

ROBUST FLICKERS OF FREEDOM*

BY MICHAEL ROBINSON

Abstract: This essay advances a version of the flicker of freedom defense of the Principle of Alternative Possibilities (PAP) and shows that it is invulnerable to the major objections facing other versions of this defense. Proponents of the flicker defense argue that Frankfurt-style cases fail to undermine PAP because agents in these cases continue to possess alternative possibilities. Critics of the flicker strategy contend that the alternatives that remain open to agents in these cases are unable to rebuff Frankfurt-style attack on the grounds that they are insufficiently robust (that is, morally significant in a way that could ground ascriptions of moral responsibility). Once we see that omissions are capable of constituting robust alternatives, even when they are not intentional, it becomes clear that agents in these cases do indeed possess robust alternative possibilities—alternatives that are ineliminable from cases of this sort. The upshot is that Frankfurt-style cases are theoretically incapable of providing us with good grounds for rejecting PAP.

KEY WORDS: PAP, moral responsibility, alternative possibilities, robust, alternatives, Frankfurt, flicker, omissions, free will

I. INTRODUCTION

According to the Principle of Alternative Possibilities (PAP), a person is morally responsible for what he has done only if he could have done otherwise. The greatest challenge currently facing this seemingly intuitive principle comes from a class of cases developed by Harry Frankfurt, commonly known as Frankfurt-style cases.¹ The original intent in offering these cases was to falsify PAP by providing possible scenarios in which agents appear to be morally responsible for performing some action that, owing to the special circumstances in which they find themselves, they were unable to avoid performing. One way to defend PAP against falsification by cases of this sort is to call into question the claim that agents in these cases are genuinely morally responsible for what they do (at least with respect to those actions whose performance is unavoidable for them). Unless the actions agents could not have avoided performing are also ones for which they are morally responsible, these cases will fail to falsify the view that the ability to do otherwise is a necessary condition for

* This essay benefited enormously from many discussions of this issue with Al Mele, Michael McKenna, and Randy Clarke. I am grateful to the other contributors to this volume—as well as to Carolina Sartorio, David Shoemaker, and Dana Nelkin—for their many helpful questions and comments on the essay. I also wish to thank an anonymous referee for this journal for the most thorough and thoughtful feedback I have ever received from a referee.

¹ Harry Frankfurt, "Alternate Possibilities and Moral Responsibility," *Journal of Philosophy* 66 (1969): 829–39.

moral responsibility. Alternatively, one could grant that agents are morally responsible for what they do in these cases but deny that they could not have done otherwise. According to this line of argument, even though Frankfurt-style cases manage to significantly *restrict* agents' ability to do otherwise, they do not *eliminate* it in the way that falsifying PAP would require. This latter strategy has become known as the *flicker of freedom defense* of alternative possibilities.²

Elsewhere, I have defended a version of the flicker strategy that focuses on the fact that agents in these cases do what they do *on their own*, just as they would have if they had been in ordinary (non-Frankfurtian) circumstances, which is something regarding which they could have done otherwise.³ Even though it is impossible for these agents not to perform the actions they do, since it remains up to them whether to perform these actions on their own, the fact that they are morally responsible for them poses no threat to PAP. Critics of the flicker strategy have argued that the alternative possibilities that remain open to agents in Frankfurt-style cases—including the ability to omit performing the relevant actions on their own—are not adequate to the task of rebuffing Frankfurt-style attack. In order to neutralize the threat facing PAP, they maintain, it is not enough for agents in these cases to have alternative possibilities; the alternatives must be *robust*. That is, they must be morally significant in a way that helps to account for the agents' moral responsibility for what they actually do. Insofar, then, as Frankfurt-style cases are able to restrict agents' ability to do otherwise to alternatives that are insufficiently robust in this sense, they will be successful in undermining PAP (assuming these agents are morally responsible for what they do). Eliminating *all* alternative possibilities, including those that are not robust, they argue, is not required.

Over the fifty years since the introduction of these cases, many philosophers have been persuaded to reject the once-dominant view that moral responsibility requires the ability to do otherwise. For some, this is because they have been convinced that (at least some versions of) Frankfurt-style cases successfully falsify PAP. For others, this is because they have thought that, even though these cases do not strictly falsify PAP, they succeed in showing that moral responsibility does not require alternative possibilities that are distinctly robust, which is enough to cast serious doubt on our reasons for thinking that moral responsibility requires any alternatives whatsoever. My aim in this essay is to show that all this has been a mistake. Even if all that is needed to demotivate (though not falsify) PAP in this way are Frankfurt-style scenarios that manage to eliminate all robust alternative possibilities, there are no (nor can there be) cases of this sort.

² The term “flickers of freedom” was coined by John Martin Fischer to refer to those alternatives that remain open to agents in Frankfurt-style cases. See John Martin Fischer, *The Metaphysics of Free Will: An Essay on Control* (Cambridge: Blackwell Publishers, 1994), 137–47.

³ See Michael Robinson, “Modified Frankfurt-type Counterexamples and Flickers of Freedom,” *Philosophical Studies* 157 (2012): 177–94.

So long as it remains up to agents in these cases whether to do on their own the things for which they are morally responsible, they will continue to possess robust alternative possibilities. And since having it be up to the agents whether they perform the relevant actions on their own is a crucial feature of Frankfurt-style cases—one which cannot be eliminated without undermining the judgment that they are genuinely morally responsible for what they do—the demand for robust flickers of freedom is one that this particular version of the flicker strategy will always be able to meet.

II. FRANKFURT-STYLE CASES AND FLICKERS OF FREEDOM

A genuine counterexample to PAP will be a case in which an agent is morally responsible for performing some action she could not have done otherwise than perform.⁴ Because PAP is generally understood to be a thesis about basic (that is, direct, or nonderivative) moral responsibility, the relevant actions in these cases tend to be mental acts of deciding, which are commonly taken to be the things for which we are basically morally responsible and from which our responsibility for overt bodily actions (and their consequences) derives.⁵ To ensure that agents cannot do otherwise than decide to act as they actually do, typical Frankfurt-style cases include the presence of a counterfactual intervener, someone who stands ready and able to force the agent to decide to act in the desired way in the event that the agent is not going to do so on her own. Crucially, however, the counterfactual intervener's presence is unknown to the agent and plays no role in the agent's actual behavior, lest it undermine the intuition that the agent is morally responsible for what she does.

The tricky part for cases of this sort comes in identifying the precise conditions under which the counterfactual intervener would intervene. If the counterfactual intervener has to wait to see how the agent actually decides, then the agent retains the ability to do otherwise, and PAP is unthreatened. In order to guarantee that the agent will decide to act in the desired way without actually compelling the agent's behavior, the

⁴ Counterexamples to PAP need not be restricted to cases involving actions. Since PAP is generally taken to apply to both actions and omissions, a case in which an agent is genuinely morally responsible for omitting to perform some action she could not have done otherwise than omit to perform would also serve as a counterexample to this principle. Here, following custom, I focus on cases involving moral responsibility for actions (as opposed to omissions) merely for the sake of simplicity.

