

UC Riverside

UC Riverside Previously Published Works

Title

Neo-Frankfurtians and buffer cases: the new challenge to the principle of alternative possibilities

Permalink

<https://escholarship.org/uc/item/9q659506>

Journal

Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition, 152(2)

ISSN

1573-0883

Author

Franklin, Christopher Evan

Publication Date

2011

DOI

10.1007/s11098-009-9472-9

Peer reviewed

Neo-Frankfurtians and buffer cases: the new challenge to the principle of alternative possibilities

Christopher Evan Franklin

Published online: 11 November 2009

© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract The debate over whether Frankfurt-style cases are counterexamples to the principle of alternative possibilities (PAP) has taken an interesting turn in recent years. Frankfurt originally envisaged his attack as an attempting to show that PAP is false—that the ability to do otherwise is not necessary for moral responsibility. To many this attack has failed. But Frankfurtians have not conceded defeat. Neo-Frankfurtians, as I will call them, argue that the upshot of Frankfurt-style cases is not that PAP is false, but that it is explanatorily irrelevant. Derk Pereboom and David Hunt’s buffer cases are tailor made to establish this conclusion. In this paper I come to the aid of PAP, showing that buffer cases provide no reason for doubting either its truth or relevance with respect to explaining an agent’s moral responsibility.

Keywords Frankfurt-style cases · Moral responsibility · Principle of alternative possibilities · Abilities · Tracing

Almost 40 years ago Harry Frankfurt (1969; reprinted 1988) crafted an ingenious case which purportedly described a scenario in which an agent was morally responsible for the action he performed and yet was unable to have done otherwise. Frankfurt’s key insight was that it seems possible for there to be circumstances which render it impossible for an agent to do otherwise and yet these same circumstances play no role in bringing about what the agent does. In such cases it appears that the agent acts on his own and thus is morally responsible, despite lacking the ability to perform any alternative action. These cases, and their variants,

C. E. Franklin (✉)

Department of Philosophy, University of California, Riverside, HMNSS Building 1604,
900 University Avenue, Riverside, CA 92521, USA
e-mail: Cfran002@ucr.edu

have come to be known as ‘Frankfurt-style cases’ or ‘FSCs’ for short. Many, however, and the present author included, resist conceding that FSCs demonstrate the falsity of the principle of alternative possibilities:

(PAP) Necessarily, an agent is morally responsible for φ only if the agent was able to have done otherwise.¹

For many, the truth of PAP is simply too obvious and secure within our moral theorizing to be overturned by some rather peculiar cases. Some have argued that PAP follows from the moral maxim ‘*ought* implies *can*’ (Copp 2003; Yaffe 2005), while others have maintained that PAP is embedded in our ordinary practices of excuses and exemptions (Glover 1970). Regardless of the route, many have and still do find PAP and its ilk immensely plausible. This commitment to PAP has led them to try to show that FSCs fail to truly describe a case in which an agent is morally responsible and yet unable to do otherwise. The debate has spawned a prolific and exceedingly interesting literature.² The hotly debated question had been whether FSCs falsify PAP. In order to defend PAP against this attack one only had to identify a single remaining alternative that the agent was able to realize. In recent years, however, proponents of FSCs have re-envisaged their attack. According to neo-Frankfurtians, as I will call them, instead of FSCs calling into question the truth of PAP, they call into question its relevance.³

Neo-Frankfurtians claim that PAP is intended to identify a condition that (partly) explains an agent’s moral responsibility.⁴ Defenders of PAP are concerned not just with its truth, with whether the ability to do otherwise is necessary for moral responsibility, but also with its relevance for explaining moral responsibility. As Leon and Tognazzini (forthcoming) have pointed out, there are two ways in which FSCs can succeed. They can succeed by showing that PAP is false—that the ability to do otherwise is not necessary for an agent’s moral responsibility. But they also can succeed by showing that PAP is irrelevant to explaining an agent’s moral responsibility—even if the ability to do otherwise is necessary.⁵ Consequently, in

¹ PAP has had numerous formulations, partly in attempt to refine it against Frankfurt’s attack, but this simple formulation will serve for present purposes.

² See Widerker and McKenna (2003) for a helpful introduction.

³ I will sometimes, for convenience, speak of PAP being explanatorily relevant, whereas strictly speaking we are concerned with whether someone’s *satisfying* PAP is relevant to explaining his moral responsibility.

⁴ Clearly no one thinks that such an ability is sufficient grounds for moral responsibility. In addition one must meet certain conditions of cognitive competence. At most, the ability to do otherwise is *part* of what explains the agent’s moral responsibility.

⁵ As I have characterized neo-Frankfurtians, they dispute not the truth of PAP, but its relevance. One might, however, think it is best to construe both traditional and neo-Frankfurtians as challenging the truth of PAP. According to this interpretation, the difference between traditional and neo-Frankfurtians concerns how they understand the content of the principle. Traditional Frankfurtians think that PAP merely claims that the ability to do otherwise is necessary for moral responsibility; whereas, neo-Frankfurtians think PAP claims, or should claim, that the ability to do otherwise is both necessary and explanatorily relevant. Nothing of substance will turn on these different characterizations of the debate.

order to defend PAP against the neo-Frankfurtian attack one must do more than show that FSCs fail to eliminate all abilities to do otherwise: one must go onto show that FSCs give us no reason to doubt that these remaining abilities (partly) explain the agent's responsibility. And it is this second task that neo-Frankfurtians dispute. Notably, Pereboom (2001), and independently Hunt (2005), have developed FSCs, known as buffer cases, which purport to show that an agent is morally responsible and this is not even partly explained by the fact that the agent was able to do otherwise. Many have thought that buffer cases are devastating to the supposed relevance of PAP (cf. Baker 2006).

In this paper I will argue that buffer cases give us no reason to judge that PAP is irrelevant to explaining moral responsibility. The structure of the dialectic is important to keep in mind. I am not here concerned with (directly) arguing that PAP is either true or relevant to moral responsibility.⁶ Everything that I argue for in this paper is consistent with both the falsity and irrelevance of PAP as I am concerned exclusively with showing that buffer cases pose no challenge to the relevance of PAP. This is an important conclusion as it does seem that we have good reasons for thinking that PAP is true and relevant to moral responsibility. As mentioned above, many have argued that PAP, or a close relative, follows from the moral maxim that '*ought implies can*'. Moreover, there is reason to think that PAP is embedded in our ordinary practices of excuses and exemptions. If it is true that someone genuinely, through no fault of her own, was unable to have done otherwise, then we normally excuse her from moral responsibility. I assume, then, that there is some weight in favor of PAP, and thus FSCs in general, and buffer cases in particular, should be construed as attempting to defeat or outweigh these initial considerations. My task is to show that buffer cases do no such thing.

I begin by clearing up what is at issue in debates concerning PAP-like principles and elucidating the neo-Frankfurtians challenge. I then develop a detailed criticism of such cases arguing that they fail to pose any threat to the explanatory relevance of the ability to do otherwise.

