

Blameworthiness and Buffered Alternatives

Note: This is the penultimate draft of an article the final and definitive version of which is forthcoming in *American Philosophical Quarterly*. The journal is available online here: <http://www.press.uillinois.edu/journals/apq.html>

1. Introduction

According to one version of the principle of alternative possibilities (PAP, for short), a person is blameworthy for what she did only if there was a *robust* alternative course of action available to her, where a robust alternative is one that is relevant per se to an explanation of why the agent is blameworthy for what she did.¹ This principle used to be widely accepted. It was seen as a plausible way of cashing out the popular idea that free will is required for blameworthiness, and was thought to provide a principled explanation of the exculpatory force of pleas like “I couldn’t help it!” and “I had no choice.” Nowadays, however, it is a matter of considerable controversy.²

In his influential 1969 article “Alternate Possibilities and Moral Responsibility,” Harry Frankfurt discussed a series of examples culminating in one that he contends is a counterexample to PAP.³ It is widely agreed that this example lacks sufficient detail to shoulder the argumentative burden placed on it. But convinced Frankfurt was on to something, a number of philosophers subsequently constructed more elaborate “Frankfurt-style examples” in an effort to produce a case that can bear the weight of the argument. The current debate over PAP is focused largely on whether any of these newer cases succeed where Frankfurt’s original example failed.

That debate continues here. I will argue that the buffer cases, which are among the most influential Frankfurt-style examples devised in recent years, are not counterexamples to PAP. As we shall see, my argument for this conclusion bears some resemblance to another criticism of the

buffer cases that has been receiving attention lately. However, I contend that whatever merit this other criticism may have, the one on offer here has important advantages over it.

2. The Buffer Cases

Consider the following Frankfurt-style case. Marla is tempted to cheat on her philosophy midterm. Given her current psychology, it is impossible for her to make any decision other than a decision to cheat on the exam, unless she first seriously attends to the reasons for not cheating on it. Only then will she be in a motivational position to make an alternative decision. In the end, Marla succumbs to temptation and decides to cheat on the exam, despite knowing this to be the morally wrong decision, and without ever seriously reflecting on the reasons not to cheat, though it was within her power at the time to do so. Now, unbeknownst to Marla, a nefarious neuroscientist implanted a device in her brain that enables him to detect her thoughts and control her behavior. Had he detected that Marla was seriously reflecting on the reasons for not cheating, he would have used this device to halt Marla's reflections and to deterministically cause her to decide to cheat, and there is nothing she could have done to stop him from doing this.

All Frankfurt-style examples have the same basic structure: an agent does something on her own that, unbeknownst to her and through no fault of her own, she would have been forced to do had she not done it on her own. What is unique about Frankfurt cases like the one presently at issue is the means by which it is determined whether the neuroscientist and his device come into play. Recall that, given Marla's psychological state, for her to decide to do something other than what she actually decided to do, she would first have had to seriously reflect on the reasons for not cheating on her midterm. Pausing to consider those reasons would not have automatically resulted in her making a different decision, even in an ordinary version of the story in which the

neuroscientist and his device are absent. But reflection on those reasons is a necessary condition for her deciding otherwise. It constitutes a mental buffer of sorts between the decision Marla actually made and any alternative decisions. In order to make an alternative decision, Marla would first have to enter (so to speak) that mental buffer zone. But if (and only if) she were to enter it by seriously contemplating the reasons not to cheat, the neuroscientist would intervene and deterministically cause her to decide to cheat before she could make an alternative decision. Accordingly, Frankfurt-style examples like this have come to be known as buffer cases.⁴

A virtue of the buffer cases is that they seem to avoid a dilemma thought to plague earlier Frankfurt-style examples. Very roughly, the dilemma is this: either the agent in the example was deterministically caused to act as she did, or she was not. If she was deterministically caused to act as she did, then claiming that she is blameworthy for acting in that way would, in the absence of further argument, beg the question against those who insist that causal determinism is incompatible with blameworthiness. If, on the other hand, the agent's behavior was not deterministically caused, then it is alleged that the example will not be one in which the agent lacks the relevant power to do otherwise. Either way, the example will fail to provide us with a case in which a person is clearly blameworthy despite lacking alternative possibilities for action.⁵

