably, even if Black is an excellent judge of what Jones is about to do, he is not infallible. He can infallibly identify neither causally necessary nor causally sufficient conditions for Jones’ choice. But the case requires that he be infallible only with respect to identifying a causally necessary condition. This is a bit more than Frankfurt has a right to claim about the case, but it is considerably less than Widerker and Kane suppose is required. The fact that most commentators on F cases have posed them in terms of an infallibly predictive device has no doubt misdirected subsequent discussion. All that F cases require is the capacity to identify a causally necessary condition for a non-determined choice in advance of the choice, and that is in principle possible in a non-deterministic world.

The best description of a case of this kind I know of is given by Derk Pereboom.¹⁸ In Pereboom’s case the machine is set to make the agent decide to declare an illegal tax deduction if he does not do so on his own. The agent’s psychology is such that a causally necessary condition for him to decide not to commit the illegal act is that a moral reason occur to him with a certain force. All the machine need do is to intervene just after such a moral reason occurs to the agent with the given force. The machine can thereby make it the case that the agent cannot do otherwise even in a non-determined universe. What makes Pereboom’s case especially promising is that many libertarians will agree that thinking of a moral reason is a causally necessary condition for choosing the morally right thing in many circumstances. Such a view not only does not presuppose determinism, but it is likely to be part of a typical libertarian account of what usually happens in cases of moral temptation. I conclude that the description of standard F cases does not presuppose determinism.

III

Some libertarian defenders of PAP contend that the agent in Frankfurt cases does have ethically important alternate possibilities, alternate possibilities entailed by her responsibility, even though she does not have the alternative of

doing otherwise. Consequently, a number of variations of PAP have been proposed and defended against Frankfurt-style examples. It is obvious that some libertarians defend PAP only because they believe that libertarian freedom requires alternate possibilities. Since they are convinced by F cases that PAP is false, they then feel compelled to look around for some variation of PAP that is true and that can withstand Frankfurt's attack. Otherwise, they think, they will be forced to give up belief in libertarian freedom. But if the argument I have given that causal contingency does not entail alternate possibilities is convincing, it may be sufficient to take away the principal motive of some libertarians for finding a true version of PAP. So for them I need not say anything more.

Other libertarians may accept my position that PAP is not entailed by causal contingency and yet continue to maintain that PAP is entailed by libertarian freedom since there is more *in* libertarian freedom than causal contingency. This position requires an argument going from a substantive account of the nature of libertarian freedom to the consequence that the agent has alternate possibilities. What would such an argument look like?

Determinism and its denial are theses about the causes of an act, what makes it happen, whereas PAP is a thesis about what the agent might have caused instead. Libertarians are not committed to a single position on the process whereby an agent causes her act, but there is agreement that the chain of causes of a free act ends in the agent. There is nothing outside the agent that is ultimately the cause of her act. One way in which this position has been explicated is via the concept of agent causation. According to this view the cause of an act for which the agent is responsible is the agent herself, not an event of any kind, not even the event of her exercising her agency. From this point of view agency is a kind of power. Thomas Reid contended that it is essentially a two-way power. An agent does not have the power to cause an act unless she also has the power not to:

Power to produce any effect implies power not to produce it. We can conceive no way in which power may be determined to one of these rather than the other, in a being that has no will.

Whatever is the effect of active power must be something that is contingent. Contingent existence is that which depended upon the power and will of its cause.¹⁹

Here Reid ties the possession of alternate possibilities to agent causation itself and he seems to be inferring the contingency of an act from the fact that it arises in a situation in which the agent had the power not to perform the act.²⁰ So on Reid's view agent causation is partially constituted by the possession of alternate possibilities. Here, then, is an argument connecting a form of the libertarian position on freedom with alternate possibilities. The argument is that agent causation entails alternate possibilities, not because causal contingency entails alternate possibilities, but because the power of agency itself demands it. Let us call this Reid's principle:

(RP) An act A is agent-caused by S only if S had the power not to agent-cause A.

