

The Robustness Requirement on Alternative Possibilities

Taylor W. Cyr

Forthcoming in *Journal of Ethics*; please cite published version.

Abstract In a series of recent papers, Justin Capes and Philip Swenson (together) and Michael Robinson (independently) have proposed new versions of the flickers of freedom reply to Frankfurt-style cases (FSCs). Both proposals claim, first, that what agents in FSCs are morally responsible for is performing a certain action *on their own*, and, second, that agents in FSCs retain *robust* alternative possibilities—alternatives in which the agent freely omits to perform the pertinent action on their own. In this paper, I argue that, by attending to the details of the omissions in question, it becomes clear that agents in FSCs lack robust alternatives of this sort, for in the alternative sequences such agents do not freely omit. Since the problem for these recent proposals arises from their attempt to show that agents in FSCs retain robust alternatives, I go on to consider whether the flicker theorist might be better off either revising or rejecting the robustness requirement on alternative possibilities. I argue that neither alternative is available to the flicker theorist, and yet I also point out that the reasons why these alternatives are unavailable serve to highlight what exactly is at issue in the debate between leeway and sourcehood theorists, namely the grounds of our freedom.

1. Introduction

Is it possible for a person to be morally responsible for an action even if that person could not have done otherwise?¹ Many philosophers have said ‘no’ and have endorsed a version of the following principle:

Principle of Alternative Possibilities (PAP): A person is morally responsible for performing an action only if she could have done otherwise than perform that action.

In his famous challenge to PAP, however, Harry Frankfurt (1969) presents a case in which an agent lacks the ability to do otherwise than some action and yet is intuitively morally responsible for that action. Here is an instance of a “Frankfurt-style case” (FSC):

Assassin: A nefarious neurosurgeon named Black wants Jones, a trained assassin, to decide to kill Smith. Black is willing to force Jones’s hand if need be, but he would prefer that

¹ Here and throughout the paper I use ‘morally responsible’ and cognates to refer only to *direct* moral responsibility, not to *derivative* moral responsibility.

Jones make the decision to kill Smith on his own. So, he secretly implants a device in Jones's brain that enables him to control Jones's thoughts and behavior. The device is rigged to deterministically cause Jones to decide at time t to kill Smith, if, but only if, Jones does not decide on his own at t to kill Smith. There is, moreover, nothing Jones can do to prevent Black from carrying out this scheme. To Black's delight, Jones decides on his own at t to kill Smith, and so the coercive device never comes into play. (Capes and Swenson 2017: 967-968)²

Successful FSCs are cases in which both 1) an agent is morally responsible for performing a certain action and 2) the agent could not have done otherwise than perform that action. *Assassin* appears to be a successful FSC. But if any FSC is successful, then PAP is false.

A popular type of response to FSCs by leeway theorists (those who defend PAP) attempts to show that, in any scenario in which the agent is intuitively morally responsible for performing a certain action, things could have gone differently for the agent, and thus the second condition on successful FSCs is not satisfied.³ In other words, according to this type of response to FSCs, agents in such cases retain what John Martin Fischer (1994: 134) has called "flickers of freedom"—alternative possibilities that are not eliminated by the presence of the counterfactual intervener. For example, in *Assassin*, while Jones makes the decision to kill Smith *on his own* in the actual sequence, it remains an open alternative possibility that he make that decision as a result of Black's coercive device (rather than on his own). In response to the "flicker of freedom" strategy (or "flicker strategy" for short), Fischer and others have objected that the alternative possibilities in question are insufficiently *robust* to ground the agent's moral responsibility (or to ground our

² See Mele and Robb (1998) for a predecessor to this case.

³ This response originated with van Inwagen (1978; 1983) and was subsequently developed by Naylor (1984). Other more recent proponents include O'Connor (2000) and Speak (2002).

judgments about the agent's moral responsibility) in the actual sequence. How exactly to characterize the robustness requirement on alternative possibilities remains contentious, but nearly everyone agrees that not just any alternative possibility will do.

In this paper, I revisit the robustness requirement on alternative possibilities. I begin, in section 2, with a summary of Fischer's original version of the requirement. I then consider Justin Capes and Philip Swenson's (2017) version of the flicker strategy in section 3 and Michael Robinson's (2012; 2014; 2019) version in section 4. While importantly different approaches, both proposals rely on the claim that agents in FSCs retain alternatives that satisfy Fischer's robustness requirement, and, in section 5, I argue that neither proposal succeeds since agents in FSCs lack robust alternatives of that sort. I go on, in sections 6 and 7, to consider whether the flicker theorist might be better off either revising or rejecting the robustness requirement on alternative possibilities. I argue that neither alternative is available to the flicker theorist, and yet I also point out that the reasons why these alternatives are unavailable serve to highlight what exactly is at issue in the debate between leeway and sourcehood theorists, namely the grounds of our freedom.

2. Fischer's Robustness Requirement

Recall that the flicker strategy aims to account for the intuitive judgment that agents in FSCs are morally responsible by pointing to residual alternative possibilities in such cases. According to Fischer, however:

...even if the possible event at the terminus of the alternative sequence (in the case of Jones and Black) is indeed an alternative possibility, it is highly implausible to suppose that it is *in virtue* of the existence of such an alternative possibility that Jones is morally responsible for what he does. 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. (1994: 140)

In order for the flicker strategy to be successful, Fischer thinks, it must be the case that the residual alternative possibilities in FSCs—the flickers of freedom—plausibly *ground* the agent’s moral responsibility for what she actually does. Unless the “flicker theorist” can show that the residual alternatives contribute to the grounding of the agent’s responsibility, then, the mere fact that *there are* flickers of freedom in FSCs does not undermine the challenge to PAP.