1 The principle of alternative possibilities

PAP's name and frequent formulation derives from Frankfurt's famous criticism of it (Frankfurt 1969/1988).⁷ However, the above formulation is subject to multiply ambiguities which need to be cleared up before we can accurately assess its truth

⁶ I am indirectly defending its truth and relevance insofar as I am arguing that a reason that is often appealed to for thinking that PAP is false and/or irrelevant is really no reason at all.

⁷ This really is an unfortunate name for the principle as it tends to lead people to think that what is at stake is whether certain alternative sequences might have occurred; whereas, what we are really concerned with is whether the *agent* was able to have brought about these alternative sequences. But the name appears too firmly embedded in the literature to have any chance of being over turned. It thus is important that we keep in mind that the defender of PAP thinks that moral responsibility does not merely require that something else could have happened, but that the agent was able to bring this something else about.

and relevance—particularly with respect to the neo-Frankfurtian’s attack. At one extreme, the principle seems too strong. If we disambiguate it in a way that makes it require that an agent be able to do otherwise at the moment of action, then there are straightforward counterexamples. Martin Luther’s famous, “Here I stand, I can do no other”, if taken at face value serves as one counterexample (cf. Dennet 1984, p. 133). Luther is morally responsible for the action of refusing to recant even though it was impossible for him to do otherwise at the moment of recanting. At the other extreme, the principle seems too weak. If we disambiguate it in such a way that makes it require only that the agent be able to have done otherwise at some moment in his life, then it is hard to understand why such an ability is necessary for moral responsibility. That Luther was able to order wine rather than beer for dinner is hardly necessary for his being morally responsible for not recanting.

I suggest that the following principle makes important progress in disambiguating PAP:

PAP₁: Necessarily, an agent S is morally responsible for φ -ing at t_2 only if S (i) was able at t_2 to avoid responsibility for φ -ing at t_2 or (ii) if S is unable at t_2 to avoid responsibility for φ -ing at t_2 , then there is some earlier time t_1 at which S ψ -ed (where ψ might be an action or omission) that explains why S is unable at t_2 to avoid responsibility for φ -ing at t_2 and S was able at t_1 to avoid responsibility for ψ -ing at t_1 .

PAP₁ makes it clear that what is at stake in principles like PAP is the ability to avoid responsibility for some event rather than being able to bring about some alternative event.⁸ What is crucial for moral responsibility is not that I am able to prevent the action from occurring, but that I am able to avoid responsibility for its occurring.⁹ Another virtue of PAP₁ is that it disambiguates PAP in a way that avoids both of the above extremes and settles on a more plausible intermediate. It weakens, but, so I would suggest, not too much, the relevance of avoidability for moral responsibility by introducing a “tracing” condition. That Luther was unable to avoid responsibility for recanting at the moment of action does not preclude him from being morally responsible. According to PAP₁, he may still be morally responsible just so long as his current state of inability traces back to—i.e., is explained by or grounded in—earlier actions he performed or failed to perform, and moreover he could have avoided responsibility at the moment of action for these earlier actions. So PAP₁ only entails that moral responsibility requires avoidability at the moment of action for certain actions. I will come back to the notion of tracing in the final section of this paper as it plays a central role in my defense of PAP₁. But let us now turn to buffer cases.

⁸ Cf. McKenna (1997), Wyma (1997), and Otsuka (1998) who argue for a similar understanding of PAP. The ability to avoid responsibility for φ is a narrower ability than the ability to do otherwise than φ . If one has the ability to do otherwise than φ , then one has the ability to avoid responsibility for φ , but not vice versa.

⁹ Of course, preventing the action from occurring is one, albeit slightly awkward, way of avoiding responsibility for it: one cannot be responsible for what did not happen.

2 Buffer cases

It is well-known that the prospects of devising a scenario in which an agent is morally responsible and yet completely unable to avoid responsibility are bleak.¹⁰ Rather than encumber the reader with this intricate debate I will simply assume that it is highly unlikely that such cases are possible and explore whether, if this is true, the Frankfurtian must concede the debate to the proponent of PAP₁. Neo-Frankfurtians think not. According to them, FSCs call into question the relevance of PAP₁. FSCs allow us to see that the ability to avoid responsibility is simply irrelevant to explaining an agent's moral responsibility. Consider David Hunt's buffer case, *Revenge*:

Jones, while attending a party, is deliberately humiliated by Smith. The first thought that occurs to Jones, after realizing what Smith has done to him, is that he would like to kill Smith. He leaves the party, escaping the immediate pressures of the situation and giving himself ample opportunity to pull back from this line of thinking. Given the kind of person Jones is, and given the situation in which he finds himself, the alternative of *not* killing Smith is not unthinkable for him; moreover, should he entertain this alternative, nothing would prevent him from deciding and acting on it. But Jones could decide (and act) otherwise only if he first considered acting otherwise, and he never does this (though he could); instead, he nurses his grievance without respite, while the idea of killing Smith becomes more and more attractive to him. Having finally decided to do the deed, he gets a gun from his car, returns to the party, and shoots Smith dead....The final element to be added to *Revenge* is the counterfactual intervener, which differs from the device in [earlier FSCs] inasmuch as it is programmed to hijack Jones's mental processes and force him to decide to kill Smith if he so much as *considers* not killing Smith. With this device in place, there is no alternative to Jones's deciding to kill Smith: Jones can decide otherwise only if he first considers doing so, but then the device will force him to decide *to* kill Smith. So an alternative decision is not available to Jones in *Revenge*...[y]et Jones, who in fact proceeds to murder Smith on his own, leaving the device untriggered, seems morally responsible for killing Smith. (Hunt 2005, pp. 132–135)¹¹

There are three features of buffer cases that need highlighting: (i) the prior sign that signals Black's need for intervention is merely a necessary condition for Jones's *not* deciding to kill Smith, (ii) this necessary condition can obtain at any time prior to his decision to kill Smith, and (iii) Jones is able to bring about the necessary condition.¹²

¹⁰ I say bleak, but not indefensible. See Fischer (2006) for just such a defense. The bleakness of contriving such an example is shown by the Kane–Widerker dilemma. See Kane (1996, pp. 142–143) and Widerker (1995).

¹¹ See also Pereboom (2001, pp. 18–37) for his *Tax Evasion* case which is structurally identical.

¹² These three features distinguish buffer cases from traditional FSCs and blockage cases. For examples of the latter kind of cases see Mele and Robb (1998, 2003).

First, Jones's considering not killing Smith, in the absence of Black, is not sufficient for his not deciding to kill Smith. There are possible worlds in which Jones considers not killing Smith, and yet still goes on to decide to kill Smith on his own and hence is responsible for such a decision.¹³ Black, however, does not want to take any chances. Even though Jones's considering not killing Smith is only necessary for his not deciding to kill Smith, Black decides to set things up so that if Jones so much as even considers not killing Smith, Black will intervene and force Jones to decide to kill Smith. Second, it is possible that Jones considers not killing Smith at any moment up to his deciding to kill Smith. This feature makes it clear that there is no condition, prior to action, which causally determines Jones's deciding to kill Smith. If there were such a condition, then Jones's responsibility would be highly questionable. Third, whether this sign is manifested is up to Jones in the sense that he is able to consider not killing Smith and able not to consider not killing Smith. These last two features of buffer cases are intended to make it clear that the circumstances which make it impossible for Jones to do other than decide to kill Smith play no role in bringing about what Jones does.