RP can be used to defend a form of PAP. The argument would go as follows:

- (1) An agent S is morally responsible for an act A only if S agent-causes A.
- (2) S agent-causes A only if S had the power not to agent-cause A. (RP)

Therefore,

(3) An agent S is morally responsible for an act A only if S had the power not to agent-cause A.

(3) is a form of PAP. A version of PAP very similar to (3) has been defended by Edward Wierenga as a response to F cases. Wierenga's version is the following:

(PAPC) Necessarily, for every person S and act of will A such that S causes A, S is morally responsible for performing A only if S could have refrained from causing A.²¹

It is clear from Wierenga's discussion that he has agent causation in mind.²² A similar defense of alternate possibilities has been given more recently by Michael McKenna, who argues that in F cases it is up to the agent whether or not her act puts her stamp upon the world.²³ McKenna does not specifically mention agent causation either, but his discussion of the power of an agent to make the act her own or not is clearly similar in spirit to the views of Reid and Wierenga.

Wierenga claims that even in F cases S could have refrained from causing A. If S makes A happen, then it was open to her to make it happen or not to make it happen. But it seems to me that even though it was open that she not make A happen, it was not open *to her* not to make A happen. It was open to the machine to make her make A happen. Of course in that case S would have exercised the power not to do A if the machine hadn't intervened. But as it turned out, the machine intervened before she could exercise that power. The most she can do is to fail to will, but even then it is the machine that makes it happen that she fails to will. *She* can't do anything to avoid being responsible for the act. The machine does it for her.

It appears, then, that "could have refrained" in PAPC does not refer to a power that S has. It is true that either the act is agent-caused or an act of the same type is machine-caused. It is also true that in most descriptions of F cases the machine's operation depends upon what the agent does.²⁴ But even then it

is not up to the agent whether or not the machine operates. It is only because of what she *would have done* that the machine goes into operation. If the machine operates, it is before she can exercise an alternate power. In Pereboom's example the machine does operate because of what the agent actually does—thinking of a certain moral reason. But having a moral thought is not exercising the right kind of power for attributing moral responsibility if indeed it is exercising a power at all. Again, the machine operates before she is in a position to exercise any contrary power.

McKenna makes another move in an attempt to preserve the idea that the agent in F cases retains a two-way power. He says an agent can have a power even though she is unable to exercise it. I have said that in F cases the machine takes away the agent's ability to exercise the power to choose otherwise. Might it still be the case that she has the power to choose otherwise even though the machine prevents her from exercising it?²⁵

I am inclined to think we should be as suspicious of unexercisable powers as of imperceptible trees. But although I think that doubt is in order, I actually agree with McKenna that there are powers one can have without being able to exercise them at every moment at which the power is possessed. The power to perform such complex activities as playing a Schubert Impromptu, making bread, or writing one's will is not compromised by the fact that a piano, kitchen equipment, or writing materials is not available at this moment. The fact that I cannot exercise the power right now does not detract in any significant way from my present possession of the power since it is not the kind of power that requires continuous exercisability. These activities take some time to perform and, in the case of making bread, the time it takes is discontinuous. The fact that I cannot initiate the process at this moment is of no consequence. In contrast, if my eyes are now bandaged shut, I do not now have the power to open my eyes. The power to open my eyes is the kind of power I need to be able to exercise now if I possess it now. That is probably because opening my eyes is a basic act or very close to a basic act, whereas playing an Impromptu and making bread are far from basic acts. My conjecture is that the possession of a power at T and the ability to exercise it at T split apart only for some acts that are clearly not basic acts. One does not at T have the power to perform a basic act unless one also possesses at T the ability to exercise that power. The power to refrain from choosing an act is like the power to open my eyes and unlike the power to make bread. Choosing and refraining from choosing are basic acts and if my conjecture is right then, like opening my eyes, they are possessed at T only when I can exercise them at T.