But why think that alternative possibilities must ground the agent’s moral responsibility, and what would an alternative possibility have to be like in order plausibly to ground an agent’s moral responsibility? Fischer explains:

The existence of *various* genuinely open pathways is alleged [by the leeway theorist] to be *crucial* to the idea that one has *control* of the relevant kind. But if this is so, I suggest that it would be very puzzling and unnatural to suppose that it is the existence of various alternative pathways along which the agent does *not* act freely render it true that the agent has the relevant kind of control (regulative control). And notice that this is precisely the situation in the Frankfurt-type cases. (1994: 141)

This passage suggests that it is because the leeway theorist (including the flicker theorist) takes alternative possibilities to be crucial to the control required for moral responsibility that Fischer

requires them to play a grounding role in an agent's responsibility in order for them to be robust.⁴ This much, as I will argue below (in section 6), should be granted by all parties to the debate. Fischer goes on, however, to claim that, in order for the alternative possibility plausibly to ground an agent's actual moral responsibility, the agent must act freely in the alternative scenario (and, of course, Fischer denies that agents act freely in the alternative sequences of FSCs). Fischer supports this claim by pointing out that it would be very puzzling to see how an alternative possibility could plausibly ground the agent's freedom in the actual sequence unless the agent acts freely in that alternative scenario.

Putting these ideas together, then, we can distill Fischer's claims about robustness into the following:

Fischer's Robustness Requirement (FRR): An agent's alternative possibility is robust only if the agent's moral responsibility for what she does in the actual sequence is plausibly grounded in the alternative possibility; and an agent's moral responsibility for what she does in the actual sequence is plausibly grounded in an alternative possibility (some alternative sequence) only if the agent acts freely in the alternative sequence.

Since agents in FSCs do not act freely in their available alternative sequences (but rather act as a result of some coercive device), the second component of FRR says that their moral responsibility in the actual sequence is not plausibly grounded by their alternative possibilities. And since their moral responsibility in the actual sequence is not plausibly grounded by their alternative possibilities, the first component of FRR says that agents in FSCs lack robust alternatives. Importantly, then, Fischer's objection to the flicker strategy requires a commitment to both

⁴ In more recent work, Fischer (2018) builds on this discussion of the role of alternatives and the grounds of responsibility in distinguishing his own actual-sequence compatibilism from a form of leeway compatibilism that has come to be known as "the new dispositionalism."

components of FRR. In the next two sections, we will consider two attempts to respond to Fischer's objection, both of which grant both components of FRR. After arguing that these responses to Fischer's objection are inadequate, I will consider whether the flicker theorist is better off rejecting (at least) the second component of FRR.

3. *Capes and Swenson's Proposal*

Recently, Capes and Swenson (2017) have responded to Fischer's robustness worry for the flicker strategy by arguing that agents in FSCs do in fact act freely (more accurately, freely refrain from acting) in the alternative sequences. In other words, they grant Fischer the truth of FRR but claim that agents in FSCs possess robust alternatives (and thus, they think, FSCs fail to undermine PAP). According to Capes and Swenson's "fine-grained" version of the flicker of freedom strategy, "a correct assessment of cases like *Assassin* requires being very precise about what agents in those examples are blameworthy for" (2017: 969). In *Assassin*, for example, what Jones is morally responsible for, on their view, is his deciding *on his own* to kill Smith, not his decision to kill Smith *simpliciter*. But notice that Jones had an alternative to deciding on his own to kill Smith; in the alternative sequence, Jones decides to kill Smith *as a result of Black's coercive device*. "So, if deciding on his own to kill Smith is what Jones is really blameworthy for," Capes and Swenson reason, "then, contrary to what Frankfurt and others claim, cases like *Assassin* do not provide us with scenarios in which someone is directly blameworthy what he did at *t* even though, through no fault of his own, the person could not have avoided doing it" (2017: 968).⁵

Now, of course, Fischer will agree with Capes and Swenson that agents in FSCs have certain alternatives—flickers of freedom. What Fischer denies, however, is that these alternatives

⁵ See also Swenson (2019).

are *relevant* alternatives, since they do not plausibly ground the agent's moral responsibility (see the first component of FRR). And to plausibly ground the agent's moral responsibility, Fischer thinks, an alternative must be one in which the agent acts freely (see the second component of FRR). Concerning the specific version of the flicker of freedom strategy endorsed by Capes and Swenson, Fischer says:

If [his deciding to kill Smith on his own] is the appropriate specification of the content of Jones' moral responsibility, then evidently there is an alternative possibility. But note again that this alternative possibility lacks robustness. After all, in the alternative sequence Jones does *not* freely refrain from "[deciding to kill Smith] on his own." Indeed, he does *not* freely behave in any fashion... Thus, again, it seems to me that the alternative possibilities so nicely generated by the strategy of redescription of the content of moral responsibility lack robustness. (1994: 143)

In order to show that Jones's alternative (deciding to kill Smith as a result of Black's coercive decide) is a relevant one, then, Capes and Swenson would need to show that Jones acts freely in this alternative sequence (i.e., that Jones freely refrains from deciding to kill Smith on his own). But, Fischer argues, clearly Jones does not freely refrain in the alternative sequence, so this alternative possibility cannot be relevant to Jones's moral responsibility in the actual sequence.

Capes and Swenson argue, however, that Jones *does* freely refrain from deciding on his own to kill Smith in the alternative sequence. To see why they think it could be the case that Jones freely refrains, despite his overt action's being a result of Black's coercive device, Capes and Swenson first point out that not all cases of free omissions are cases in which the omission results from a prior choice. Consider the following case:

Indecision: Marla is deliberating about whether to attend a party this evening. Part of her wants to go; it will be a fun party, and she knows that she will have a good time. Another part of her, though, would prefer a quiet evening at home. At t , where t is some instant during the period of time that Marla is deliberating, Marla omits to decide to attend the party. To be clear, she does not decide at t not to attend the party, nor does she decide not to make a decision at t . She simply fails to decide at t one way or another. (2017: 971)

If we fill in the details of the case, Capes and Swenson explain, *Indecision* illustrates the possibility of freely omitting to make a certain decision. They continue:

Typically when an agent omits or refrains from making a specific decision at t , the agent does not first decide not to make that decision at t . Instead, the agent either makes a different decision at t , or, as in cases like *Indecision*, makes no decision at all. And, as in *Indecision*, as long as the agent had it within his power at the time to A and his failure to A was not the result of any freedom-subverting factor, it appears that the agent freely avoids A -ing. (2017: 971)

As long as an agent's refraining from making a certain decision does not result from some freedom-subverting factor, Capes and Swenson think, the agent may freely refrain from making the decision even if the agent's omission to make the decision did not result from some prior choice.