What should we make of such cases? It does seem plausible to judge that Jones is morally responsible for deciding to kill Smith, but then again, it is also clear that Jones was able to avoid responsibility for deciding to kill Smith. Recall that Jones was able to consider not killing Smith and if he did this Black would have intervened, forced Jones to decide to kill Smith, and thus Jones would escape responsibility for this decision.¹⁴ So although normally (i.e. in the absence of Black) Jones's considering not killing Smith is consistent with Jones's being responsible for deciding to kill Smith, given Black's presence, Jones's exercising his ability to consider not killing Smith is sufficient for his avoiding responsibility for deciding to kill Smith. Does this pose a problem for buffer cases? Hunt and Pereboom think not since, after all, the neo-Frankfurtian challenge is to the relevance of PAP₁, not its truth. Pereboom explains:

it may be that a condition is necessary for some phenomenon A but sometimes holds by virtue of features that do not illuminate A, features that are irrelevant to explaining the nature of A. I believe that even if it turns out that [the availability of alternative possibilities] are necessary for moral responsibility, they can hold by virtue of features of a situation that are explanatorily irrelevant to what would make an agent morally responsible, and that as a result the condition at issue fails to illuminate the nature of the phenomenon. (Pereboom 2001, p. 25)

To illustrate that these two things—a necessary condition and an explanatory condition—can come apart, notice that many conditions which are necessary for moral responsibility appear to be explanatorily irrelevant. For example, that $2 + 2 = 4$ is a necessary condition for an agent's responsibility for φ -ing since

¹³ Of course Black will be absent in such a world.

¹⁴ I will move seamlessly between 'the ability to consider not killing Smith' and 'the ability to avoid responsibility for deciding to kill Smith' throughout the paper since, in this context, the latter ability is constituted by the former ability.

there is no possible world in which the agent is responsible for φ -ing and that $2 + 2 \neq 4$. Nonetheless, it is clear that this mathematical truth does not even partly explain why the agent is morally responsible. The relevant question, then, is not whether Jones was able to avoid responsibility, but whether it is plausible to think that this ability partly explains why Jones is responsible for deciding to kill Smith.¹⁵

Hunt and Pereboom think it is clear that it does not. They suggest the following obstacle to this ability's relevance for explaining Jones's responsibility. Consider Jones's epistemic situation. Black is a covert manipulator. Jones does not understand that his considering not killing Smith will trigger the intervener, subsequently causing him to decide to kill Smith and thus avoid responsibility for this decision. In the absence of Black, exercising this ability is merely necessary for avoiding responsibility for not deciding to kill Smith—Jones's considering not killing Smith is consistent with his being responsible for deciding to kill Smith. Hence Jones does not understand that by making such a consideration seem salient he will *thereby* avoid responsibility for deciding to kill Smith, he does not understand that, in the present context, his considering not killing Smith is inconsistent with his being responsible for deciding to kill Smith. Because of this, the ability to consider not killing Smith is irrelevant to explaining his responsibility. Let us refer to this as the 'epistemic condition for relevance': an ability to do otherwise is relevant to explaining why an agent is morally responsible for φ -ing only if the agent understands that in exercising this ability he would thereby avoid moral responsibility for φ -ing.¹⁶

¹⁵ I am skeptical that mathematical truths are relevantly analogous to the ability to do otherwise with regard to being necessary conditions for moral responsibility, and thus while I concede that a necessary condition and explanatory condition can come apart, it is a mistake to construe defenders of the relevance of PAP as conflating these two notions. First, it is part of our pre-theoretical understanding of moral responsibility that it requires avoidability. The same cannot be said for mathematical and logical truths such as that $2 + 2 = 4$. If the person who stepped on our toes convinces us that he was not able to avoid responsibility for stepping on our toes, then we are inclined to excuse him from responsibility. Second, that Jones was able to consider not killing Smith and hence avoid responsibility for deciding to kill Smith is only contingently true, unlike the supposed necessary truths of mathematics and logic. The set of possible worlds in which $2 + 2 \neq 4$ is empty and hence it is far from surprising or interesting to point out that there are no worlds in which $2 + 2 \neq 4$ and an agent is morally responsible. There are no worlds in which $2 + 2 \neq 4$ and so clearly there are no possible worlds in which $2 + 2 \neq 4$ and someone is morally responsible. There are, however, worlds in which Jones is unable to avoid responsibility for deciding to kill Smith and so it is interesting and perhaps surprising that there are no worlds in which Jones is responsible for deciding to kill Smith and yet unable to avoid responsibility for deciding to kill Smith. This should make it clear that defenders of the explanatory relevance of PAP do not mistakenly infer the relevance of the ability to do otherwise from the fact that it is a necessary condition. Rather, it is the fact that the ability to do otherwise is a *distinctive* kind of necessary condition that leads people to think that it is explanatorily relevant.

¹⁶ This principle is borrowed from Pereboom (2001, p. 26), in particular the notion of understanding. Pereboom never states exactly what he has in mind by this notion. I will assume that all that is required for understanding that p is that someone possesses a certain set of dispositions. So an agent understands that p if and only if she stands ready to assert that p or take p as a premise in further reasoning and so on. In this way one can understand things that one never thinks about. I understand that if the sun is yellow all over it cannot simultaneously be orange all over. I understand this, not because I have ever entertained this proposition, but because were you to ask me if this is true, I would say 'yes,' I would employ this proposition as a premise in an argument and so on.

To bring out the plausibility of the epistemic condition consider an analogous case. Imagine that Jones was able to take a sip of coffee prior to deciding to kill Smith, and he did not understand that by drinking the coffee he would thereby avoid responsibility since, unbeknownst to him, there is a deadly form of poison in the coffee that will instantly kill him upon consumption. Jones was able to take a sip of the coffee and had he done this he would have avoided responsibility for deciding to kill Smith since he would have died and thus not made the decision. But it seems simply incredible to think that in this case Jones's responsibility is explained by that fact that he was able to take a sip of coffee since he did not understand the consequences of this action (cf. Pereboom 2001, p. 26).¹⁷ The epistemic condition does indeed seem plausible.