⁵ An agent is basically morally responsible for some action she performs when her responsibility for that action is not inherited from (or in virtue of) her responsibility for other actions. Moral responsibility that is inherited in this way is derivative, or indirect. See Alfred R. Mele, *Free Will and Luck* (New York: Oxford University Press, 2006), 86. For arguments that agents can be basically morally responsible for things other than certain mental actions, see Robert Adams, "Involuntary Sins," *Philosophical Review* 94 (1985): 3–31, and Angela Smith "Responsibility for Attitudes: Activity and Passivity in Mental Life," *Ethics* 115 (2005): 236–71, and "Control, Responsibility, and Moral Assessment," *Philosophical Studies* 138 (2006): 367–92.

counterfactual intervener must have some way of knowing in advance what the agent will do if left to act on her own. Toward this end, traditional Frankfurt-style cases posit the existence of a *prior sign*—something beyond the agent's voluntary control (such as a blush or a cluster of neural firings) that indicates how the agent would freely act, were no intervention to occur.

The trouble with including indicators of this sort is that they are capable of doing the work set out for them in traditional Frankfurt-style cases only by presupposing a deterministic relation between an agent's free action and its causal antecedents—a relation incompatibilists flatly reject.⁶ This problem can be seen as posing the following dilemma. Either the occurrences of these prior signs (or that which they indicate) make it merely probable that agents in these cases will decide to act in the desired ways, or they deterministically lead to the agents so deciding. If the occurrences of these signs make it merely probable that the agents will decide to act in the desired ways, then it remains causally open to them not to decide to act in this way after all. Insofar as this is the case, however, the agents will continue to possess alternative possibilities with respect to their actions, and so their bearing moral responsibility for these actions is perfectly compatible with PAP. If, on the other hand, the occurrences of these signs (or that which they indicate) will deterministically lead to the agents deciding to act in the desired ways, then it will no longer remain causally open to the agents not to decide to act as they do. But since having one's actions causally determined precludes one from being morally responsible for them, according to incompatibilists, the claim that agents in these cases are nevertheless morally responsible for what they do cannot be asserted without begging any questions.⁷ Either way, then, traditional Frankfurt-style scenarios fail to constitute genuine counterexamples to PAP.⁸

This particular line of objection has prompted a number of attempts to modify traditional Frankfurt-style cases in ways that will enable them to do the work set out for them by Frankfurt without employing the use

⁶ See David Widerker, "Libertarianism and Frankfurt's Attack on the Principle of Alternative Possibilities," *Philosophical Review* 104 (1995): 247–61, Robert Kane, *The Significance of Free Will* (New York: Oxford University Press, 1996), and Carl Ginet, "In Defense of the Principle of Alternative Possibilities: Why I Don't Find Frankfurt's Argument Convincing," *Philosophical Perspectives* 10 (1996): 403–17.

⁷ For arguments to the effect that there is in fact nothing illicit in the way traditional Frankfurt-type counterexamples make use of prior signs, see John Martin Fischer, "Recent Work on Moral Responsibility," *Ethics* 110 (1999): 93–139, and Ishtiyaque Haji and Michael McKenna, "Dialectical Delicacies in the Debate about Freedom and Alternative Possibilities," *Journal of Philosophy* 101 (2004): 299–314.

⁸ It should be noted that, relying as it does on the decidedly incompatibilist intuition that determinism and moral responsibility are mutually exclusive, this is not a line of argument available to compatibilist defenders of PAP. Compatibilists who wish to resist the conclusion of Frankfurt's argument cannot do so on the grounds that it is question-begging to assert that agents can be morally responsible for their actions despite their being causally determined.

of (determinative) prior signs.⁹ Here is a modified Frankfurt-style case offered by Alfred Mele and David Robb:

Our scenario features an agent, Bob, who inhabits a world at which determinism is false . . . At t_1 , Black initiates a certain deterministic process P in Bob's brain with the intention of thereby causing Bob to decide at t_2 (an hour later, say) to steal Ann's car. The process, which is screened off from Bob's consciousness, will deterministically culminate in Bob's deciding at t_2 to steal Ann's car unless he decides on his own at t_2 to steal it or is incapable at t_2 of making a decision (because, for example, he is dead by t_2) . . . The process is in no way sensitive to any 'sign' of what Bob will decide. As it happens, at t_2 Bob decides on his own to steal the car, on the basis of his own indeterministic deliberation about whether to steal it, and his decision has no deterministic cause. But if he had not just then decided on his own to steal it, P would have deterministically issued, at t_2 , in his deciding to steal it. Rest assured that P in no way influences the indeterministic decision-making process that actually issues in Bob's decision.¹⁰

Rather than rely on the presence of some prior sign concerning what Bob will freely do if left to act on his own, Mele and Robb employ a particular kind of causal preemption to guarantee that an intervention will take place if and only if Bob does not decide on his own to steal Ann's car (and is capable of making a decision at that time). The significance of the preemption's occurring precisely at t_2 (the moment of action) is that it ensures that at no point leading up to and including t_2 is it possible for Bob to do otherwise than decide at t_2 to steal Ann's car. Because the process Black initiates in Bob's brain is insensitive to anything Bob does prior to t_2 , Mele and Robb avoid the need to controversially presuppose a deterministic relation between Bob's action and any of its causal antecedents. And because the process initiated by Black is causally preempted by Bob at t_2 , and so plays no causal

⁹ See, for example, Isthiyaque Haji, *Moral Appraisability* (New York: Oxford University Press, 1998), David P. Hunt, "Moral Responsibility and Unavoidable Action," *Philosophical Studies* 97 (2000): 195–227, Michael McKenna, "Robustness, Control, and the Demand for Morally Significant Alternatives: Frankfurt Examples with Oodles and Oodles of Alternatives," in Widerker and McKenna, eds., *Moral Responsibility and Alternative Possibilities* (Aldershot, UK: Ashgate, 2003), 201–18, Alfred Mele and David Robb, "Rescuing Frankfurt-Style Cases," *The Philosophical Review* 107 (1998): 97–112, Mele and Robb, "Bbs, Magnets, and Seesaws: The Metaphysics of Frankfurt-Style Cases," in Widerker and McKenna, *Moral Responsibility and Alternative Possibilities*, 127–38, Derk Pereboom, "Alternative Possibilities and Causal Histories," *Philosophical Perspectives* 14 (2000): 119–38, and Eleonore Stump, "Libertarian Freedom and the Principle of Alternative Possibilities," in Daniel Howard-Snyder and Jeff Jordan, eds., *Faith Freedom and Rationality* (Totowa, NJ: Rowman and Littlefield, 1996), 73–88.

¹⁰ Mele and Robb, "Rescuing Frankfurt-Style Cases," 101–102 (footnotes omitted).

role in Bob's actual decision to steal Ann's car, there appears to be no reason to deny that Bob can be genuinely morally responsible for this decision.