But, Capes and Swenson argue, the forgoing points suggest a response to Fischer's objection to the flicker strategy. In *Assassin*, Capes and Swenson argue:

Jones had it within his power to decide on his own to kill Smith, and his failure in the alternative sequence to decide on his own was not a result of coercion, manipulation, or any other freedom-subverting factor. [Footnote omitted] But once we acknowledge all this, it seems we should also acknowledge that Jones freely avoids deciding on his own in the

alternative sequence, even though he does not choose the possibility of not deciding on his own to kill Smith. (2017: 971-972)

In the alternative sequence in *Assassin*, Jones is coerced into deciding to kill Smith, and thus, everyone should agree, Jones does not freely make this decision. Nevertheless, according to Capes and Swenson, we should acknowledge that (in the alternative sequence) Jones's *omission to decide on his own to kill Smith* is a free omission. Thus, Capes and Swenson think that, in the actual sequence, Jones "could have freely avoided deciding on his own to kill Smith, which, in turn, strongly suggests that that alternative possibility is sufficiently robust to ground the agent's responsibility for deciding on his own to kill Smith" (2017: 972).

4. Robinson's Proposal

Like Capes and Swenson, Robinson (2012; 2014; 2019) argues that responding to FSCs requires paying close attention to what exactly agents in these cases are morally responsible for. Unlike Capes and Swenson, Robinson is willing to grant that, in a case like *Assassin*, Jones is morally responsible for the unavoidable decision (*simpliciter*), but, contrary to proponents of FSCs, Robinson argues that Jones is only *derivatively* morally responsible for this decision. What Jones is *basically* (or *directly*) morally responsible for, according to Robinson, is his deciding *on his own* to kill Smith. Using "A" as an action variable, Robinson writes:

Perhaps the lesson from [certain variations of FSCs] is not that agents in Frankfurt-type counterexamples (as well as other similar scenarios) are not at all morally responsible for A-ing, even though they are responsible for A-ing on their own, but merely that their responsibility for A-ing *derives* from their responsibility for A-ing on their own. On the face of it, this certainly seems to have the ring of plausibility...If a person is not morally

responsible for doing a thing on her own, it is difficult to see how she could be morally responsible for doing it simpliciter. (2012: 186-187; cf. 2019: 216-217)

Applying this view to *Assassin*, Jones is basically morally responsible for deciding on his own to kill Smith, and he is only derivatively morally responsible for deciding simpliciter. On Capes and Swenson's proposal, by contrast, Jones is not morally responsible for deciding simpliciter, though he *would* be if the counterfactual intervener were to be subtracted from the case.

While this difference between the two proposals is important, we can safely ignore it here.⁶ Relevant for our purposes is that these two proposals have in common, which is their agreement about the importance of the alternative possibility in which Jones omits to decide on his own to kill Smith. As Robinson points out, "omitting to perform the relevant action on one's own is a genuine alternative possibility in all Frankfurt-style scenarios" (2019: 219). Moreover, like Capes and Swenson, Robinson takes this alternative to be robust, since "agents can see to it that the world takes one path rather than another, not only by acting but also by omitting to act" (2019: 224). Robinson provides his own example, similar to Capes and Swenson's *Indecision*, which I will label *Bus*:

Bus: Consider, for instance, 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

⁶ For discussion of the differences and relative merits of these proposals, and for other reasons for rejecting them, see Sartorio (2019).

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 is not the only way that 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. (2014: 439-440; cf. 2019: 224-225)

And, just as Capes and Swenson say that reflection on *Indecision* suggests a response to Fischer’s robustness objection, Robinson (2014: 443) goes on to make a similar response, arguing that agents in FSCs have it within their control not to decide on their own to do the bad action, despite the presence of the counterfactual intervener.⁷ This, in turn, suggests that agents in FSCs have alternative possibilities that are sufficiently robust to ground their moral responsibility.

5. Problem: Jones Doesn’t Freely Omit to Decide on His Own

Is it plausible, though, that agents in FSCs have alternatives that satisfy FRR? I will argue that it is not, and, for simplicity’s sake, I will focus on the *Assassin* case. For Jones’s alternative—refraining from deciding on his own to kill Smith—to be sufficiently robust, it must be the case

⁷ Robinson (2014) is responding primarily to McKenna (2003), who uses an example, Brain Malfunction, that I have not discussed here but that is relevantly similar to *Assassin*.

that Jones *freely* refrains from deciding on his own to kill Smith. But Jones doesn't freely omit to decide on his own to kill Smith.⁸

The first thing to note is that Jones is clearly not morally responsible for this omission, for he does not satisfy the epistemic condition on moral responsibility with respect to that omission. Prior to refraining from deciding on his own to kill Smith, Jones had no idea that, if he did not decide on his own to kill Smith then, Black's device would force him to decide to kill Smith then. Given this ignorance, it is implausible that Jones was aware (or should have been aware) of the salience of omitting to decide *on his own* to kill Smith then. Unlike Marla in *Indecision* or Robinson in *Bus*, both of whom could reasonably be expected to be aware of the salient options—deciding to go to the party or not deciding to go to the party, deciding to get on the bus or not deciding to get on the bus—, Jones could not reasonably be expected to know that his salient options were to decide on his own at *t* to kill Smith or to omit to decide on his own at *t* to kill Smith. So if Jones's freely omitting to decide on his own to kill Smith in the alternative sequence entails that Jones is morally responsible for this omission, then Jones does not freely omit to decide on his own to kill Smith.