Jones's considering not killing Smith appears clearly to violate the epistemic condition and so it seems that *Revenge* succeeds at showing that PAP₁ is irrelevant to explaining Jones's responsibility. Although Jones was able to avoid responsibility for deciding to kill Smith, and thus Jones satisfies PAP₁, his ability fails to explain why he is responsible for deciding to kill Smith. Buffer cases provide, as it were, a buffer zone between explanatorily irrelevant and relevant abilities. In order for Jones to possess an ability that is relevant to explaining his moral responsibility—such as an ability to refrain from deciding to kill Smith—he must first exercise his ability to consider not killing Smith. But once he does this Black will immediately intervene and force him to decide to kill Smith. Black's presence, therefore, appears to sever any access Jones has to alternatives which are relevant to explaining his responsibility without Black's presence itself bringing about what Jones does. If this is correct, then buffer cases provide good reason to conclude that PAP₁ is irrelevant to explaining responsibility.

3 The failure of buffer cases

Even though buffer cases mount an important challenge to the relevance of PAP₁, I believe that they ultimately fail: they give us no reason to doubt that Jones's ability to consider not killing Smith, and thus his ability to avoid responsibility, partly explains why he is responsible for deciding to kill Smith. But before I argue for this contention I need to make salient a feature of buffer cases that is often overlooked.

Returning to *Revenge*, suppose that Jones decided to kill Smith at t_3 .¹⁸ Was he able at t_3 to do otherwise than decide to kill Smith? Perhaps surprisingly the

¹⁷ Hunt also suggested to me a similar line of argument during the Q/A time when I presented an earlier version of this paper at 2007 Society of Christian Philosophers, Mountain-Pacific Region Meeting. I take this to indicate that Hunt accepts (or at least is inclined to accept) the epistemic condition for relevance.

¹⁸ The time index in 'Jones decided to kill Smith at t_3 ' is ambiguous: it might mean 'Jones [decided to kill Smith at t_3]' or 'Jones [decided to kill Smith] at t_3 '. In the latter case the index attaches to the time of the decision and in the former it attaches to the content of the decision. Similar points apply to 'Jones considered not killing Smith at t_3 .' In both cases the time index is to be disambiguated in the latter way. Thus, 'Jones decided to kill Smith at t_3 ' is equivalent to 'Jones [decided to kill Smith] at t_3 ' and 'Jones considered not killing Smith at t_3 ' is equivalent to 'Jones [considered not killing Smith] at t_3 '. Although these disambiguations are clearer I find them awkward and thus will continue to employ the ambiguous, albeit more natural, wording.

answer is ‘Yes’. Recall that the second and third salient features of buffer cases guarantee that Jones was able, at any moment up to his making the decision, to consider not killing Smith. These features of buffer cases make it clear that at no time prior to Jones’s decision was there a sufficient condition for his making the decision and hence his decision to kill Smith was not determined. Whether an agent can be responsible for a decision that is determined by factors beyond his control is a hotly contested question and much of the force of FSCs would be lost if they depended on a particular answer to it. But, then, this means that instead of deciding to kill Smith at t_3 , Jones was able at t_3 to consider not killing Smith at t_3 . I believe this is often overlooked since, if Jones had considered not killing Smith at t_3 , Black would have intervened and forced Jones to decide to kill Smith at some later time. But this does not diminish the importance of this point: at the moment of decision, at t_3 , Jones was able to do otherwise and thus avoid responsibility for deciding to kill Smith at t_3 . Jones then satisfies clause (i) of PAP₁. This point will become important below.

But remember that neo-Frankfurtians dispute the relevance, not the truth of PAP₁ and so they need not object to my claim that Jones satisfies (i) of PAP₁. What they must insist on is that this ability does not even partly explain why Jones is responsible for deciding to kill Smith. The rest of my paper will be concerned with rebutting this claim. My strategy will be to show that Jones’s ability to consider not killing Smith satisfies Hunt and Pereboom’s epistemic condition for relevance. As this is the *only* reason Hunt and Pereboom offer for thinking that the remaining ability to avoid responsibility in buffer cases cannot explain Jones’s moral responsibility, by showing that this ability *does* satisfy the epistemic condition I will thereby show that buffer cases pose no threat to the relevance of PAP₁.

So what reason is there to think that Jones’s considering not killing Smith satisfies the epistemic condition? Once we realize that Jones was able, at the moment of decision, to consider not killing Smith it becomes clear—at least when we insert the relevant time indices—that his ability satisfies the epistemic condition. The first thing to notice is that if Jones considered not killing Smith at t_3 , he would *thereby* have avoided responsibility for deciding to kill Smith at t_3 and this has nothing to do with the presence of Black. Deliberating at t_3 about whether to kill Smith is inconsistent with deciding to kill Smith at t_3 . Consequently, if Jones exercised his ability to consider not killing Smith at t_3 , he would not have decided to kill Smith at t_3 and therefore would have avoided responsibility for deciding to kill Smith at t_3 . After all, one cannot be morally responsible for something that never happened. Moreover, Jones understands all this. He understands that deliberating about whether to φ at time t is inconsistent with deciding to φ at t . Presumably Jones understands that if he considered not killing Smith at t_3 , then he would thereby avoid responsibility for deciding to kill Smith at t_3 . Had we asked Jones later if it was possible for him to be responsible for deciding to kill Smith at t_3 even though he was considering not killing Smith at t_3 he would say ‘No’. Upon further prompting he might go on to explain that his considering not killing Smith at t_3 is inconsistent with his deciding to kill Smith at t_3 . Since, in such a situation, he could not and would not have decided to kill Smith at t_3 , then he could not and would not be

responsible for deciding to kill Smith at t_3 .¹⁹ “One cannot be responsible for what doesn’t happen,” he might add. Therefore, Jones’s ability to avoid responsibility—i.e. his ability to consider not killing Smith—intuitively satisfies Hunt and Pereboom’s epistemic condition for relevance: his exercising this ability is sufficient to avoid responsibility and he understands this.

Hunt and Pereboom fail to distinguish two different senses in which Jones’s considering not killing Smith is sufficient for his avoiding responsibility for deciding to kill Smith. The first sense of avoidance is that Jones’s considering not killing Smith at t_3 is sufficient for his avoiding responsibility for deciding to kill Smith at the time of action, namely t_3 , and at any later time. The second sense of avoidance is that Jones’s considering not killing Smith at t_3 is sufficient for his avoiding responsibility for deciding to kill Smith at the time of action. When avoidance is understood in the first sense, it is true that Jones failed to understand that by considering not killing Smith at t_3 he would thereby avoid responsibility for deciding to kill Smith at t_3 and at any later time. Jones’s considering not killing Smith at t_3 is sufficient for his avoiding moral responsibility for deciding to kill Smith at t_3 and at any later time *only* because of Black’s presence and powers, and Jones is unaware of these facts. By exclusively focusing on the first sense of avoidance, Hunt and Pereboom’s claim that Jones’s ability to avoid responsibility is not relevant to explaining responsibility gains plausibility. But once we notice the availability of a second sense of avoidance, we are able to see that this ability does indeed satisfy the epistemic condition. Jones’s considering not killing Smith at t_3 is sufficient for avoiding moral responsibility for deciding to kill Smith at t_3 and this connection holds independent of Black’s presence. Moreover, Jones understood that his considering not killing Smith at t_3 precluded the possibility of deciding to kill Smith at t_3 . Therefore, Jones’s ability to consider not killing Smith, his ability to avoid responsibility, does indeed satisfy the epistemic condition for relevance.