In any case, even if I am wrong in this conjecture and an agent can possess the power to refrain from choosing in F cases without having the ability to exercise that power, what is gained is not very helpful to the defender of (3) or PAPC. What is the point of insisting that an agent is responsible for her act A only if she has the power to agent-cause A, but the power need not be exercisable? An unexercisable power is a pretty thin reed upon which to hang moral

responsibility. What McKenna may be trying to point to is a deeper, more enduring property of the agent that is possessed even in some situations in which she cannot do otherwise. I have argued that her status as a libertarian agent can remain even in such cases and is that deeper property.

I conclude that PAPC is false because it is false that *the agent* could have refrained from causing A. If S refrains from causing A, the designer of the machine is the cause of S's not causing A. S cannot exercise a power to refrain because she can never complete the act of deciding to make what the designer of the machine considers the wrong choice. And if she cannot exercise the power to refrain she does not have the power to refrain. She fails to exercise agent causation, but failure to exercise agent causation is not the same thing as exercising agent-refraining. It is not open to her to make it not happen, although it is open that she not make it happen.

McKenna says he's willing to concede that the agent does not have it within her power to do otherwise. Nevertheless, he says she could have avoided the particular thing that she did.²⁶ Now my remarks above imply that there is a sense in which this is true. Even though *she* can't avoid agent-causing her act, the state of affairs of her agent-causing her act is avoidable. It is not avoidable through the exercise of her own power, although it is avoidable because of something about her or something she does. Since the most we have a right to say in F cases is that her agent-causing A is avoidable, and since she does agent-cause A in F cases, the strongest principle we can affirm is not (2) but (2')

(2') S agent-causes A only if S's agent-causing A is avoidable.

And we can conclude, not (3), but (3')

(3') An agent S is morally responsible for an act A only if S's agent-causing A is avoidable.

(3') does affirm an alternate possibility of a weak kind. But the avoidability of S's agent-causing A in (3') amounts to nothing more than that it might not happen, and that does not add anything to the thesis that it is causally contingent. The avoidability is not due to a power that S has, although it is due to a power she would have had and exercised if the machine had not intervened.

Another approach to connecting alternate possibilities and responsibility focuses on the conditions under which an agent is blameworthy. Michael Otsuka argues that an agent is blameworthy only if there is something she could have voluntarily done instead that would have rendered her entirely blameless.²⁷ The advantage of Otsuka's proposal is that he does not make it a requirement for blameworthiness for performing an act of a given type that the agent could have refrained from performing an act of the type for which she is blameworthy.²⁸ But she must have had an option open to her that gives her a way out of blame. Now if the machine intervenes after a necessary condition for the

choice the machine is designed to prevent, the intervention is too early to satisfy Otsuka's requirement. Suppose that we use a case like Pereboom's in which the precipitating event is a moral thought which is necessary but not sufficient for the agent to make the decision the machine wants to prevent. If the thought is a kind that always occurs involuntarily, there is nothing the agent could have done before the intervention that would have rendered her blameless, and yet we think she is to blame if the machine does not need to intervene. If the agent could have entertained the thought voluntarily, then it is true that there is something she could have done that would have precipitated the action of the machine, thereby rendering her blameless, but her blamelessness would have little to do with what she did. Surely, when the machine does not need to intervene and in ordinary situations in which there is no machine in existence we do not blame her because she could have had such a thought. She is to blame because her act was the result of a libertarian free choice, not because she might have had a moral thought which would have been necessary but not sufficient for deciding not to commit the act. Having a moral thought is not exercising the kind of power relevant to her blameworthiness. Otsuka is right that her blameworthiness is avoidable, but its avoidability is not due to her own power. Like the other defenders of alternate possibilities, he can affirm (3'), but he cannot attribute any significant form of two-way power to the agent.