Elsewhere I have argued at length that even if the kind of modified Frankfurt-style case just described successfully avoids the dilemma facing traditional prior-sign cases, it nevertheless fails to falsify PAP. I will not rehearse all that I said there, but the core of my argument can be put as follows.¹¹ It is generally accepted by those on both sides of this debate that PAP is properly understood as a thesis about *basic* (that is, direct, or nonderivative) moral responsibility. Indeed, it is only when PAP is understood in this way that the "need" for Frankfurt-style examples arises in the first place. Were it taken to apply to instances of derivative (or indirect) moral responsibility as well, falsifying this principle would not require so much talk of nefarious neuroscientists and their science-fictional devices; standard drunk-driver cases would do the trick.¹² So understood, then, PAP claims that a person is basically morally responsible for what he has done only if he could have done otherwise. Importantly, it is the referent of the antecedent which fixes that of the consequent, not the other way around. That is to say, that than which a person must have been able to do otherwise, according to PAP, is determined by that (whatever it is) for which he is basically morally responsible.

In order for a Frankfurt-style case to count as a genuine counterexample to PAP, it must be one in which an agent could not have done otherwise than that for which he is basically morally responsible. The example provided by Mele and Robb is not such a case. Even supposing it is true both that Bob could not have done otherwise than decide to steal Ann's car and that the deterministic process initiated by Black played no causal role in bringing about Bob's decision, it is not deciding to steal Ann's car for which Bob is basically morally responsible.¹³ This is not

¹¹ The remainder of this section is based on Robinson, "Modified Frankfurt-type Counterexamples."

¹² See Mele, *Free Will and Luck*, 84–86.

¹³ Some philosophers have argued that the particular modified Frankfurt-style scenario offered by Mele and Robb, insofar as it involves actually blocking off every alternative to Bob's deciding to steal Ann's car, is one in which the agent's action appears to be causally determined. See Robert Kane, "Responses to Bernard Berofsky, John Martin Fischer and Galen Strawson," *Philosophy and Phenomenological Research* 60 (2000): 157–67, Derk Pereboom, *Living Without Free Will* (Cambridge: Cambridge University Press, 2001), and Laura Ekstrom, "Libertarianism and Frankfurt-Style Cases," in Robert Kane, ed., *The Oxford Handbook of Free Will* (Oxford and New York: Oxford University Press, 2002), 309–22. Others have challenged the coherence of the occurrent preemption on which their case relies. See David Widerker, "Theological Fatalism and Frankfurt-Counterexamples to the Principle of Alternative Possibilities," *Faith and Philosophy* 17 (2000): 249–54. Here I set these worries aside in order to show that, even if we grant that these worries can be allayed, PAP is able to accommodate cases of this sort. For their reply to these concerns, as well as a detailed account of the type of mechanism they have in mind, see Mele and Robb, "Bbs, Magnets, and Seesaws."

to deny, I hasten to add, that Bob is genuinely morally responsible for deciding to steal Ann's car, only that he is *basically* morally responsible for this. It is deciding *on his own* to steal Ann's car for which Bob is basically morally responsible and from which his moral responsibility for so deciding *simpliciter* derives.¹⁴

One reason for thinking that Bob is basically morally responsible not for deciding *simpliciter* to steal Ann's car but for so deciding on his own is implicit in Mele and Robb's case itself. Like virtually all other Frankfurt-style scenarios (both in their traditional and modified forms), Mele and Robb's case focuses on a mental act of deciding rather than an overt bodily action. This is no accident. Instead, it is the product of two widely shared assumptions. The first is that PAP is properly understood as a thesis about basic moral responsibility. The second is that responsibility for overt bodily actions (such as raising one's arm) derives from responsibility for certain mental actions (such as deciding to raise one's arm) much in the same way that responsibility for the consequences of one's actions derives from one's responsibility for the actions themselves. On this view, if one is morally responsible both for deciding to raise one's arm and for raising one's arm, it is the former for which one is basically morally responsible, whereas one is only derivatively responsible for the latter. This variation in the kind of responsibility agents bear for certain mental actions as compared to overt bodily actions is best explained by a difference in control. The successful performance of even simple overt actions like raising one's arm is much more vulnerable to environmental impediments than the successful performance of simple mental actions like deciding to raise one's arm. It is because agents have increased control over certain

¹⁴ At this point, some readers might find themselves asking just what it means here for an agent to decide on his own to A and how this is related to an agent's deciding to A *simpliciter*. Are these intended to designate distinct actions or events? Are they distinct act *types*, of which there could be indefinitely many act *tokens*? Is deciding on one's own to A just a *way* of deciding to A and is it the *way* the agents decide to A for which they are basically morally responsible? I think there are plausible arguments that can be offered in favor of each of these ways of viewing the relationship between deciding on one's own to A and deciding to A *simpliciter*—some of which can be found in this essay, others of which I have defended elsewhere. (See Robinson, "Modified Frankfurt-type cases," 184–85.) Because I think the central line of argument advanced in this essay is flexible enough to accommodate different ways of answering these questions, I will not attempt to settle this issue here. As for the meaning of the "on one's own" locution, the first thing to say is that this is a question best posed to Frankfurt and the other defenders of Frankfurt-style cases. Recall that they are the ones who have been at such pains to construct elaborate scenarios precisely to ensure that *either* the agents will *decide on their own* to act in the desired ways or else they will be forced to so decide. They are also the ones who have argued in support of the agents being morally responsible for what they do in these cases primarily by appealing to the fact that the agents in these cases made the relevant decisions *on their own*, just as they would have had the special ensuring conditions not been present. Nevertheless, what I mean (and what I presume they mean) in saying that an agent decides on his own to A is simply that he decides to A and that he does so without being causally forced or interfered with (in the way that he would have been in the alternative scenario in a Frankfurt-style case).

mental actions, such as deciding, that they are basically morally responsible for these, rather than for overt bodily actions.¹⁵

Things are a bit different in Frankfurt-style scenarios. Looking at Mele and Robb's case, we can see that Bob has the same degree of control over whether he will steal Ann's car as he does over whether he will *decide* to steal Ann's car—namely, none. It is guaranteed that Bob will decide to steal Ann's car and that he will steal it, and there is nothing he can do to prevent either of these things from occurring. Even in this case, however, whether Bob will decide on his own to steal Ann's car, and whether he will steal Ann's car on his own, remains entirely up to him. Thus, the same reasons typically taken to support the view, implicit in Frankfurt-style examples, that one's responsibility for performing some overt action A derives from one's responsibility for performing some mental action, such as deciding to A, equally suggest that, at least in Frankfurt-style cases, one's responsibility for deciding to A derives from one's responsibility for so deciding on one's own, for which one is basically morally responsible.¹⁶

So, even if Bob is genuinely morally responsible for deciding to steal Ann's car, his responsibility for this derives from his responsibility for *deciding on his own* to steal Ann's car. Of course, since deciding on his own to steal Ann's car is something it was within Bob's power not to have

¹⁵ An anonymous referee for this journal suggests that another way to account for this difference in responsibility could be to focus on the *directness* of control (rather than on susceptibility to external interference). Recall that, following Mele, I say that an agent is indirectly (or derivatively) responsible for something when that agent is responsible for it in virtue of being responsible for something else, and an agent is directly (or basically) responsible for something when the agent is responsible for it but not in virtue of being responsible for anything else (see n. 5). Analogously, we might say that an agent has indirect control over something when she controls it *by* controlling something else, and an agent has direct control over something when she controls it *but not by* controlling anything else. One who is willing to accept this would then be in a position to argue that agents are directly responsible for deciding on their own to A (rather than for so deciding *simpliciter*) on the grounds that (i) agents are directly morally responsible only for that which is in their direct control and that (ii) just as agents control whether they will A by controlling their decisions, so too do they control whether they will decide to A by controlling whether they will so decide on their own.