Perhaps this is too quick, however, for there may be cases in which agents are ignorant of morally relevant consequences of their omissions and who nevertheless seem morally responsible for omitting. Suppose that Jane falsely believes that she can kill Dan just by thinking *let Dan be dead* and yet omits to do so. Although the stakes involved in omitting are vastly different than Jane thinks, and although she is ignorant of which options are salient (from our perspective), since she

⁸ A different worry for the flicker strategy, developed in Stockdale (2022), is that it is apparently possible to construct variants of FSCs in which agents could not have done otherwise than decide on their own (even in the alternative sequence). As Stockdale notes, this objection to the flicker strategy sidesteps arguments about robustness, and for that reason I will set aside this objection here, where my focus is on robustness.

thinks one of her salient options involves killing Dan, Jane still seems morally responsible for omitting. If Jane's ignorance does not excuse, why think that Jones's ignorance that Black's device would force him to decide to kill Smith excuses him for the omission to decide on his own?⁹

In my view, however, there is an important and relevant difference between Jones's ignorance and Jane's ignorance, and attending to this difference will help to clarify what I take to be required in order to meet the epistemic condition on moral responsibility for omissions. Because of his ignorance of Black's device, the options Jones considers, and the only reasonably accessible options for Jones to consider, we may suppose, are to decide to kill Smith, to decide not to kill Smith, and to put off making the decision until later. Given the presence of Black's device, however, only two options are really possible: Jones decides to kill Smith on his own and Jones is forced by Black's device to decide to kill Smith. Key here is that the second of these true options is not reasonably accessible to Jones—it's not on his radar, nor could he reasonably have been expected to have it on his radar. By contrast, the options Jane considers are to think *let Dan be dead* (which she thinks will result in Dan's death) or to omit to think this. Despite her mistaken belief that thinking *let Dan be dead* will bring about Dan's death, the options she considers (to think this thought or to omit to think it) are indeed her options, and so, unlike Jones, her options (including the one she takes) are reasonably accessible to her. This explains how Jane can be morally responsible for her omission despite her ignorance of a morally relevant consequence of that omission.

So far I have been arguing that Jones does not meet the epistemic condition on moral responsibility in the alternative sequence, but one might respond by pointing out that this is not relevant. In particular, one might claim that it is possible for an agent to act freely (thus satisfying

⁹ Thanks to an anonymous reviewer for raising this worry and for providing the case.

one of the necessary—but not sufficient—conditions for moral responsibility, namely the freedom condition) without satisfying the epistemic condition on moral responsibility. On this view, while Jones freely omits in the alternative sequence, he is not morally responsible for his omission since he does not satisfy the epistemic condition. And since, as we saw above, Fischer’s robustness requirement (FRR) requires only that an alternative be one in which the agent *acts freely* (in order for the alternative to be robust), *not* that she is *morally responsible* in that alternative sequence, perhaps it is plausible to maintain that Jones has a robust alternative after all.

In my view, however, there are two reasons to think that this response to my criticism cannot ultimately succeed, and thus that Jones lacks a robust alternative. First, as Alfred Mele (2010) argues, it is unclear whether one can fail to satisfy some epistemic condition on moral responsibility without also failing to satisfy the freedom condition on moral responsibility. Consider a case from Fischer and Mark Ravizza (1998) that is meant to illustrate how failing to meet the epistemic condition can excuse one of moral responsibility: a person “backs his car out of his garage unaware that a tiny kitten is snoozing beneath the rear tire” (1998: 12). According to Fischer and Ravizza, the reason that “it would be odd to judge that the driver is morally responsible for the kitten’s untimely death...is that he was (through no fault of his own) unaware that he was driving over the little kitten” (1998: 12). Fischer and Ravizza are clearly right that the driver is not morally responsible for driving over the kitten, but does he nevertheless freely drive over the kitten? Not only is an affirmative answer to that question highly counterintuitive, but, as Mele argues, Fischer and Ravizza’s own theoretical commitments imply that this agent does not freely kill the kitten.¹⁰ Similarly, I take it that it would be counterintuitive to maintain that Jones freely

¹⁰ It remains to be seen, of course, whether there is any epistemic condition on moral responsibility that is not also a condition on acting freely, but since the cases typically invoked to motivate the

omits to decide *on his own* at *t* to kill Smith, and even on Fischer and Ravizza's own account of acting freely (i.e., their account of "guidance control," which is less demanding than many alternatives—especially libertarian—accounts of acting freely), Jones does not count as acting (or omitting) freely in the alternative sequence.

Second, even if there are cases in which an agent acts freely without satisfying some epistemic condition on moral responsibility, there is a reason to be skeptical that Jones genuinely *omits* to decide on his own at *t* to kill Smith in the alternative sequence. To see why, consider the number of things you are not doing right now. You are not reading the newspaper, whistling a tune, or doing a handstand, among countless other things. For any of these things that you are not now doing, are you *omitting* to do them? Arguably you are not, for, as Randolph Clarke argues:

There are vastly many instances of our not doing things—even things that we can do—that don't count as omissions, given the ordinary meaning of the term. Setting aside cases of intentionally not doing a certain thing, the difference between the instances that count and those that don't lies...in whether an action of the sort that is absent was called for by some norm, standard, or ideal. (2014: 33)

Since there was no norm, standard, or ideal calling for you to do a handstand just now, it would be odd to describe your not doing one as an omission of yours (especially considering that you presumably had not even thought about doing one then). But notice that the same is true of Jones; while he does not decide on his own at *t* to kill Smith, there is no norm, standard, or ideal calling for him to decide on his own at *t* to kill Smith, so it would be strange to describe Jones's behavior

epistemic condition are also cases in which agents are unfree, we have reason to be skeptical that any such epistemic condition will be forthcoming.

in the alternative sequence as a case of his omitting to decide on his own to kill Smith, much less a case of his *freely* omitting to decide on his own.