In short, the challenge facing defenders of any version of PAP is that the version defended not only must not be falsified by F examples, but it has to be motivated. The problem is that the motivation is likely to go through an account of agency that utilizes something like (2). However, F cases can be interpreted as counterexamples to (2) as well as to most versions of PAP. On the other hand, (2') can be supported and used to conclude (3'), but that tells us no more than that a responsible act must be avoidable, not that the agent must have been able to do something to avoid it. Hence, even though I think it is right to focus on agent causation in identifying the ground of responsibility, it is probably wrong to think of the power of agent causation as involving alternate powers. Even though there almost always *are* alternate powers of some sort in cases of agent causation, it is doubtful that they are an intrinsic part of the power of agency itself.

IV

I have claimed it is likely that no significant version of PAP can be devised that is not falsified by Frankfurt-style counterexamples. I have also argued that it is likely that successful versions of F cases can be devised that do not require that the machine be a perfect predictor of the agent's choices, contra Kane and Widerker. Nonetheless, my real reaction to F cases does not depend upon their being literal counterexamples to PAP. The beauty of these thought experiments is that they force us to confront what it is in a situation in virtue of which we judge the agent responsible. What we see in an F situation, I believe, is that we

don't care what Black's mechanism is capable of doing because it doesn't actually do anything at all. And since the lack of alternate possibilities is tied to what the mechanism is capable of doing rather than to what it actually does, we see that we don't care whether or not the agent has alternate possibilities. Either way she is responsible. Suppose that Wierenga, McKenna, Otsuka, and others are right in thinking that she does, in fact, have alternate possibilities of some kind in Frankfurt situations. Since the point of the F cases is to show us the irrelevance of the mechanism, it does not matter how far the power of the mechanism actually goes. There is no doubt that the mechanism has *some* significant power of constraint on the agent. There are some alternate possibilities which she normally possesses in a non-deterministic world which she does not possess on the machine. But the intuition of most people who hear of F cases is that the machine has no effect at all on her level of responsibility or on that for which she is responsible.

What the F cases do is to make us think about responsibility in a novel way, a way that calls our attention to the difference between the features of a responsibility situation that are salient and those that are not. What we see is that Black's machine plays no role in our judgment about Jones's responsibility. We do not evaluate Jones any differently with the machine than without it, nor would we evaluate Jones any differently if, *per impossibile* (perhaps), the machine became a perfect predictor. Nor need we work to identify the respect in which the agent retains an alternative since our judgment of her responsibility does not depend upon the success of that search, nor do we attribute responsibility to her in any different way once we think we have identified the precise sense in which she has alternate possibilities.

V

The intuition that the agent is responsible in F cases does not refute the libertarian thesis LTR, nor does it falsify (3'). Whatever reason the libertarian has for thinking a person agent-causes his act in an ordinary non-Frankfurt-style situation applies just as well to a Frankfurt situation, and if the former act is causally contingent, so is the latter. The avoidability of the one is the same as the other. This means that F cases do not refute the libertarian position. Furthermore, they do not give any support to determinism. That is because a certain kind of libertarian will say that the reason why the agent is responsible in F cases is that she agent-causes the act and agent-causation requires that the act be causally contingent. Notice that this point is not the same as the first. The first and weaker point is simply that the libertarian is not forced to give up LTR when she accepts that the agent is responsible in F cases. Libertarianism is compatible with the rejection of PAP. The second and stronger point is that she will accept that the agent is responsible in F cases *for libertarian reasons*. She accepts F cases because the libertarian condition still obtains in those cases. She can reject PAP, then, because doing so has no effect on the deeper libertar-

ian intuition operative in these cases. Of course, the determinist will see F cases differently, but that is just because the determinist and the non-determinist see every case differently. The point is that F cases do not give the libertarian any more reason to accept determinism than ordinary cases do.