Buffer cases therefore fail to show that PAP₁ is explanatorily irrelevant to moral responsibility. These cases are designed to constitute a counterexample to the explanatory relevance of the ability to avoid responsibility by showing that an agent can be responsible even though he does not *understand* that any of his abilities are such that by exercising one of them he will thereby avoid responsibility. *Revenge* seemed to fit the bill for neo-Frankfurtians. At first glance it appeared that Jones was responsible for deciding to kill Smith at t_3 even though none of his abilities were such that Jones understood that by exercising it he would thereby avoid responsibility. But appearances can be deceiving. I argued that Jones’s ability to consider not killing Smith at t_3 was just this kind of ability: Jones understood that by exercising it he would thereby avoid responsibility for deciding to kill Smith at t_3 . I also appealed to the ambiguity inherent in the term ‘avoidance’ to explain why Hunt and Pereboom failed to consider this response. There are two senses in which an agent can avoid responsibility. The first, and stronger, sense of avoidance is that

¹⁹ If Jones lacks this rather minimal level of understanding about the nature of action, deliberation and responsibility, then one might doubt whether he meets the minimal level of competence to be responsible in the first place.

an agent's ability to φ at t is sufficient for his avoiding responsibility for φ at t and at any later time. The second, and weaker, sense of avoidance is that an agent's ability to φ at t is sufficient for his avoiding responsibility for φ at t . It is by exclusively focusing on the first and stronger sense of avoidance that *Revenge* appears to be a counterexample to the explanatory relevance of PAP₁. However, once we distinguish these senses of avoidance, we realize that *Revenge* is no such thing: it gives us no reason to doubt that Jones's ability to consider not killing Smith at t_3 , his ability to avoid responsibility, is relevant to explaining his responsibility for deciding to kill Smith at t_3 .

4 Objections and replies

4.1 Objection 1

A possible objection that arises at this juncture maintains that I am mistaken in assuming that the second sense of avoidance is the relevant sense in the present context. Rather, according to this objection, it is the first sense that is relevant and since I admit that Jones does not understand that his ability to consider not killing Smith is sufficient to avoid responsibility for deciding to kill Smith in this sense of avoidance, it follows that this ability fails to satisfy the epistemic condition for relevance.

4.2 Reply to objection 1

But even if this objection is sound, it raises no problem for my contention that buffer cases fail to show that PAP₁ is explanatorily relevant. Hunt and Pereboom need there to be a single sense of avoidance under which the epistemic condition is true *and* Jones's ability to consider not killing Smith at t_3 does not satisfy it. But there is no such sense of avoidance: Hunt and Pereboom face a dilemma. Either they adopt the first or second sense of avoidance. If they adopt the first sense, then Jones's ability to consider not killing Smith at t_3 does indeed violate the epistemic condition, but this turns out to be unproblematic since the epistemic condition under this reading of 'avoidance' is false. If they adopt the second sense of avoidance, then the epistemic condition is true but Jones's ability to consider not killing Smith at t_3 satisfies it.

I have defended the second horn above. There I argued that Jones understood that his ability to consider not killing Smith at t_3 is such that by exercising it Jones would have avoided responsibility for deciding to kill Smith at t_3 . Let me now defend the first horn. According to this horn, Hunt and Pereboom can insist that the first sense of avoidance is the relevant sense only at the expense of undermining any plausibility that their epistemic condition seemed to have. This sense of avoidance is too strong as can be seen by the fact that it becomes absurdly easy to show that the PAP₁ is not explanatorily relevant: one can do so without even having to appeal to FSCs. Rarely if ever do we possess an ability such that we understand that by exercising that ability we will thereby avoid responsibility at the time of action and

at any later time. We are often in situations in which we have the opportunity to exercise an ability that is sufficient for avoiding responsibility for an action at the time of action, but not for avoiding responsibility for that action at any later time. Consider someone who in a burst of anger is tempted to make a demeaning remark to his friend and acts on this temptation. In normal circumstances we would think this person possessed the ability to refrain from making this remark and thus is able to avoid responsibility for making the remark. But is his ability to refrain from making the demeaning remark relevant to explaining his moral responsibility for making this remark? Clearly not, at least so long as we assume that the first sense of avoidance is the relevant sense. Even if this person had exercised his ability to refrain from making the demeaning remark, it is still possible at some later time that he will again become angry but this time succumb to the temptation and make the demeaning remark. Consequently, his ability to refrain from making the remark is not such that if he exercised it he would thereby avoid responsibility for making the demeaning remark at the time of action and at *any* later time. And, by hypothesis, PAP₁ is not explanatorily relevant since none of the agent's abilities satisfy the epistemic condition.

This result suggests that the first sense of avoidance is too strong. Abilities such as the friend's ability to refrain from making the demeaning comment may indeed fail to be explanatorily relevant, but not because they fail to satisfy the epistemic condition under the first reading of avoidance. It is too stringent to require for explanatory relevance that exercising the ability results in one's avoiding responsibility at the time of action and at any later time. Instead, it is sufficient that exercising the ability allows one to avoid responsibility at the present time—what happens later is immaterial. Therefore, buffer cases face serious problems regardless of the sense of avoidance that we adopt.

4.3 Objection 2

According to this objection, my argument above does indeed show that *Revenge* fails to cast the explanatory relevance of PAP₁ into doubt, but maintains that there is a way to modify *Revenge* in order to avoid my above argument. So, this objection concedes that my argument succeeds against a particular instance of a buffer case, but denies that it shows that buffer cases in general fail to show that PAP₁ is not explanatorily relevant. Hunt and Pereboom can modify *Revenge* by placing a temporal gap between when Jones is able to consider not killing Smith and his actual decision to kill Smith. Hunt and Pereboom might stipulate that, necessarily, if Jones fails to consider not killing Smith by t_1 , then he will decide to kill Smith at t_3 . Not only is Jones's considering not killing Smith a necessary condition for not deciding to kill Smith at t_3 , but Jones's considering not killing Smith *by* t_1 is a necessary condition for not deciding to kill Smith at t_3 . But, just as in *Revenge*, if Jones does consider not killing Smith by t_1 , then Black will intervene at t_2 and force Jones to decide to kill Smith at t_3 . Hence, at t_3 it is impossible for Jones to do other than decide to kill Smith. The only way he can try to do otherwise is if he, by t_1 , considers not killing Smith and yet if he does this Black will intervene and force Jones to decide to kill Smith at t_3 . Everything else proceeds just as it does in

Revenge. Jones never considers not killing Smith and hence at t_3 is unable to avoid responsibility for deciding to kill Smith. Nonetheless we might still judge that Jones is responsible for making this decision. After all he chose to do it for his own reasons, he was not coerced, and so on. Let us refer to this case as *Revenge*_{TG}.²⁰

4.4 Reply to objection 2

Although Hunt and Pereboom do not actually offer this version of a buffer case, it strikes me as a very plausible way to refashion this kind of case in order to avoid my above objection. Consequently, if we are to have any confidence that buffer cases do not pose a threat to the explanatory relevance of PAP₁, it must be shown that *Revenge*_{TG}, like *Revenge*, fails. I turn now to developing an argument that shows just this.