It is worth considering two possible responses to this last point about whether Jones genuinely omits to decide on his own to kill Smith. One response is to say that, while there is no norm calling for Jones to decide on his own to kill Smith, there is a norm calling for him to *omit* to decide on his own to kill Smith (after all, that's the best Jones can do under the circumstances), and this may be enough for Jones's not deciding on his own to kill Smith to count as a genuine omission. Recall that Jane falsely believes that she can kill Dan just by thinking *let Dan be dead* and yet omits to do so; here it seems that Jane genuinely omits, yet there is no norm calling for her to think *let Dan be dead*, only a norm calling for her to refrain from so thinking.¹¹ But, again, I think that there is an important difference between Jones and Jane, and it makes a relevant difference here too. Jane's not thinking *let Dan be dead* is on her radar, whereas Jones's best option, to omit to decide on his own to kill Smith, could not be reasonably expected to be on his radar. If he knew about Black's device, and so was aware that he would either decide on his own or as a result of Black's device, then I would agree that his not deciding on his own to kill Smith would count as a genuine omission, but without such awareness his not deciding on his own appears to be a mere absence of action, not a genuine omission.

A different response would be to opt for a less narrow view of what counts as an omission. Suppose that June does not murder her neighbors when she sees them. It may seem that she omits to murder them even though there is no norm, standard, or ideal that calls for her to commit murder, and it may seem that she is morally responsible for this omission.¹² My response to this suggestion

¹¹ Thanks to an anonymous reviewer for raising this response.

¹² Thanks to another anonymous reviewer for raising this response.

will depend on the details. If June desires to kill her neighbors, or has some other (pro tanto) reason for committing murder, then I am inclined to agree that she omits (and perhaps is morally responsible for the omission, though that may depend on how exactly the details are filled out); but in this version of the case it seems to me that there *is* a norm, standard, or ideal calling for the action, namely whatever reason she has for committing the murder. On the other hand, if June's murdering her neighbors is not something that she ever considers (or ever takes herself to have a reason for doing), then her not committing murder looks more like your not doing a handstand now; it's an absence of action that does not rise to the level of a genuine omission. And that's how I am thinking of Jones's not deciding on his own to kill Smith.¹³

6. *Robustness Revisited*

I have argued that Capes and Swenson's and Robinson's proposals fail because it is implausible that agents in the alternative sequences of FSCs act (or omit) freely. If we accept Fischer's robustness requirement (FRR), as Capes and Swenson and Robinson do, then the problem for the flicker strategy remains intact.¹⁴ It seems, then, that the flicker theorist should

¹³ One might think that it does not really matter whether Jones *omits* to decide on his own to kill Smith; perhaps an agent can freely abstain from performing some action even if that absence of action does not rise to the level of a genuine omission. Clarke (2014: 117) says that this is possible and gives an example in which it seems like an agent is morally responsible for not getting milk when, due to their inability to procure the milk at a certain time, they do not *omit* to get the milk then. It seems to me that moral responsibility for absences like this must always be indirect, or derivative, and that, for an agent to be directly morally responsible for not doing something, the absence must be an omission, otherwise we lack a way of delimiting the absences for which we may be morally responsible from the much larger class of absences of action in general, which is enormous. And it seems to me that the most plausible approach for delimiting those absences for which we may be morally responsible is exactly along the lines Clarke suggests, namely when the absent action is called for by some norm, standard, or ideal.

¹⁴ Capes and Swenson make it explicit that they accept FRR, and Robinson's attempt to show that agents in FSCs have alternatives that are within their control (2014: 439) strongly suggests that he accepts FRR as well.

reject FRR. This leaves the flicker theorist with two options: either to opt for a more minimal robustness requirement or to reject robustness requirements altogether. I will consider the first option in this section and the second in the next.

In light of the previous section, the most natural way to modify FRR into a more minimal robustness requirement would be to reject the second component of FRR—the component that says that an agent’s moral responsibility for what she does in the actual sequence is plausibly grounded in an alternative possibility (some alternative sequence) only if the agent acts freely in the alternative sequence. As we saw above, Fischer supports that component by pointing out that it would be very puzzling to see how an alternative possibility could plausibly ground the agent’s freedom in the actual sequence unless the agent acts freely in that alternative scenario. But the leeway theorist may deny that this is puzzling, maintaining instead that, even if the agent fails to act freely in the alternative sequence, this alternative can nonetheless ground the agent’s *ability to avoid being morally responsible for what she does in the actual sequence* (since the mere presence of the alternative guarantees the possibility that the agent does otherwise), and that this *avoidability* in turn grounds the agent’s moral responsibility for what she actually does.

The idea that alternative possibilities are significant insofar as they ground avoidability is not new.¹⁵ Leeway theorists are often motivated by the idea that, if an agent couldn’t help but perform some action, then her performing that action is not really up to her. If, on the other hand, an agent really could do otherwise, then it is not the case that she couldn’t help but perform the action, and this plausibly (partly) grounds her moral responsibility for what she actually does by providing her with the ability to avoid being responsible. But, while an agent’s ability to avoid being morally responsible requires that there be some alternative sequence available to her, it does

¹⁵ For more on this idea, see especially Otsuka (1998) and Widerker (2003).

not require that she act freely in that alternative sequence, so leeway theorists have not been given reason to accept the second component of FRR.

Given that leeway theorists need not accept as strong a robustness requirement as Fischer's, and given that they may take alternatives to be important only insofar as they ground an agent's ability to avoid being morally responsible for what she actually does, leeway theorists could adopt the following robustness requirement:

Minimal Robustness Requirement (MRR): An agent's alternative possibility is robust only if the agent's moral responsibility for what she does in the actual sequence is plausibly grounded in the alternative possibility; and an agent's moral responsibility for what she does in the actual sequence is plausibly grounded in an alternative possibility (some alternative sequence) only if that alternative grounds the agent's ability to avoid being morally responsible for what she does in the actual sequence.