But the libertarian can go farther. She can capitalize on F cases to support non-determinism. Ironically, one of Frankfurt's own remarks provides the argument. Frankfurt says:

Even though the person was unable to do otherwise...it may not be the case that he acted as he did *because* he could not have done otherwise. Now if someone had no alternative to performing a certain action but did not do it because he was unable to do otherwise, then he would have performed exactly the same action even if he *could* have done otherwise. The circumstances that made it impossible for him to do otherwise could have been subtracted from the situation without affecting what happened or why it happened in any way. Whatever it was that actually led the person to do what he did, or that made him do it, would have led him to do it or made him do it even if it had been possible for him to do something else instead.²⁹

In a deterministic world the agent's act is caused by events outside the agent, so in one clear sense of "because," the act occurs because of those events. He acts because he could not do otherwise. Hence, he does not satisfy the condition for responsibility in the passage just quoted. Frankfurt remarks that it is important that there is no morally relevant difference between the situation in which the agent has alternate possibilities and one in which they have been eliminated by a Frankfurt-type mechanism. This means that the comparison of the F situation with one in which the agent *can* do otherwise is important to the intuition that the agent is responsible on Frankfurt's own account. He says that the circumstances that make it the case that the agent could not have done otherwise could have been subtracted from the situation without affecting what happened or why it happened in any way. But notice that this implies that what the agent would have done *if* he had been able to do otherwise is relevant to his responsibility in the actual situation since the subtraction of the conditions that make it the case that he cannot do otherwise turns the situation into one in which he *can* do otherwise. A curious feature of Frankfurt's point here is that as a determinist he cannot claim that a possible world in which the agent can do otherwise is close to the actual world, so the comparison of the actual situation with one in which he can do otherwise is a comparison of the actual world with a very distant world, one in which the basic structure of natural law differs from that which actually obtains.

On the page following the passage just quoted Frankfurt offers an amendment to his claim and a new principle of alternate possibilities which he can support: "a person is not morally responsible for what he has done if he did it *only because* he could not have done otherwise." [emphasis added].³⁰ Notice that this is not simply a clarification or innocuous addition to the passage quoted above; it is a different and weaker position. To say that an agent is not responsible for an act if he did it only because he could not have done otherwise

leaves open the possibility that when he is responsible he did the act *partly* because he could not have done otherwise. In that case, if the circumstances that make it the case that the agent could not have done otherwise were subtracted from the situation, it presumably *would* affect what happened or why it happened in some way, contrary to what Frankfurt says in the quoted passage. As a determinist, Frankfurt must countenance dependence of the agent's act on the conditions that make it the case that he cannot do otherwise since on his view those conditions are the cause of the act. But it is worth noting that this position requires a partial retraction of his claim that in a responsibility situation the conditions that make it the case that the agent cannot do otherwise are irrelevant to why he did what he did.

Libertarians will make a different judgment. We think that Frankfurt was right in his point that the agent is responsible because what he does in an F situation is relevantly similar to what he would have done if he been able to do otherwise. But it is also important that Black's meddling is a quirk that in the normal course of events would not have entered the picture at all. The point could be put as follows: It is only an accident that Black exists, and if he had not existed the agent would have had alternate possibilities. And if he had had alternate possibilities he would have done the very same thing in the same way. He is, therefore, just as responsible as he would have been if he had had alternate possibilities. To say otherwise is to permit the agent too great a degree of positive moral luck. He can't get off the moral hook *that* easily.

Therefore, we get the conclusion Frankfurt wants precisely because of our interest in what would have happened if Jones had been able to do otherwise, along with our inclination to think that but for the short-lived existence of an inoperative device, Jones would have actually been able to do otherwise. If the universe is causally determined, however, the second condition is not satisfied and the first is logically problematic since any proposition of the form *If Jones had been able to do otherwise, then...* has an antecedent that contradicts the laws of nature in such a world.