¹⁶ Some might wonder whether proponents of this strategy wish to maintain that, as a general matter and not just in Frankfurt-style cases, it is deciding on their own to A (rather than deciding to A) for which agents are basically morally responsible. Indeed, they might worry that there is something suspiciously ad hoc about the proposal otherwise. For my own part, I am inclined to think that, on balance, the reasons offered in support of this view favor holding that it is normally the case that agents' responsibility for A-ing derives from their being morally responsible for A-ing on their own. Importantly, however, this is not a stand that is required in order to adopt the line of argument advanced in this essay. It remains open to those who might wish to maintain that agents are sometimes basically morally responsible not only for A-ing on their own but also for A-ing *simpliciter*. Moreover, those convinced to adopt this stance on the grounds that the control agents have in normal (non-Frankfurt) cases over whether they will A on their own is matched by the control they have over whether they will A *simpliciter* can do so in a principled way that is well poised to defend against charges of being ad hoc.

done, the judgment that he is (basically) morally responsible for doing this is consistent with the truth of PAP.¹⁷

It is important to note that none of the foregoing hinges upon anything peculiar to Mele and Robb's particular example. Rather, it capitalizes on two features that are essential to any Frankfurt-style scenario: the agent performs the relevant action on his own (without being forced), and it was possible for the agent not to have done so. The fact that agents in these cases act on their own, just as they would have if no ensuring mechanism had been in place, is what is supposed to drive the intuition that they are morally responsible for what they do, without which there is no threat to PAP. Furthermore, omitting to perform the relevant action on one's own is a genuine alternative possibility in all Frankfurt-style scenarios (and one which in modified cases lies safely within the scope of the agents' direct voluntary control). Indeed, it is precisely because it is possible for agents in these cases not to perform the relevant actions on their own that the ensuring mechanisms are there to begin with; otherwise these mechanisms would be superfluous. Since this alternative possibility is ineliminable, the particular version of the flicker strategy employed here can be brought to bear against any modified Frankfurt-style example. Insofar as it is successful, the upshot is not only that PAP has yet to be falsified by any of the Frankfurt-style cases currently on offer but that cases of this kind are theoretically incapable of falsifying PAP.

III. WHY MUST FLICKERS BE ROBUST?

The chief objection to the flicker strategy comes from John Martin Fischer. The problem with the various versions of this approach, he argues, is that the alternative possibilities on which they focus are not sufficiently *robust*:

I suggest that it is not enough for the flicker theorist to analyze the relevant range of cases in such a way as to identify an alternative possibility. Although this is surely a first step, it is not enough to establish the flicker of freedom view, because what needs also to be shown is that these alternative possibilities *play a certain role* in the appropriate understanding of the cases. That is, it needs to be shown that these alternative possibilities *ground* our attributions of moral responsibility.¹⁸

¹⁷ I take it to be a comparative virtue of this particular version of the flicker strategy that it does not commit its proponents to maintaining that agents in Frankfurt-style cases are morally responsible only for A-ing on their own and not at all for A-ing *simpliciter*. (Cf. Marjory Bedford Naylor, "Frankfurt on the Principle of Alternate Possibilities," *Philosophical Studies* 46 [1984]: 249–58.) Nor does it require reformulating PAP in a way that is prone to spur complaints of gerrymandering or diminish its intuitive force (by making it overly complex). (See, for example, Seth Shabo, "Flickers of Freedom and Modes of Action: A Reply to Timpe," *Philosophia* 35 [2007]: 63–74. Cf. Kevin Timpe, "A Critique of Frankfurt-Libertarianism," *Philosophia* 34 [2006]: 189–202.)

¹⁸ Fischer, *The Metaphysics of Free Will*, 140.

According to Fischer, neutralizing the threat to PAP from Frankfurt-style cases requires not only that agents in these cases could have done otherwise than that for which they are (basically) morally responsible but that the existence of these alternatives is able to help account for why agents would, in virtue of possessing these alternatives, be morally responsible for what they actually do. As he puts it, “the flicker theorist must claim that the addition of the sort of possibility he has identified would transform a case of lack of responsibility into one of responsibility.”¹⁹

Fischer grants that it remains open to agents in Frankfurt-style cases not to decide on their own to act as they do. The reason this alternative is not sufficiently robust to figure into an account of why these agents are morally responsible for what they do, he maintains, is that whether this alternative is realized depends on whether the agents display some prior sign indicating that they are not going to so decide on their own, thereby prompting the ensuring mechanism to kick in and force them to decide to act in the desired way. By hypothesis, however, the display of these signs (such as a blush or a neurological pattern in the brain) is involuntary. Since it is involuntary, it is implausible to suppose that the addition of the possibility of displaying one of these signs is itself sufficient to transform a case of lack of responsibility into one of responsibility. This, Fischer says, would “appear to involve a kind of alchemy, and it is just as incredible.”²⁰

One thing to notice about this line of response is just how dependent it is upon the presence of some prior sign indicating how agents in these scenarios will “freely” behave if left to act on their own. It is this very feature of traditional Frankfurt-style cases, however, that has been the target of heavy criticism in recent years, prompting the creation of modified Frankfurt-style examples (such as Mele and Robb’s) that do not rely on prior signs of this sort. As noted earlier, one of the effects of eliminating determinative prior signs from Frankfurt-style examples is that, right up until the moment of choice (the locus of free action), it remains within the agents’ direct voluntary control whether they decide on their own to perform the relevant action. Thus, whatever the merits of this particular attempt to undermine the flicker defense of PAP against traditional Frankfurt-style cases, in the context of modified Frankfurt-style scenarios it has no force. Those wishing to preserve this line of objection to the version of the flicker strategy defended here will need to show why omitting to decide on one’s own is not a robust alternative possibility even when it lies within the scope of an agent’s direct voluntary control.

It is also important to see that, since PAP makes no mention of the necessity of robustness (nor of anything that implies the necessity of robustness),

¹⁹ *Ibid.*, 141.

²⁰ *Ibid.*, 141. This argument can also be found in John Martin Fischer, “Responsibility and Self-Expression,” *The Journal of Ethics* 3 (1999): 278–86, and in Fischer, “Recent Work,” 109–23.

merely defending PAP against falsification by Frankfurt-style cases does not require that agents in these cases possess distinctly robust alternative possibilities.²¹ There is, however, another way that these cases can be used to challenge PAP, one for which the demand for robustness would be appropriate. Here I make a distinction between two kinds of Frankfurt-style attack, each of which corresponds to a different way of conceiving the role of Frankfurt-style examples. In the first kind of Frankfurt-style attack, which I shall call *direct Frankfurt-style attack*, the role of Frankfurt scenarios is to act as counterexamples that falsify PAP. In the second kind of attack, which I shall call *indirect Frankfurt-style attack*, the aim of these cases is not to falsify PAP but to falsify some other principle to which it is connected in some important way.²²

This second kind of Frankfurt-style attack is suggested by something Fischer says. He writes, "Of course, I do not accept the alternative-possibilities control model of moral responsibility. But my contention is that, if you do buy into this traditional picture, then you should also accept that the alternative possibilities must be of a certain sort—they must be sufficiently robust."²³ According to Fischer, then, one who endorses PAP should also endorse something like the following general principle (call it the Principle of Robust Alternative Possibilities [PRAP]):

PRAP: A person is morally responsible for what he has done only if he could have done otherwise and his doing otherwise was a robust alternative possibility.