Notice first that in *Revenge*_{TG}, unlike in *Revenge*, Jones clearly does not satisfy (i) of PAP₁: there was nothing Jones was able at t_3 to do to avoid responsibility for deciding to kill Smith at t_3 . However, Jones does satisfy (ii) of PAP₁. As (ii) of PAP₁ makes clear, an agent can be responsible for φ -ing at t even though he was unable at t to avoid responsibility for φ -ing just so long as there is some earlier action or omission ψ that explains why he is unable at t to avoid responsibility for φ -ing and the agent, at this earlier time, was able to avoid responsibility for ψ -ing. In *Revenge*_{TG}, Jones lacks the ability at t_3 to avoid responsibility for deciding to kill Smith, but this is explained by his omitting to consider not killing Smith at t_1 and Jones was able at t_1 to avoid responsibility for failing to consider not killing Smith at t_1 —just at (ii) requires.

So Jones satisfies PAP₁, but as before this is unproblematic for Hunt and Pereboom since their claim is that buffer cases pose a threat to the relevance, not the truth of PAP₁. As before, their argument that Jones's satisfying this clause cannot possibly explain why he is responsible for deciding to kill Smith at t_3 is that he fails to satisfy the epistemic condition for relevance: Jones does not understand that by considering not killing Smith by t_1 he will thereby avoid responsibility for deciding to kill Smith at t_3 .²¹

I believe, however, that even this more elaborate case fails, but for very different reasons. Unlike in *Revenge*, Jones's ability to consider not killing Smith does not satisfy the epistemic condition with respect to his decision to kill Smith in *Revenge*_{TG}. It is true that if Jones considered not killing Smith at t_1 , then he would thereby avoid responsibility for deciding to kill Smith at t_3 . By deliberating about what to do at t_1 , Jones forces Black's hand and thus excuses himself from responsibility. But Jones does not understand any of this. It is hard to imagine that a rational agent could fail to understand that deliberating about whether to φ at t precludes the possibility of φ -ing at that exact same moment. But it is not

²⁰ 'TG' stands for 'temporal gap'.

²¹ Of course Hunt and Pereboom do not maintain this in the sense that these claims never appear in their work. I will continue to speak of 'Hunt and Pereboom maintaining this' and 'Hunt and Pereboom arguing for that' but it should be kept firmly in mind that I am merely speculating here on what they might say. These speculations notwithstanding, I do believe that *Revenge*_{TG} poses an interesting threat to PAP₁ and so it is worth showing why it also fails.

implausible to think that Jones was unaware that his considering not killing Smith at t_1 was inconsistent with his deciding to kill Smith at t_3 . The connection between considering not killing Smith at t_1 and avoiding responsibility for deciding to kill Smith at t_3 obtains only in virtue of Black's presence—of which Jones is quite unaware. Hence, it does seem that Jones fails to understand that by considering not killing Smith at t_1 he will thereby avoid responsibility for deciding to kill Smith at t_3 .

But although I concede that Jones's ability to consider not killing Smith at t_1 does not satisfy the epistemic condition with respect to his decision to kill Smith at t_3 , I deny that he must satisfy it in order for this ability to be relevant to explaining his responsibility. In this case Hunt and Pereboom's condition is too strong. Interestingly, the epistemic condition for the relevance of an ability appears to differ depending on whether the ability satisfies clauses (i) or (ii) or PAP₁. When an ability that satisfies (i) partly explains an agent's moral responsibility for φ -ing, it must satisfy Hunt and Pereboom's epistemic condition for relevance. However, when an ability that satisfies (ii) partly explains an agent's moral responsibility for φ -ing, then this epistemic condition need not be satisfied. This is a difference that is worth trying to understand in its own right, but in the present context it will suffice to show that there is such a difference.²²

Part of the motivation for including clause (ii) in PAP₁ is the realization that we play an important role in shaping who we become and the kind of circumstances we are in.²³ If character and circumstances render us unable to avoid responsibility, then it might seem that those features of the world assail us in a way that render us excused from responsibility. However, if we brought about our character and/or circumstances and we were able to avoid responsibility for doing so, then it is plausible to think we are responsible for actions that issue forth from our character and circumstances. Luther's 'Here I stand I can do no other' fits the latter mold. Visiting Rome and seeing the abuses, composing and posting the 95 theses, combating his various theological opponents, publically burning the papal bull and so on all served to mold him into the kind of person who would not, indeed could not recant and renounce his disagreements with the Roman Catholic Church.

According to (ii) of PAP₁, if the ability to consider not killing Smith at t_1 is relevant to explaining Jones's responsibility for deciding to kill Smith at t_3 at all, then it is only *indirectly* relevant. There are two steps here. The first is to show that Jones's responsibility for deciding to kill Smith at t_3 is grounded in or traces back to his responsibility for omitting to consider not killing Smith at t_1 . The core idea behind "tracing" is that an agent can be responsible for φ -ing at t even though he does not satisfy certain conditions for moral responsibility at t —e.g. cannot do otherwise—just so long as his failure to satisfy these conditions at t is grounded in or explained by an earlier action(s) for which he is responsible. An often cited example of where tracing is appropriate is the case of a drunk driver running over a pedestrian. In such a case it is reasonable to suppose that the drunk driver is

²² I am thankful to Paul Hoffman for pointing out the possibility of this asymmetry.

²³ For recent discussions of this important theme see Kane (1996) and Ekstrom (2003).

responsible for killing the pedestrian even though he lacked the needed control for moral responsibility at the time of the killing. It is reasonable to suppose this because he is responsible for getting drunk and his responsibility for killing the pedestrian traces back to his responsibility for getting drunk.

The second step is to show that Jones's being able to consider not killing Smith at t_1 partly explains why he is responsible for omitting to consider not killing Smith at t_1 . It is in this sense that Jones's ability to consider not killing Smith is indirectly relevant to explaining his responsibility for deciding to kill Smith at t_3 : it partly explains his responsibility for omitting to consider not killing Smith at t_1 which in turn partly explains his responsibility for deciding to kill Smith at t_3 .²⁴ In cases that involve tracing, an ability can be relevant to explaining the agent's responsibility for φ -ing even if it fails to be such that if he exercised this ability he would thereby avoid responsibility. Such an ability can still be relevant if it satisfies the epistemic condition with respect an earlier action for which he is responsible and his responsibility for this earlier action grounds his responsibility for φ -ing.

There are then two steps I need to establish. The first step consists in showing that Jones's responsibility for deciding to kill Smith traces back to his omitting to consider not killing Smith. The second step consists in showing that his ability to consider not killing Smith partly explains his responsibility for omitting to consider not killing Smith. If both of these steps can be established, then it will be clear that Jones's ability to consider not killing Smith by t_1 , even though it does not satisfy the epistemic condition with respect to his decision to kill Smith at t_3 , is nonetheless relevant to explaining his moral responsibility for deciding to kill Smith at t_3 .