At the beginning of this paper I remarked that it is tempting to bypass the difficult concept of what it is to make an act happen and to take various digressions. There is nothing wrong with digressions as long as they illuminate our target concept and we eventually get back to it. Talk of counterfactual conditions for the attribution of responsibility is a digression. It can help illuminate the idea of causing or making something happen, but eventually we have to get back to the idea those conditions aimed to elucidate, and that is the idea of causing an act, making it happen. Once we have identified who or what makes an act happen we have identified the potential bearer of responsibility for it. In Frankfurt cases in a non-deterministic world the agent makes her act happen in as ultimate a sense as you like; in a deterministic world it is not the agent that ultimately makes her act happen. Frankfurt cases do nothing to lead us to rescind the view that this is a significant difference. What they do show is that PAP is a false path. The presence of alternate possibilities may be a reliable sign of the presence of the agency needed for responsibility, but it is not necessary for it.

I suspect that the moral of Frankfurt cases can be generalized to apply to many properties that have counterfactual conditions for their application. If I am right that such conditions are usually proposed because they are a sign of what we think is really important, it can happen that they fail because of the presence of a counterfactual manipulator even when the target property still obtains. For example, several well-known definitions of knowledge include counterfactual conditions. In another place I have argued that epistemic Frankfurt-style cases show that these conditions can fail for reasons parallel to the reasons for the failure of PAP.³¹ In general, what might have happened is often a good indication of what *did* happen, but it is not essential to it.³² PAP is a thesis about what might have happened; causal determinism and non-determinism are theses about what did happen.

Notes

1. I use “act” for the event for which an agent is primarily responsible. It may be an act of will, an overt act, or even a firing of neurons, provided that it is the kind of event for which we attribute responsibility to the agent in the basic sense.
2. The term “accidentally necessary” comes from Ockham’s use of the idea of necessity *per accidens*, a form of temporal necessity. The necessity is said to be “accidental” because the event it is not necessary at all times but only after the event occurs. See *Predestination, God’s Foreknowledge, and Future Contingents*, trans. by Marilyn McCord Adams and Norman Kretzmann (Indianapolis: Hackett, 1969).
3. I made this point in *The Dilemma of Freedom and Foreknowledge* (Oxford University Press, 1991), p. 119, and in “Foreknowledge and Freedom,” *Companion to Philosophy of Religion*, edited by Philip L. Quinn and Charles Taliaferro (Blackwell, 1997). For a recent proponent of this position see David Hunt, “On Augustine’s Way Out,” *Faith and Philosophy* 16, #1 (January 1999), and “Moral Responsibility and Unavoidable Action,” forthcoming in *Philosophical Studies*.
4. Harry Frankfurt, “Alternate Possibilities and Moral Responsibility,” *Journal of Philosophy* 66 (December 1969), pp. 828–839; reprinted in *Moral Responsibility*, edited by John Martin Fischer (Ithaca, N.Y.: Cornell University Press, 1986), pp. 143–152. Subsequent page references will be taken from the Fischer reprint.
5. I have argued this before in *The Dilemma of Freedom and Foreknowledge*, chap. 6.
6. This adaptation of Frankfurt’s example using a neurological device is similar to some of the cases described by John Martin Fischer. An early use of this type of example appears in “Responsibility and Control,” *Journal of Philosophy* 89 (January 1982), pp. 24–40.
7. “Alternate Possibilities and Moral Responsibility,” in Fischer, p. 149, note 3. Frankfurt does not conclude the truth of determinism in this paper, but he thinks that PAP is the only thing standing in the way of the acceptance of the compatibility of determinism and responsibility. It can then be argued that once compatibilism is accepted, there is nothing blocking the acceptance of determinism.
8. Robert Kane, *The Significance of Free Will* (N.Y.: Oxford University Press, 1996); David Widerker, “Libertarianism and Frankfurt’s Attack on the Principle of Alternate Possibilities,” *Philosophical Review* 104 (April 1995), pp. 247–261, and “Libertarian Freedom and the Avoidability of Decisions,” *Faith and Philosophy* 12 (January 1995).