Since PAP does not entail PRAP, proponents of the former are not forced to accept the latter upon pain of contradiction. Why, then, should proponents of PAP endorse PRAP? An earlier passage seems to provide the answer:

The intuitive picture behind the alternative-possibilities control requirement is that moral responsibility requires that the agent select one from among various genuinely open paths the world might take. There are two important ideas here. One is that there must be various paths genuinely available to the agent (at least at some times suitably related to the time of the behavior under consideration). The second idea is that the agent (and not some outside force or mere chance) selects which path will be the path into the future. It seems to me that both ideas are important components of the traditional conception of the sort of control associated with moral responsibility—alternative-possibilities control.²⁴

²¹ For more on this point, see Robinson, "Modified Frankfurt-type Counterexamples."

²² The remainder of this section derives from Michael Robinson, "The Limits of Limited-blockage Frankfurt-style Cases," *Philosophical Studies* 169 (2014): 429–46.

²³ Fischer, "Recent Work," 121.

²⁴ *Ibid.*, 99. See also Fischer, *The Metaphysics of Free Will*, 140–41, and Fischer, "Responsibility and Self-Expression," 283–84.

In Fischer's view, it is the intuition that moral responsibility requires a certain sort of control—namely, the type of control that involves possessing not just any alternative possibilities but ones that are distinctly robust—which lies behind the demand for alternative possibilities. It is PRAP that is the fundamental intuition, not PAP. Of course, PRAP entails PAP, so one who accepts the former should also accept the latter; but, again, PRAP is fundamental.

Mounting an indirect Frankfurt-style attack against PAP involves two steps. The first step is to establish a connection between PAP and some other principle, in this case PRAP. This is where Fischer's point about the basic picture underlying the traditional alternative-possibilities view of moral responsibility comes into play. I suggest that the most plausible way to think about what Fischer says concerning the relationship between PAP and robustness is that PRAP (or some other similar principle) *grounds* PAP in such a way that PAP is plausible only if PRAP is plausible. The second step in indirect Frankfurt-style attack is to falsify PRAP (or at least show that it is implausible). This is where Frankfurt-style cases come in. The role of these cases in indirect Frankfurt-style attack is to act as counterexamples to PRAP.

One advantage of indirect Frankfurt-style attack (over against direct Frankfurt-style attack) is that Frankfurt-style cases need not eliminate all alternative possibilities in order to do the work set out for them. In order for a case of this sort to count as a genuine counterexample to PRAP it need only be one in which an agent has no *robust* alternative possibilities to doing that for which he is basically morally responsible. (Like PAP, PRAP should be taken to be a thesis about basic moral responsibility.) A further advantage of using Frankfurt-style cases in this way is that, here, the demand for robust alternative possibilities is entirely fitting. Any flicker-of-freedom-type defense of PRAP against Frankfurt-style cases, if it is to be successful, will need to show both that agents in these cases continue to possess alternative possibilities and that these alternatives are sufficiently robust.

Of course, since PAP does not entail PRAP, falsifying PRAP is insufficient to demonstrate the falsity of PAP. Thus, indirect Frankfurt-style attack is not able to show that PAP is false (then again, as I have argued, neither is direct Frankfurt-style attack). Rather than falsifying PAP, the aim of indirect Frankfurt-style attack is to *demotivate* this principle by showing that the reasons for endorsing PAP in the first place ought not to be accepted. This in itself would be significant. And given that Frankfurt-style examples are theoretically incapable of falsifying PAP, I suggest that this is the most that opponents of PAP can reasonably hope to achieve with cases of this sort. To accomplish even this much, however, it must be shown both that PAP is plausible only if PRAP is plausible and that there are possible Frankfurt-style scenarios in which agents without any robust alternative possibilities are nevertheless (basically) morally responsible for what they do.

For present purposes, I am willing to grant that PAP is plausible only if PRAP is also plausible, and that falsifying PRAP would have the result of demotivating PAP in the way suggested above. In what follows, I want to consider whether there could be Frankfurt-style scenarios in which agents without any robust alternatives are nevertheless (basically) morally responsible for what they do. First, however, we must get clearer on the nature of robustness.

IV. WHAT DOES ROBUSTNESS REQUIRE?

The robustness challenge, originally offered by Fischer, has been echoed by a number of prominent defenders of Frankfurt to address the flicker strategy. It is generally agreed that a robust alternative is one that is capable of grounding ascriptions of moral responsibility—that is, one that is relevant in and of itself to explaining why an agent is morally responsible for what she has actually done. There is some variation, however, when it comes to spelling out just what is required for an agent to have alternatives that are robust in this sense. Michael McKenna maintains that, in addition to being within an agent's voluntary control, a robust alternative must be such that "It would tell us something (different from what we are told in the actual world) about the moral quality of the agent's conduct were she to have so acted in this alternative scenario."²⁵ To do this, he says, an alternative must be both *morally significant* (that is, morally better or worse than the action the agent actually performs) and *deliberatively significant* (that is, reasonable to expect the agent to have considered as an alternative to what she actually does, given her agent-relative perspective). According to Derk Pereboom, an agent's having a robust alternative requires that the agent "could have willed something other than what she actually willed such that she understood that by willing it she would thereby have been precluded from the moral responsibility she actually has for the action."²⁶

For present purposes, I want to grant that robustness requires everything that Fischer, McKenna, and Pereboom say it does—with one possible exception. At certain points, some of their characterizations of robustness seem open to being interpreted as claiming that a robust alternative must involve an action (even if only a mental act of deciding). This is true of the passage by Pereboom quoted above, and Fischer has at times expressed a similar sentiment:

²⁵ Michael McKenna, "Robustness, Control," 204.

²⁶ See Derk Pereboom, "Source Incompatibilism and Alternative Possibilities," in Widerker and McKenna, *Moral Responsibility and Alternative Possibilities*, 188. David Widerker has offered a similar account of what are here referred to as robust alternatives. See Widerker, "A Defense of Frankfurt-friendly Libertarianism," *Philosophical Explorations* 12 (2009), 87–108.

On the traditional alternative-possibilities picture, it is envisaged that an agent has a choice between two (or more) scenarios *of a certain sort*. In one scenario, he deliberates and forms an intention to perform an act of a certain kind and then carries out this intention in an appropriate way. In at least one other possible scenario, he deliberates and forms an intention to perform a different kind of act (or no act at all) and carries out this intention in an appropriate way. This is what is involved in having robust alternative possibilities, and certainly this is the natural way to think about the sort of alternatives that allegedly ground moral responsibility.²⁷

Like Fischer, McKenna and Pereboom maintain that the mere possibility of having something different occur cannot plausibly contribute to an explanation of an agent's being morally responsible for what she does. To do this, an alternative must be within an agent's control.