My task here is not to show that each one of these steps is undoubtedly true, but that buffer cases provide no reason for doubting their truth. Consider the second step first. Does Jones's ability to consider not killing Smith at t_1 partly explain why he is responsible for omitting to consider not killing Smith at t_1 ? It would seem so and buffer cases give no reason to the contrary. There is *prima facie* reason to think that the ability to avoid responsibility partly explains an agent's moral responsibility. We often cite just such a factor when explaining an agent's responsibility. Moreover, buffer cases fail to call this ability's relevance into doubt since this ability satisfies the epistemic condition with respect to Jones's omitting to consider not killing Smith at t_1 . Jones understood that by considering not killing Smith at t_1 , he would thereby avoid responsibility for omitting to consider not killing Smith at t_1 . Jones understands that his considering not killing Smith at t_1 would make it

²⁴ Nothing here turns on the nature of omissions. Some have thought that only non-occurrences of which an agent considered count as omissions. Under this definition of omissions Jones's not considering not killing Smith would not be an omission, but only a mere non-occurrence since Jones does not consider this action. I am inclined to a more liberal understanding of omissions, but regardless of one's view of omissions, we can trace responsibility back to a non-occurrence even when the agent has not considered this option as cases of negligence show. Suppose I negligently forget to set my alarm clock and as a result am late for a very important meeting the next morning. My boss might reasonably hold me responsible for being late precisely because I forgot to set my alarm and so proved myself to be irresponsible. However, in this case I did not deliberate about whether to set the alarm. The problem is that setting the alarm never occurred to me. Fischer and Ravizza (1998, pp. 124–125) make a similar point. Whether or not mere non-occurrences count as omissions we can trace responsibility back to mere non-occurrence and so I need not take a stand on the nature of omissions.

impossible for him simultaneously to omit to consider not killing Smith at t_1 . Moreover, he understands that if he considered not killing Smith at t_1 , then he cannot be responsible for omitting to consider not killing Smith at t_1 : how can one be responsible for something that did not happen? Thus, Jones's ability to consider not killing Smith does satisfy the epistemic condition with respect to omitting to consider not killing Smith. And as this ability's failing to satisfying the epistemic condition is the only reason that Hunt and Pereboom offer for doubting its relevance, we have reason to think that Jones's ability at t_1 to consider not killing Smith is explanatorily relevant for his responsibility for omitting to consider not killing Smith at t_1 .²⁵ Consequently, if we can establish the first step, PAP₁ is vindicated against the neo-Frankfurtian challenge.

Establishing this step is less straightforward, but again I think buffer cases give no reason to doubt its truth. One might explain that the reason Jones's moral responsibility for deciding to kill Smith traces back to his responsibility for omitting to consider not killing Smith is that this omission is *tantamount* to deciding to kill Smith. Jones's omitting to consider not killing Smith by t_1 is tantamount to deciding to kill Smith at t_3 because this omission, together with certain uncontested necessary conditions for Jones's deciding to kill Smith at t_3 , like Jones's not dying before he gets a chance to make this decision, is *sufficient* for Jones's deciding to kill at t_3 .²⁶ Even though Jones was unable to decide not to kill Smith at t_3 , he is still responsible for deciding to kill Smith because he is responsible for an earlier omission and this earlier omission is tantamount to his deciding to kill Smith at t_3 .

Supposing that Jones's omission is tantamount to his deciding to kill Smith, does this make it reasonable to think that Jones's responsibility traces back? Hunt and Pereboom might argue that it does not by appealing to a *new* epistemic condition: an agent S's responsibility for φ -ing can trace back to his responsibility for ψ -ing only if were S to avoid responsibility for φ -ing he would thereby avoid responsibility for ψ -ing. This is an epistemic condition not for when an ability is explanatorily relevant to responsibility, but for when moral responsibility for an action/omission can trace back to earlier actions/omissions. Jones does not understand that by considering not killing Smith at t_1 that he will thereby avoid responsibility for deciding to kill Smith at t_3 since he is ignorant of Black's presence. Therefore, his responsibility for the decision cannot trace back to his responsibility for the omission.

Whatever plausibility Hunt and Pereboom's epistemic condition for the explanatory relevance of an ability has, it is lacking for this epistemic condition concerning tracing. Return to the drunk driver case. Does the drunk driver understand that by not getting drunk he will thereby avoid responsibility for killing the pedestrian? Certainly not since his failing to get drunk is not sufficient for his avoiding responsibility for killing the pedestrian. For example, suppose after finally succeeding in resisting his temptation to have another drink, he drives towards home

²⁵ Although I have suggested both that we have reason to think that Jones's ability at t_1 to consider not killing Smith is explanatorily relevant *and* that buffer cases give us no reason to doubt this, strictly speaking I only need to establish the latter claim since my goal is to defend PAP₁ against buffer cases and not to show directly that PAP₁ is explanatorily relevant.

²⁶ I am grateful to Derk Pereboom for providing helpful suggestions on how to make the tantamount relation more precise.

completely sober. Unfortunately he is incredibly tired and falls asleep at the wheel, veers off the road and runs over the same poor pedestrian. It is reasonable to judge that he is responsible for killing the pedestrian because he negligently drove his car when he knew he shouldn't have. Consequently, failing to get drunk is not sufficient for avoiding responsibility for running over the pedestrian, but it is nonetheless something to which we can trace the drunk driver's responsibility. Hunt and Pereboom's new epistemic condition for when the agent's earlier responsibility for an action/omission can ground an agent's later responsibility for an action/omission is too strong.

This becomes especially clear when we acknowledge the relevant time indices. Does the drunk driver understand that by not getting drunk at t_1 he will thereby avoid responsibility for running over the pedestrian at t_3 ? It seems not since he might get drunk at t_2 , start driving at $t_{2.5}$ and run over the pedestrian at t_3 . But again this possibility does not show that we cannot trace the drunk driver's responsibility back to his responsibility for getting drunk at t_1 . So Jones's failing to satisfying this new epistemic condition for tracing presents no obstacle to his responsibility for the decision tracing back to his responsibility for his omitting to consider not killing Smith.

Now, exactly what epistemic condition an agent must satisfy in order for it to be fair to trace his responsibility back to his responsibility for earlier actions or omissions is a controversial and difficult matter. But even without knowing precisely what the epistemic condition is, I believe that we still can conclude that buffer cases pose no threat to tracing Jones's responsibility back to this omission. First, note that it seems that the cases of Jones and the drunk driver are relevantly similar so that if one case involves tracing it is likely that the other case does too. And since the case of the drunk driver is a paradigm case of tracing, Jones's case also seems to fit the mold. Second, often the notion of 'could have reasonably foreseen' is invoked: the drunker driver's responsibility traces back only if he could have reasonably foreseen that by getting drunk he might cause harm to someone.²⁷ Likewise, we might say that Jones could have reasonably foreseen that not deliberating about whether to kill Smith might lead to his killing Smith. Since Jones could have reasonably foreseen this consequence, he is responsible for deciding to kill Smith because he failed to consider not killing Smith.