9. "Rescuing Frankfurt-Style Cases," *Philosophical Review* vol. 107, #1 (January 1998), pp. 97–112.
10. P. 103.
11. Mele and Robb, footnote 11.
12. *ibid.*
13. Eleonore Stump, "Libertarian Freedom and the Principle of Alternate Possibilities," in *Faith, Freedom, and Rationality*, edited by Jeff Jordan and Daniel Howard-Snyder (Lanham, Md: Rowman and Littlefield, 1996), pp. 73–88.
14. "Stumping for Widerker," *Faith and Philosophy*, vol. 16, #1 (January 1999), pp. 83–89.
15. This interpretation of Stump is confirmed in "Dust, Determinism, and Frankfurt: A Reply to Goetz," *Faith and Philosophy* 16, #3 (July 1999), pp. 413–422. In this paper Stump makes her position clearer. She argues that all it takes to falsify PAP is a device that interrupts the neural sequence correlated with an act of will after it begins but before it ends. The precise nature of the relationship between the act of will and the agent's brain states does not matter for the argument, and it is even possible that they are identical. This scenario falsifies PAP, but it does not require that an act of will is determined, even if it is identical with a neurological event. Stump maintains that Aquinas also rejected both PAP and causal determinism.
16. See Eleonore Stump, "Intellect, Will, and the Principle of Alternate Possibilities," in *Christian Theism and the Problems of Philosophy*, edited by Michael D. Beaty (University of Notre Dame Press, 1990). In this paper Stump describes an episode from Dostoevsky's *The Possessed* in which Verkhovensky is prepared to intervene with a surgical procedure to insure that Fedya murders the Lebyatkins if there is "any chance" that he won't do so when offered a bribe (p. 255). This implies that the intervention would occur after a causally necessary but insufficient condition for deciding not to commit the murder has occurred. Of course, Verkhovensky is not in a position to identify causally necessary conditions with 100% accuracy, but the point is sound since we can easily imagine how someone could identify a causally necessary condition in a non-deterministic universe.
17. "Alternate Possibilities and Moral Responsibility," pp. 148–149.
18. "Alternate Possibilities and Causal Histories," this volume.
19. "Of Active Power in General," Essay 1 of *Essays on the Active Power of the Mind* (MIT Press, 1969), p. 35.
20. In fact, Reid thought that the idea of a cause derives from that of a power and concluded that in the strict sense only conscious beings are causes. See "Of the Liberty of Moral Agents," chap. 2, "Of the Words Cause and Effect, Action and Active Power," in *Essays on the Active Powers of the Human Mind*, (MIT Press, 1969), pp. 267ff. What is especially interesting about this view is that it makes agency and the possession of alternate possibilities the prior idea. Causation and contingency are derivative.
21. Edward Wierenga, *The Nature of God: An Inquiry into Divine Attributes* (Ithaca: Cornell University Press, 1989), p. 85. An earlier version of the same type of response is given by Margery Bedford Naylor in "Frankfurt on the Principle of Alternate Possibilities," *Philosophical Studies* 46 (1984), pp. 249–258.
22. Wierenga says his view is similar to that of Reid in the footnote on p. 85.
23. Michael McKenna, "Alternate Possibilities and the Failure of the Counterexample Strategy," *Journal of Social Philosophy*, vol. 28, #3 (Winter 1997), p. 75.

24. An exception is causal over-determination scenarios in which the machine is set to cause a certain choice but the agent causes the same choice by her own power at the same time.
25. McKenna, p. 81.
26. McKenna, p. 77.
27. Michael Otsuka, "Incompatibilism and the Avoidability of Blame," *Ethics* 108 (July 1998), pp. 685–701.
28. Otsuka makes this explicit on p. 690.
29. Frankfurt, pp. 150–151.
30. Frankfurt, p. 152.
31. In "Must Knowers Be Agents?", presented at conference on Virtue and Duty in Epistemology, Santa Barbara, Calif, Nov. 13, 1999; forthcoming in *Virtue Epistemology*, Oxford University Press.
32. I am grateful to David Hunt, Andrew Eshleman, and Daniel Speak for comments on an earlier draft of this paper.

Copyright of Philosophical Perspectives is the property of Wiley-Blackwell and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.