The point that alternatives can be robust only if they are within an agent's control is well-taken. It is certainly plausible that whatever intuitive force accompanies PAP stems from the view that moral responsibility requires what Fischer describes as "the ability to make a difference in the sense of selecting one from various paths the world could take, where these various paths are all genuinely available to the agent."²⁸ Unless the availability of various alternatives enhances an agent's control over the way things go, it is unclear how this would help ground attributions of moral responsibility.

Importantly, however, agents can see to it that the world takes one path rather than another, not only by acting but also by omitting to act.²⁹ Consider a normal case of me sitting at the bus stop, waiting for my bus to arrive. Once it arrives there are various paths the world could take that are genuinely available to me. One path involves me getting on the bus and making it home in time for dinner, and I can see to it that the world takes this path by deciding to get on the bus, which it is within my direct voluntary control to do. There are also a number of other paths that do not involve me getting on the bus or making it home in time for dinner, and I can see to it that the world takes one of these paths rather than the first one. One way I can do this is by making a different decision—either a decision not to get on the bus or a decision to do something else that is incompatible with my getting on the bus then. Importantly, however, this

²⁷ Fischer, *The Metaphysics of Free Will*, 142. Though Fischer seems to allow that an agent's performing no act at all could qualify as a robust alternative, this is conditioned on its being the result of a decision to do so. Thus, he grants that robust alternatives can at least partly consist in omissions, but only insofar as they also involve an act of intention formation (that is, a decision).

²⁸ Fischer, "Responsibility and Self-Expression," 283.

²⁹ Here I borrow from my argument for this point in Robinson, "The Limits of Limited-blockage."

is not the only way I can control which path the world takes. Another way I can see to it that the world takes a path that does not involve me getting on the bus is by simply omitting to (decide to) get on the bus. This too is within my direct voluntary control. Just as there is nothing else I need to do in order to decide (not) to get on the bus—I simply decide—there is nothing I need to do in order to omit to (decide to) get on the bus—I simply omit to (decide to) do so.

Given that omissions may equally lie within agents' direct voluntary control, there appears to be no justification for insisting that only actions are capable of constituting robust alternatives. As long as their claims are understood in a way that is compatible with that, let us assume that robustness requires everything that Fischer, McKenna, and Pereboom say it does. Again, this will be helpful in demotivating PAP only if it can be shown that robust alternatives of this sort are unnecessary for moral responsibility.

V. ROBUST FLICKERS OF FREEDOM

We are now in a position to see why indeterministic Frankfurt-style cases are theoretically incapable of falsifying PRAP and, hence, why indirect Frankfurt-style attack (like direct Frankfurt-style attack) cannot possibly succeed.³⁰ To begin, assume that, like PAP,

- (1) PRAP is a thesis about basic moral responsibility.

Also like PAP, PRAP is a conditional in which the referent of the consequent is fixed by that of the antecedent (and not the other way around). That is, PRAP requires that a person had a robust alternative to *doing that for which he is basically morally responsible*. Consequently,

- (2) A counterexample to PRAP must be a case in which a person had no robust alternative to doing that for which he is basically morally responsible.

The first major step in this argument is to show that

- (3) Agents in Frankfurt-style cases are basically morally responsible for deciding on their own to act as they do (not for so deciding *simpliciter*).³¹

³⁰ Deterministic Frankfurt-style cases are unable to serve as good counterexamples to PRAP for the same reason they are unable to do so for PAP (see Sec. II).

³¹ Recall that, following custom, the focus of this discussion has been restricted to cases of moral responsibility for actions (as opposed to omissions) for the sake of simplicity.

Earlier I argued that the same reasons that support the view that moral responsibility for overt actions (such as raising one's arm) derives from responsibility for certain mental actions (such as deciding to raise one's arm) equally support the view that, at least in Frankfurt-style cases, agents' responsibility for deciding to perform the relevant actions derives from their responsibility for so deciding on their own (for which they are basically morally responsible). There is another reason for thinking that agents in Frankfurt-style cases are basically morally responsible for deciding on their own to act as they do, and not for so deciding *simpliciter*, which we can see by considering the following question. Assuming that Bob (from Mele and Robb's case) is morally responsible both for deciding on his own to steal Ann's car and for deciding to steal Ann's car, what should we think about the relationship between his responsibility for the former and his responsibility for the latter? Three possibilities present themselves:

- (i) Bob's moral responsibility for deciding to steal Ann's car derives from his responsibility for deciding on his own to steal Ann's car.
- (ii) Bob's moral responsibility for deciding on his own to steal Ann's car derives from his responsibility for deciding to steal Ann's car.
- (iii) Neither Bob's moral responsibility for deciding on his own to steal Ann's car nor his responsibility for deciding to steal Ann's car derives from the other; each is something for which Bob is basically morally responsible.

Of these options, (ii) appears to be the least plausible. Indeed, it is difficult to imagine what reason there could be for thinking that Bob's moral responsibility for deciding on his own to steal Ann's car derives from his responsibility for so deciding *simpliciter*—particularly since it seems so natural to explain the latter in terms of the former. This is evinced by the fact that so many proponents of Frankfurt-style argument (including Mele and Robb, and even Frankfurt himself) have sought to support the judgment that agents are genuinely morally responsible for their decisions in these cases precisely by pointing out that these are decisions the agents made on their own. For this reason, (iii) too seems problematic insofar as it suggests that there is *no* connection between Bob's responsibility for deciding to steal Ann's car and his responsibility for so deciding on his own. Simply on the face of it, then, (i) appears to provide the most (if not the only) plausible way of thinking about the relationship between Bob's responsibility for deciding on his own to steal Ann's car and his responsibility for so deciding *simpliciter*.

From (2) and (3) it follows that

- (4) A counterexample to PRAP must be a case in which a person had no robust alternative to deciding on his own to act as he did.

As noted earlier, here I want to grant that having robust alternatives involves everything Fischer, McKenna, and Pereboom say it does (again, so long as their claims are not interpreted as requiring robust alternatives to be actions). Thus, let us assume that

- (5) An alternative is robust if and only if it was (i) morally significant, (ii) deliberately significant, (iii) an alternative that the agent would understand to be a way of precluding himself from being (basically) morally responsible for what he did, and (iv) within the agent's voluntary control.

For the sake of simplicity, most (if not all) of the debate surrounding PAP, Frankfurt-style cases, and robustness has limited the focus to cases involving blameworthiness. Following that custom, let us assume that the relevant actions in these cases are morally bad actions for which the agents are meant to be not only morally responsible but blameworthy.

The next major step is to show that the alternative of *voluntarily omitting to decide on one's own to do that which one knows to be wrong* satisfies all the conditions for robustness identified in (5). To begin, I take it to be noncontroversial that

- (6) Omitting to decide on one's own to do that which one knows to be wrong is a morally significant alternative to so deciding on one's own (in all Frankfurt-style cases).