So it seems reasonable to trace Jones's responsibility for deciding to kill Smith back to his responsibility for omitting to consider not killing Smith. This omission is tantamount to deciding to kill Smith and although Jones may not have understood all the details, he could have reasonably foreseen that omitting to consider not killing Smith might lead to his deciding to kill Smith. Moreover, Jones's responsibility for omitting to consider not killing Smith is itself partly explained by Jones's ability to consider not kill Smith. I conclude then that buffer cases do not throw doubt on either Jones's satisfying PAP₁ or this fact's partly explaining why Jones is responsible for deciding to kill Smith.

²⁷ See Fischer and Tognazzini (2009). In this paper they provide a range of possible epistemic conditions that govern tracing cases in an attempt to respond to some skeptical challenges raised by Vargas (2005). The interesting thing to note is that Jones satisfies them all.

5 Conclusion

Buffer cases thus fail to pose any threat to PAP₁. First, they clearly fail as counterexamples to the truth of PAP₁: in both *Revenge* and *Revenge*_{TG} Jones satisfies PAP₁. Moreover, in both of these cases it is reasonable to cite Jones's ability to consider not killing Smith in an explanation of his moral responsibility for deciding to kill Smith, or at least buffer cases give us no reason to doubt its relevance. With regard to *Revenge* we saw that Hunt and Pereboom faced a dilemma: either they adopt the strong reading of avoidance and so their epistemic condition is false, or they adopt the weak sense of avoidance and Jones satisfies their epistemic condition. But either way there is no single sense of avoidance under which the epistemic condition is true and Jones fails to satisfy it. In *Revenge*_{TG} things were a bit more complex but none the more problematic. We saw that Jones's ability to consider not killing Smith was indirectly relevant to explaining his responsibility for deciding to kill Smith. Jones's ability to consider not killing Smith satisfied Hunt and Pereboom's epistemic condition with respect to avoiding responsibility for omitting to consider not killing Smith and hence this ability is directly relevant to explaining his responsibility for omitting to consider not killing Smith. Moreover, we saw that Jones's responsibility for deciding to kill Smith traced back to his responsibility for omitting to consider not killing Smith and hence Jones's ability to consider not killing Smith was indirectly relevant to explaining his responsibility for deciding to kill Smith. That buffer cases fail to pose any threat to the relevance of PAP₁ is an important conclusion for proponents of PAP₁ and similar principles. Although I have not shown, and have not tried to show, that PAP₁ is true and relevant to explaining moral responsibility, I have cleared the ground from one of the most threatening obstacles to the truth and relevance of this principle. The door is thus wide open, or at least more widely open, for proponents of PAP₁ to run their usual arguments for the truth and relevance of this principle.

Acknowledgments Earlier versions of this paper were presented at the 2007 Society of Christian Philosophers, Mountain-Pacific Region Meeting and University of California, Riverside Agency Workshop. I am grateful to the audiences for their helpful discussion, especially Justin Coates, John Fischer, Paul Hoffman, David Hunt, Michael Nelson, Garret Pendergraft, Ben Mitchell-Yellin, Phillip Swenson, Patrick Todd, Neal Tognazzini, and Gary Watson. I would like to thank David Hunt and Derk Pereboom for their willingness to read and provide very helpful feedback on earlier drafts of this paper as well as an anonymous referee for this journal. I am especially indebted to John Fischer who provided invaluable feedback on this paper through its many variations.

Open Access This article is distributed under the terms of the Creative Commons Attribution Non-commercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Baker, L. R. (2006). Moral responsibility without libertarianism. *Nous*, 40, 307–330.
- Copp, D. (2003). 'Ought' implies 'can', blameworthiness, and the principle of alternative possibilities. In D. Widerker & M. McKenna (Eds.), *Moral responsibility and alternative possibilities: Essays on the importance of alternative possibilities*. Burlington, VT: Ashgate Publishing Co.

- Dennet, D. (1984). *Elbow room*. Cambridge: MIT Press.
- Ekstrom, L. W. (2003). Free will, chance, and mystery. *Philosophical Studies*, 113, 153–180.
- Fischer, J. M. (1994). *The metaphysics of free will: An essay on control*. Malden: Blackwell Publishers.
- Fischer, J. M. (2006). Frankfurt-style compatibilism. In *My way: Essays on moral responsibility*. New York: Oxford University Press.
- Fischer, J. M., & Ravizza, M. (1998). *Responsibility and control: A theory of moral responsibility*. New York: Cambridge University Press.
- Fischer, J. M., & Tognazzini, N. (2009). The truth about tracing. *Nous*, 43, 531–556.
- Frankfurt, H. (1969). Alternate possibilities and moral responsibility. (Reprinted from *The importance of what we care about*, 1988, New York: Cambridge University Press.)
- Glover, J. (1970). *Responsibility*. New York: Humanities Press Inc.
- Hunt, D. (2005). Moral responsibility and buffered alternatives. *Midwest Studies in Philosophy*, 29, 126–145.
- Kane, R. (1996). *The significance of free will*. New York: Oxford University Press.
- Leon, F., & Tognazzini, N. (Forthcoming). Why Frankfurt examples don't need to succeed to succeed. *Philosophy and Phenomenological Research*.
- McKenna, M. (1997). Alternative possibilities and the failure of the counterexample strategy. *Journal of Social Philosophy*, 28, 71–85.
- Mele, A., & Robb, D. (1998). Rescuing Frankfurt-style scenarios. *Philosophical Review*, 107, 97–112.
- Mele, A., & Robb, D. (2003). Bbs, magnets and seesaws: The metaphysics of Frankfurt-style cases. In D. Widerker & M. McKenna (Eds.), *Moral responsibility and alternative possibilities: Essays on the importance of alternative possibilities*. Burlington, VT: Ashgate Publishing.
- Otsuka, M. (1998). Incompatibilism and the avoidability of blame. *Ethics*, 108, 685–701.
- Pereboom, D. (2001). *Living without free will*. New York: Cambridge University Press.
- Vargas, M. (2005). The trouble with tracing. *Midwest Studies in Philosophy*, 29, 269–291.
- Widerker, D. (1995). Libertarianism and Frankfurt's attack on the principle of alternative possibilities. *Philosophical Review*, 104, 247–261.
- Widerker, D., & McKenna, M. (2003). *Moral responsibility and alternative possibilities: Essays on the importance of alternative possibilities*. Burlington, VT: Ashgate Publishing.
- Wyma, K. (1997). Moral responsibility and leeway for action. *American Philosophical Quarterly*, 34, 57–70.
- Yaffe, G. (2005). More on 'ought' implies 'can' and the principles of alternative possibilities. *Midwest Studies in Philosophy*, 29, 307–312.