The first component of MRR is identical to the first component of FRR, and I see no reason why the flicker theorist should reject it. The second component of MRR replaces Fischer's "freedom requirement" with the "avoidability requirement" motivated in the previous paragraph. I take it that many flicker theorists will already find this second component of MRR attractive. Moreover, because its second component is less demanding than FRR's, MRR is a simpler and more natural requirement for leeway theorists.

It is important to note that it will be impossible to construct successful FSCs without leaving the agent with alternatives that count as robust, according to MRR. If it is going to be intuitive that the agent is morally responsible for acting in a particular way, according to the flicker theorist, it will remain open that the agent does something else instead (or does the same action

but not on their own). And yet any such alternative sequence will inevitably be one that grounds the agent's ability to avoid being morally responsible for what she does in the actual sequence.

If the flicker theorist can accept a more minimal robustness requirement (like MRR), this would allow her to reject two alternative robustness requirements that have received a good deal of attention in the recent literature. The first is proposed by McKenna and says that a robust alternative must meet two conditions:

- (1) The alternative is morally significant. 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 the alternative scenario.
- (2) The alternative has to be within the control of the agent. (2003: 204)

Important for our purposes is McKenna's second condition. Are the alternative sequences in FSCs within the control of agents in FSCs? Arguably, they aren't, since these agent lack freedom in the alternative sequences. And yet, as we saw above, it is open to the flicker theorist to maintain that, insofar as they ground such agents' avoidability of responsibility, these alternatives can nonetheless play the required role in grounding agents' responsibility in the actual sequence.

A second alternative robustness requirement (which is itself a modified version of an earlier robustness requirement) is proposed by Pereboom:

Robustness (B): For an agent to have a robust alternative to her immoral action A, that is, an alternative relevant per se to explaining why she is blameworthy for performing A, it must be that

- (i) she instead could have voluntarily acted or refrained from acting as a result of which she would be blameless, and

(ii) for at least one such exempting acting or refraining, she was cognitively sensitive to the fact that she could so voluntarily act or refrain, and to the fact that if she voluntarily so acted or refrained she would then be, or would likely be, blameless. (2014: 13)

Just as flicker theorists should reject McKenna's second condition as too demanding (assuming that we have good reason for rejecting the second component of FRR), flicker theorists should reject Pereboom's first condition on robust alternatives as too demanding, for, flicker theorists should contend, agents in FSCs need not have access to alternatives in which they act or refrain from acting *voluntarily* in order for those alternatives to ground the ability to avoid responsibility. And since, as we saw above, it appears to be open to the flicker theorist to maintain that robust alternatives need only play the role of grounding agents' avoidability of responsibility.¹⁶

However, as I will now argue, because the flicker theorist is committed to a certain picture of the control required for moral responsibility (one that requires that agents could have done otherwise), flicker theorists must require more of robust alternatives than MRR does and so must ultimately reject MRR. The basic point, which takes us back to Fischer's initial support for FRR, is that the leeway theorist's account of freedom says that, in order for an action to be *up to* an agent (in the sense required for moral responsibility), that agent must have the power to select from more than one option, each of which can be brought about freely. By attempting to ground agents' actual responsibility by appealing to alternatives in which they do not act freely, the flicker theorist is failing to capture the picture of control that motivated acceptance of PAP in the first place.

For another statement of this objection (or something close to it), consider the following from Fischer:

¹⁶ For a more detailed discussion of Pereboom's robustness requirement, see Kittle (2019).

The point might be put as follows. The proponent of the idea that regulative control is required for moral responsibility insists that there can be no moral responsibility, if there is but one path leading into the future: to get the crucial kind of control, we must add various alternative possibilities. Now it seems that the flicker theorist must claim that the addition of the sort of alternative possibility he has identified would transform a case of lack of responsibility into one of responsibility. But this seems mysterious in the extreme: how can adding an alternative scenario (or perhaps even a set of them) in which Jones does not *freely* [decide to kill Smith] make it true that he actually possesses the sort of control required for him to be morally responsible for his [deciding to kill Smith]? This might appear to involve a kind of *alchemy*, and it is just as incredible. (1994: 141)

It may be thought that, if an agent has the power to select from more than one option, each of which can be brought about freely, then she may possess a certain sort or degree of control that would not be possible without alternative possibilities. Fischer's worry for the flicker theorist, however, is that alternatives in which agents do not act freely could not give rise to that sort or degree of control. The worry for MRR, then, is that it is so minimal a robustness requirement on alternatives that it gives up the picture of control with which PAP is associated.

To see why the flicker theorist (*qua* leeway theorist) must accept a more substantial robustness requirement than MRR, it is worth asking why a leeway theorist should accept any robustness requirement on alternative possibilities at all—a point to which we'll return in the next section. Here it will be helpful to consider Carolina Sartorio's discussion of the role of the alternatives on the leeway (or alternative-possibilities) view:

...the robustness requirement is more than a well-motivated condition that an alternative-possibilities view would have to meet in order to be an adequate view of freedom. Rather,

it's a *constitutive* condition of any such view...the alternatives that such views require will automatically have to be robust, in the sense that they will have to be something in virtue of which the agent is free. "Robustness" may sound like a mysterious metaphysical term but, in fact, it's nothing more than a reminder that the alternatives required by an alternative-possibilities view are supposed to be such that the agent can be free in virtue of having them. (2016: 15)

As we saw above (in section 2), Fischer thinks that the leeway theorist should require alternatives to play a grounding role in an agent's responsibility because such a theorist takes alternative possibilities to be crucial to the control required for moral responsibility. As Sartorio points out, there is more to say here, for it is a constitutive condition of the leeway theorist's view that alternative possibilities ground agents' freedom and responsibility. And this is a constitutive condition on the leeway theorist's view because non-leeway views (sourcehood, or actual-sequence views) are distinguished from the leeway position by *not* grounding agents' freedom and responsibility in alternative sequences.¹⁷ Everyone should agree, then, that an agent's alternative possibility is robust only if the agent's moral responsibility for what she does in the actual sequence is plausibly grounded in the alternative possibility.