Arguably, omitting to decide on one's own to do that which one knows to be wrong (all things considered) is always a morally significant alternative to (because it is morally better than) deciding on one's own to do that which one knows to be wrong. Possible exceptions to this might be cases in which a person knows that her deciding to A is unavoidable, regardless of whether she decides on her own to A (say, because she knows that she will be forced to decide to A unless she so decides on her own). Of course, in such cases it is no longer obvious that performing the action would be morally wrong, as opposed to just morally bad. Regardless, since it is an essential feature of Frankfurt-style cases that the agents neither believe nor ought to believe that there are counterfactual interveners about, ready to force them to decide to A unless they do so on their own, these considerations do not apply in Frankfurt-style cases.

I also take it to be evident that

- (7) Omitting to decide on one's own to do that which one knows to be wrong is a deliberately significant alternative to so deciding on one's own (in all indeterministic Frankfurt-style cases).

Assuming that agents in these cases are morally responsible for what they do (as we are to suppose), omitting to decide on their own to do that which they know to be wrong is a morally significant alternative that it is reasonable for us to expect them to consider as an alternative to deciding on their own to do that which they know to be wrong.³²

Similarly, any agent satisfying the epistemic conditions for morally responsible agency will understand that she can avoid being (basically) morally responsible for doing something by simply omitting to do so on her own. For any morally competent agent will understand that one is not (basically) morally responsible either for actions one is forced to perform nor for actions one does not perform at all, and these are the only possibilities when one omits to perform an action on one's own.³³ Thus,

- (8) Omitting to decide on one's own to perform a given action is an alternative that one would understand to be a way of precluding one from being (basically) morally responsible for so deciding (in all Frankfurt-style cases).

Next, it is trivially true that

- (9) Voluntarily omitting to decide on one's own to do that which one knows to be wrong is within one's voluntary control whenever one voluntarily omits to do so.

From (5) through (9), it follows that

- (10) Voluntarily omitting to decide on one's own to do that which one knows to be wrong is a robust alternative to so deciding on one's own (in all indeterministic Frankfurt-style cases).

That is to say, voluntarily omitting to decide on one's own to do that which one knows to be wrong is an alternative that, *when it remains open to agents in indeterministic Frankfurt-style cases*, is a distinctly robust alternative possibility. Premise (10) should *not* be understood as saying that this alternative *in fact* remains open to agents in all indeterministic Frankfurt-style cases. (That claim comes next.)

³² For an extended discussion of this point, see Robinson, "The Limits of Limited-blockage." Cf. McKenna, "Robustness, Control."

³³ Certainly, there may be cases in which agents are morally responsible for the performance of actions that they are forced to perform, provided that they are morally responsible for the fact that they are forced to perform them. For instance, if Jones freely arranged for Black to force him to perform a given action, Jones could be morally responsible for the performance of this action even though this is not an action he performs on his own. In all such cases, however, agents will be only derivatively morally responsible for the performance of those actions they do not perform on their own. Since PAP and PRAP are theses about *basic* moral responsibility, these cases have no bearing on the present point.

The final major step in the argument is to show that

- (11) It is an essential feature of indeterministic Frankfurt-style cases that it is possible for agents in these cases to voluntarily omit to decide on their own to do that which they know to be wrong.

It is precisely because it is possible for these agents not to perform the relevant actions on their own that the ensuring mechanisms are there in the first place; otherwise, they would be superfluous, which they are not. Without a mechanism that stands ready to force the agents to make the relevant decisions unless they do so on their own, there is nothing to guarantee that the agents make the relevant decisions at all. Crucially, however, even if these ensuring conditions are able to guarantee that the agents make the relevant decisions, they are not able to guarantee that the agents make the relevant decisions *on their own*. There is no way to guarantee that an agent will make a certain decision on his own because it is simply impossible (indeed, incoherent) to *force* an agent to do something *on his own* (that is, without being forced). As a result, it is always possible for an agent not to do a thing on his own. In indeterministic cases, this option is available right up until the moment of choice, within the scope of agents' voluntary control.³⁴ The only way to prevent an agent from *voluntarily* omitting to decide on her own to perform an action she is considering performing, without causally determining what she does, is to render her unconscious or otherwise not in control of her behavior, thereby precluding her responsibility. Thus, provided they are morally responsible for what they do, agents in these cases will be able to voluntarily omit to decide on their own to do so.

To be clear, I allow that it is possible to construct indeterministic Frankfurt-style cases in which agents are morally responsible for deciding to perform some action despite being unable to *intentionally* omit to decide on their own to perform that action.³⁵ But an omission need not be intentional in order to constitute a robust alternative possibility. It must be within an agent's control, to be sure, but an omission can be within an agent's control (even direct voluntary control) without being intentional.³⁶

³⁴ Returning this alternative to within the scope of agents' voluntary control was one of the results of the general move away from Frankfurt-style cases that relied on determinative prior signs. See McKenna and Widerker, 6–10.

³⁵ It is generally accepted that in order for an action to be intentional it must be suitably related to an intention of a certain sort (e.g., an intention to perform that action, or an intention to try to perform that action). Plausibly, the same is true of intentional omission. See Randolph Clarke, "Intentional Omissions," *Nous* 44 (2010): 158–77. Insofar as this is correct, (indeterministic) Frankfurt-style scenarios can eliminate agents' ability to intentionally omit to decide on their own to act as they do simply by blocking them from coming to have the relevant intention(s).

³⁶ I have also argued for this point in Robinson, "The Limits of Limited-blockage."

The kind of omission I here have in mind corresponds to a class of action that some philosophers have dubbed “nonintentional action.”³⁷ To illustrate, consider a dentist who performs a necessary procedure on a patient and, in so doing, causes the patient to be in pain. Suppose that, as she is performing the procedure, the dentist is consciously aware of the fact that she is causing the patient pain, even though she has no intention of doing so and is doing everything she can to minimize the patient’s pain and discomfort.³⁸ Because the dentist is in control of her actions and is aware that she is causing the patient pain as she is doing it, it would seem to be a mistake to say that she does this unintentionally. At the same time, however, it would also be inaccurate to say that she intentionally causes the patient to be in pain since causing the patient pain in no way figures into the content of any intention she has with respect to performing the procedure. Fortunately, as some theorists point out, there is a middle ground: “insofar as an agent who is A-ing is neither aiming at A-ing nor trying to A, either as an end or as a means to (or constituent of) an end, she is not intentionally A-ing; insofar as an agent is A-ing knowingly and nonaccidentally, she is not unintentionally A-ing; and actions that are neither intentional nor unintentional are nonintentional.”³⁹ I suggest that, like actions, omissions too can be nonintentional in this way. Just as one can knowingly and voluntarily perform an action and yet not perform it intentionally, one can knowingly and voluntarily omit to perform an action without intentionally omitting to perform it.

To see this, think back to the ordinary case of my sitting at the bus stop from earlier. Once my bus pulls up to the stop, there are various paths the world can take. Some paths involve my getting on the bus and making it home in time for dinner, while many others do not. Of those that do not, some paths involve my getting up to do something else then (climb a tree, go to a bar, stroll through the park), while other paths do not involve my doing anything other than remaining seated until the bus pulls away and simply omitting to get on the bus. Consider this last path. If my remaining seated until the bus pulls away were caused by (or otherwise appropriately related to) a decision or an intention not to get on the bus, then we could reasonably say that I intentionally omitted to get on the bus. If, instead, my remaining seated until the bus pulls away were the result of my getting lost deep in thought and failing to recognize that the bus had arrived, then this would seem to be a clear case of my unintentionally omitting to get on the bus (and not something I did voluntarily).

³⁷ See Alfred Mele and Paul Moser, “Intentional Action,” *Nous* 28 (1994): 39–68, and Alfred Mele and Steven Sverdlik, “Intention, Intentional Action, and Moral Responsibility,” *Philosophical Studies* 82 (1996): 265–87.

³⁸ This example comes from Mele and Sverdlik, “Intention, Intentional Action, and Moral Responsibility,” 274.

³⁹ *Ibid.* See also Mele and Moser, “Intentional Action,” 45.

But what would nonintentionally omitting to get on the bus amount to? What would it look like to knowingly and voluntarily, yet not intentionally, omit to get on the bus? Suppose that, when the bus arrives, I am still uncertain whether to get on the bus. I recognize that the bus has arrived at the stop and will soon depart without me unless I (decide to) get on the bus. As the seconds tick by, I continue to remain seated at the bus stop. I have not decided (and do not intend) to get on the bus, nor have I decided (or otherwise come to intend) not to get on the bus. Perhaps I continue to deliberate about whether to get on the bus, perhaps I do not. Either way, as each moment passes, I am aware, of course, that I am not currently getting on the bus and have not as of yet decided to do so. Eventually, the bus pulls away without me.

This is a prime example of nonintentional omission, which is when one knowingly and voluntarily omits to do something but without intentionally omitting to do it. Because I did not intend not to get on the bus in this example (nor to do anything else incompatible with my getting on the bus then), it is implausible to maintain that my omitting to do so was intentional. And because I was fully aware that I was omitting to (decide to) get on the bus every moment I remained seated, it would seem to be a mistake to say that my omitting to get on the bus was unintentional or accidental. What remains is nonintentional omission. Although I was aware that I was not getting on the bus, I had not settled on not getting on the bus—I did not intend not to do so. And even though I did not decide or intend not to get on the bus, I was nevertheless fully in control of my behavior. I was not subject to threat or coercion of any kind, and my getting on or not getting on the bus remained entirely up to me. Thus, there appear to be no grounds for denying that my omitting to get on the bus in this case was a voluntary omission on my part.⁴⁰

⁴⁰ While I have argued that omissions can be nonintentional in a way that is analogous to actions, there is a potential asymmetry that is worth highlighting here. Typical examples of nonintentional action involve side-effects of other actions performed intentionally: a runner who wears down the soles of her shoes in the course of running a marathon, a sniper who alerts the enemy to his presence in the course of shooting at his target, a dentist who inflicts pain upon her patient in the course of performing a needed medical procedure. As such, they are only indirectly within an agent's voluntary control. Although these nonintentional actions are indeed performed voluntarily, their voluntariness derives from their relation to other, intentional actions, which the agent performs voluntarily. By contrast, there is nothing *else* an agent needs to do *in order* to knowingly and voluntarily (that is, nonintentionally) omit to perform some action; he simply omits to do so. To be clear, this is not to say that omitting to A always involves knowingly and voluntarily omitting to A. Certainly there are other conditions that must be met in order for an agent who omits to A to knowingly and voluntarily omit to A—namely, that the agent is aware that he is omitting to A and that his omitting to A is within his voluntary control (and perhaps free from coercion). The main point here is that there is nothing *else* he needs to *do* in order for his omitting to A to qualify as an instance of knowingly and voluntarily omitting to A. Thus, whereas nonintentional actions might only ever be indirectly within agents' voluntary control, nonintentional omissions are (or at least can be) within an agent's direct voluntary control.

From (10) and (11), it follows that

- (12) It is an essential feature of indeterministic Frankfurt-style cases that agents in these cases have robust alternatives to deciding on their own to act as they do.

And from (4) and (12), it follows that

- (13) Indeterministic Frankfurt-style cases are theoretically incapable of serving as genuine counterexamples to PRAP.

So long as it remains open to agents in these cases to knowingly and voluntarily omit to decide on their own to do that which they know to be wrong, they will continue to possess robust alternative possibilities. And since there is no way to eliminate the possibility of agents knowingly and voluntarily omitting to decide on their own to act as they do without undermining the judgment that they are morally responsible for their actions, agents in these cases will continue to possess robust alternative possibilities so long as they are morally responsible for what they do. Thus, either agents are not morally responsible for what they do in these cases or they continue to possess robust alternative possibilities. Either way, these cases will fail to falsify PRAP and so will not be useful in mounting a successful indirect Frankfurt-style attack on PAP.

VI. CONCLUSION

Toward the end of the 1968 film, *The Lion in Winter*, there is a scene where Richard, Geoffrey, and John have been locked in the dungeon by their father, King Henry II, and are waiting to be executed. Eventually, upon hearing footsteps coming down the stairs and believing they had finally seen sunlight for the last time, Richard stands to his feet, defiantly declaring, "He'll get no satisfaction out of me. He isn't going to see me beg." "Why, you chivalric fool," Geoffrey chides, "as if the way one fell down mattered." "When the fall is all there is," Richard replies, "*it matters.*"

Critics of the flicker strategy have argued that the alternatives that remain open to agents in Frankfurt-style examples, including the option of omitting to perform the relevant actions on their own, are unable to do the work set out for them, that they are not the kinds of alternative on which proponents of the ability to do otherwise have traditionally been focused, and that they are too diminutive to matter. Granted, generally believing that it remains up to us not only *how* we perform the actions we freely perform but *whether* we perform them at all, this is not the full-fledged freedom with which we are normally concerned and which we commonly take ourselves to possess. Still, when this flicker is all there is—when the

ability to knowingly and voluntarily omit to do on one's own that which one knows to be wrong is all the freedom that remains—it matters.

Indeed, that such flickers matter—and, moreover, that they matter in such a way that they are able to help ground ascriptions of moral responsibility—is, I suggest, the main lesson to be drawn from Frankfurt-style cases. Although many philosophers have been convinced by these cases that we ought to reject entirely the ability to do otherwise as a necessary condition for moral responsibility, this is not something that Frankfurt-style cases are (even in principle) capable of showing. Instead, what they show is that there are *certain ways of understanding this condition* that ought to be rejected. In particular, proponents of an alternative-possibilities condition for moral responsibility have commonly held that the relevant alternative must be an alternative intentional action (or at least an intentional omission), so that an agent will be morally responsible for deciding to A only if it was possible for her to have intentionally done something other than decide to A (for instance, decide not to A or decide to do something else that is incompatible with A-ing). But just as Locke's locked room demonstrates that one can be morally responsible for performing an overt (intentional) action even if one could not have avoided performing that overt action, Frankfurt-style cases show us that one can be morally responsible for making a decision even if it was not possible for one to have avoided making that decision. These cases are instructive, not because they show that PAP is false or ought to be rejected, but because they serve to identify the alternative possibilities that can plausibly be thought to be intrinsically relevant to moral responsibility. Their value is in showing us how best to understand this principle, not in demonstrating why we should abandon it.

Philosophy, Chapman University