But the mere possibility of something else happening to an agent in an alternative sequence is insufficient to ground the agent's moral responsibility in the actual sequence. To see this, consider the following scenario. Suppose that causal determinism precludes leeway freedom, and suppose that an agent who is causally determined to perform a particular action satisfies all conditions on moral responsibility minus having leeway freedom (if this is a condition on moral

¹⁷ For a full discussion of this point, see Sartorio (2016: 7-16; cf. Forthcoming), and see Fischer (2018) for related points.

responsibility). Now suppose that the case is tweaked such that the agent's action is not causally determined, but suppose that the reason the action is not determined is because there is a slight but non-negligible chance that the agent will be incapacitated then.¹⁸ If the agent whose action is causally determined is not morally responsible for that action, then neither is the agent whose action is undetermined in the way just described. Because mere alternatives, construed as the possibility of something else happening to the agent, are not such that agents are free in virtue of having them, they are insufficient to ground freedom on the leeway model.

Now, perhaps the sort of alternatives needed for genuine avoidability of responsibility in the actual sequence must be plausibly construed as *behavior* of the agent's (whether an action or omission), not simply an alternative in which something else *happens* to the agent. Even this, however, is too minimal, for we can suppose that an agent's action is undetermined and that, in the alternative sequence, the agent is forced by some other agent to act against their will (perhaps to perform some alternative action which would have been unthinkable for them); and, in such a case, the possibility of the agent *acting* otherwise (at the hands of some other agent) does not plausibly ground the agent's freedom in the actual sequence any more than the possibility of something else *happening* to the agent would.

What is needed to ground an agent's freedom on the leeway model, then, is an alternative sequence in which the agent freely does (or omits) otherwise than they do in the actual sequence. But this is exactly the requirement that is codified in Fischer's robustness requirement (FRR)—as well as McKenna's and Pereboom's. We have been exploring the possibility of tinkering with these stronger robustness requirements in order to circumvent the difficulties I raised for the

¹⁸ This scenario is similar to some of Mele's (2005) variations on Pereboom's (2001) original manipulation cases, and also to "Fischer-type examples" (Fischer 2006: 150-151).

versions of the flicker strategies considered earlier in this paper, but it turns out that the flicker theorist (*qua* leeway theorist) cannot accept a more minimal robustness requirement like MRR.

7. Robustness as Misplaced Demand?

So far I have considered how the flicker theorist might supplant previously proposed robustness requirements with a more minimal one, but another strategy the flicker theorist might try is to reject requirements of robustness altogether. In fact, in addition to developing his robust flickers of freedom proposal, Robinson has in earlier work argued the demand that alternatives be robust is misplaced given the dialectic:

PAP maintains that agents are (basically) morally responsible for what they have done only if they could have done otherwise. Proponents of Frankfurt-type counterexamples (both in their modified and traditional forms) argue that these cases falsify PAP. The aim of the flicker strategy is to rebuff this attack, which it does by arguing that, since agents in these cases could have done otherwise than that for which they are (basically) morally responsible, these cases are not genuine counterexamples to PAP. But, then, the demand at this point that flicker strategists show not only that agents in these cases do in fact possess alternative possibilities—that is, not only that Frankfurt-type scenarios fail to falsify PAP—but that these alternatives are relevant to agents’ moral responsibility for what they do in these cases seems misplaced. (2012: 191; cf. 2019: 220-221)

If the flicker theorist’s aim is only to rebuff an attack on PAP, then showing that agents in FSCs inevitably have some alternative sequence open to them seems to suffice. Admittedly, it would be nice for the flicker theorist to go on to spell out why alternative possibilities are relevant to moral

responsibility, but that (positive) project is independent, Robinson thinks, of the (negative) project of showing that FSCs fail.¹⁹

But, given what we have seen about the rationale for robustness requirements in the previous section, it is fairly straightforward why the robustness requirement is not dialectically inappropriate. Flicker theorists aim to defend a conception of freedom—the leeway model—according to which alternative possibilities are such that agents are free in virtue of having them. If FSCs can succeed in showing that agents may be morally responsible without having alternatives that could plausibly ground the agents’ moral responsibility, this is a problem for the flicker theorist just as a full-fledged counterexample to PAP would be a problem. To put the point another way, while much of the dialectic since the publication of Frankfurt’s classic paper has focused on the (sometimes very technical) details concerning whether any FSCs provide a genuine counterexample to PAP, some have used the debate about FSCs as an opportunity to weigh the two competing models of freedom, namely leeway and sourcehood models; and, insofar as FSCs bring out the plausibility that the grounds of an agent’s freedom and responsibility lie within the actual sequence (and not in her power to select from among alternative possibilities), this counts in favor of the sourcehood model and against the leeway model. Flicker theorists, then, in order to preserve not only PAP but their very model of freedom, must accept a (more-than-minimal) robustness requirement on alternative possibilities.

¹⁹ It may be that FSCs do some work in demotivating the leeway model of freedom even if they do not succeed in providing counterexamples to PAP. In fact, Frankfurt (2003) has suggested this point, and several others have developed it, including Leon and Tognazzini (2010), McKenna (2008), Sartorio (), and Zagzebski (2000). If this is right, it is not clear that showing that agents in FSCs have some alternative suffices to rebuff the attack on PAP (and the leeway conception of freedom with which it is associated).

8. Conclusion

Interestingly, as the attentive reader will have noticed, all of the proposals for robustness requirements on alternative possibilities discussed here have been proposed by *sourcehood* theorists (Fischer, McKenna, Pereboom, and myself)—theorists who don't think robust alternatives are required for moral responsibility. After proposing their accounts of robustness, these theorists go on to argue that agents in FSCs lack relevantly robust alternatives, and so the flicker of freedom strategy fails. The typical response from *leeway* theorists attracted to the flicker of freedom strategy (e.g., Capes and Swenson, Robinson) is to maintain that agents in FSCs have alternatives that are robust in the sense specified by their dialectical interlocutor's robustness requirement. In this paper, after objecting to specific proposals from such flicker theorists, I have considered whether flicker theorists might want to take a different strategy, one that involves spelling out an alternative, and more minimal, robustness requirement and then showing that this requirement can do the work that the flicker of freedom strategy requires of it. In the end, I followed in the tradition of proposing a robustness requirement and then arguing that this proposed robustness requirement will not ultimately allow the flicker theorist to escape the challenge raised by FSCs. But the problems identified here for a more minimal robustness requirement highlights what exactly is at issue in the debate between leeway and sourcehood theorists, namely the grounds of our freedom.

Acknowledgements

Thanks to Zac Bachman, Dave Beglin, Justin Capes, Luca Ferero, John Fisher, Andrew Law, Meredith McFadden, Jonah Nagashima, Debbie Nelson, Jeremy Pober, Jared Smith, Philip Swenson, and two anonymous reviewers for *Journal of Ethics*.

References

- Capes, J., and P. Swenson. 2017. Frankfurt cases: The fine-grained response revisited. *Philosophical Studies* 174: 967-981.
- Clarke, R. 2014. *Omissions: Agency, metaphysics, and responsibility*. New York: Oxford University Press.
- Fischer, J. M. 1994. *The metaphysics of free will: An essay on control*. Oxford: Blackwell.
- Fischer, J. M. 2006. *My way: Essays on moral responsibility*. New York: Oxford University Press.
- Fischer, J. M. 2018. The freedom required for moral responsibility. In *Virtue, happiness, and knowledge: Themes from the work of Gail Fine and Terence Irwin*, eds. D. Brink, S. Meyer and C. Shields. Oxford: Clarendon Press.
- Fischer, J. M., and M. Ravizza. 1998. *Responsibility and control: A theory of moral responsibility*. Cambridge: Cambridge University Press.
- Frankfurt, H. 1969. Alternate possibilities and moral responsibility. *Journal of Philosophy* 66: 829-839.
- Frankfurt, H. 2003. Some thoughts concerning PAP. In *Moral responsibility and alternative possibilities: Essays on the importance of alternative possibilities*, eds. D. Widerker and M. McKenna. Burlington, VT: Ashgate Publishing Company.
- Kittle, S. 2019. When is an alternative possibility robust? *European Journal of Philosophy* 27: 199-210.
- McKenna, M. 2003. Robustness, control, and the demand for morally significant alternatives: Frankfurt examples with oodles and oodles of alternatives. In *Moral responsibility and alternative possibilities: Essays on the importance of alternative possibilities*, eds. D. Widerker and M. McKenna. Burlington, VT: Ashgate Publishing Company.
- McKenna, M. 2008. Frankfurt's argument against alternative possibilities: Looking beyond the examples. *Nous* 42: 770-793.
- Mele, A. 2005. A critique of Pereboom's 'Four-Case Argument' for incompatibilism. *Analysis*, 65: 75-80.
- Mele, A. 2010. Moral responsibility for actions: Epistemic and freedom conditions. *Philosophical Explorations* 13: 101-111.
- Mele, A., and D. Robb. 1998. Rescuing Frankfurt-style cases. *Philosophical Review* 107: 97-112.
- Naylor, M. 1984. Frankfurt on the principle of alternate possibilities. *Philosophical Studies* 46: 249-258.
- O'Connor, T. 2000. *Persons and causes: The metaphysics of free will*. New York: Oxford University Press.
- Otsuka, M. 1998. Incompatibilism and the avoidability of blame. *Ethics* 108: 685-701.
- Pereboom, D. 2001. *Living without free will*. Cambridge: Cambridge University Press.
- Pereboom, D. 2014. *Free will, agency, and meaning in life*. New York: Oxford University Press.
- Robinson, M. 2012. Modified Frankfurt-type examples and the flickers of freedom. *Philosophical Studies* 157: 177-194.
- Robinson, M. 2014. The limits of limited-blockage Frankfurt-style cases. *Philosophical Studies* 169: 429-446.
- Robinson, M. 2019. Robust flickers of freedom. *Social Philosophy and Policy* 36: 211-233.
- Sartorio, C. 2016. *Causation and free will*. New York: Oxford University Press.
- Sartorio, C. 2017. Frankfurt-style examples. In *The Routledge Companion to Free Will*, eds. K. Timpe, M. Griffith, and N. Levy. New York: Routledge.

- Sartorio, C. 2019. Flickers of freedom and moral luck. *Midwest Studies in Philosophy* 43: 93-105.
- Sartorio, C. Forthcoming. The grounds of our freedom. *Inquiry*.
- Speak, D. 2002. Fanning the flickers of freedom. *American Philosophical Quarterly* 39: 91-105.
- Stockdale, B. 2022. Moral responsibility, alternative possibilities, and acting on one's own. *Journal of Ethics* 26: 27-40.
- Swenson, P. 2019. Luckily, we are only responsible for what we could have avoided. *Midwest Studies in Philosophy* 43: 106-118.
- van Inwagen, P. 1978. Ability and responsibility. *The Philosophical Review* 87: 201-224.
- van Inwagen, P. 2000. Free will remains a mystery. *Philosophical Perspectives* 14: 1-19.
- Widerker, D. 2003. Blameworthiness and Frankfurt's argument against the principle of alternative possibilities. *Moral responsibility and alternative possibilities: Essays on the importance of alternative possibilities*, eds. D. Widerker and M. McKenna. Burlington, VT: Ashgate Publishing Company.
- Zagzebski, L. 2000. Does libertarian freedom require alternate possibilities? *Philosophical Perspectives* 14: 231-248.