The buffer cases are thought to circumvent this dilemma by providing us with genuinely indeterministic Frankfurt-style examples in which the agent lacks any robust alternatives but is blameworthy for what she did nonetheless. Nothing about the examples requires us to assume that the agent's behavior was deterministically caused. So let us assume that it was not. Let us assume in particular that prior to t , where t is the moment at which Marla decided to cheat, it was undetermined whether Marla would decide at t to cheat on her midterm or whether she would instead reflect at t on the reasons for not cheating. Given this assumption, the claim that Marla is

blameworthy for her decision to cheat on the exam does not beg the question against incompatibilists. At the same time, though, the assumption does nothing to change the fact that it was not within Marla's power to avoid deciding to cheat on the exam, for if she had not decided on her own at t to cheat, she would have been forced to decide to cheat a few seconds later by the neuroscientist and his device. It seems, then, that Marla could not have avoided deciding to cheat, even though her actual decision to cheat was not deterministically caused. But because Marla decided on her own to cheat, without any "help" from the neuroscientist, it also seems that she could be blameworthy for deciding as she did, her lack of alternatives notwithstanding.

It is important to note that Marla did have one alternative course of action available to her at the time. Rather than deciding at t to cheat on her midterm, she could have instead voluntarily paused at t to seriously contemplate the reasons for not cheating on the exam. However, according to proponents of the buffer cases, it is not even partly in virtue of the fact that Marla could have performed this alternative action at t that she is blameworthy for what she actually did at t . In other words, they claim that pausing at t to reflect on the reasons for not cheating is not a *robust* alternative, that is, it is not relevant per se to an explanation of why Marla is blameworthy for her decision, and thus is not the sort of alternative action called for by PAP.⁶

One of their main arguments for this claim is due to Derk Pereboom.⁷ He contends that an alternative action is robust only if the agent was cognitively sensitive to the fact that if she voluntarily performed that action she would then likely be blameless for the action she actually performed (2014, p. 13).⁸ He then argues that the alternative action available to an agent in Marla's position does not satisfy this requirement. If Marla had paused at t to consider the reasons for not cheating, the neuroscientist would have used his coercive mechanism to force her to decide a few seconds later to cheat, in which case she presumably would not have been

blameworthy for her decision. So, if Marla had paused at t to contemplate the reasons for not cheating, she would have avoided blameworthiness for deciding to cheat on her test. The difficulty, however, is that Marla was oblivious to all this. She had no idea that pausing at t to contemplate the reasons for not cheating would render her blameless for deciding to cheat. For all she knew, she could pause briefly at t to reflect on the reasons not to cheat, freely decide a moment later to cheat, and still be blameworthy for her decision, which of course would be true had the neuroscientist not been lurking in the background. Marla was therefore not cognitively sensitive to the fact that if she considered the reasons for not cheating she would then be blameless for deciding to cheat. Hence, Pereboom and other defenders of the buffer cases conclude that pausing at t to reflect on the reasons against cheating is not a robust alternative.

I will evaluate this argument in due course. The important thing to note for now is that if it is sound, if the only alternative course of action available to Marla in this story is not robust, then it would appear that we have a counterexample to PAP, a case in which someone is blameworthy for what she did even though there was no robust alternative available to her.

3. A Robust Alternative

The buffer cases are among the more promising Frankfurt-style examples, to be sure. But are they really counterexamples to PAP? As I mentioned earlier, I will argue that they are not. More precisely, I will argue that while the agent in buffer cases may be blameworthy for what she did, the alternative action that was available to her is relevant per se to an explanation of why she is blameworthy for her behavior. The buffer cases therefore do not provide us with examples in which someone is blameworthy for an action despite lacking a robust alternative, and thus they

pose no threat to PAP. My discussion will focus on the story about Marla, but since that story is representative of extant buffer cases, what I have to say about it applies to those other cases too.

Let us remind ourselves of some pertinent details. Marla knew that cheating on her philosophy midterm was immoral, but she decided at t to cheat on it anyway. Now, this was not something she *had* to do at t . She could have easily paused at t to reflect on the reasons for not cheating instead. It was, moreover, morally permissible for Marla to perform this alternative action at t , and we may assume that she was cognitively sensitive to that fact as well. Indeed, it is plausible that seriously reflecting on the reasons for not cheating is exactly what Marla should have done in this situation. The upshot of all this is that Marla deliberately did something at t which she knew to be immoral, despite being cognitively sensitive to the fact that it was within her power to do the right thing at t instead. This is not an insignificant feature of the story.

Proponents of the buffer cases claim that there was no robust alternative available to Marla, no alternative the availability of which helps to explain why she is blameworthy for what she did. Evidently, then, they are committed to the following claim as well: that Marla was cognitively sensitive to the fact that she could have easily done the right thing at t rather than doing something at t which she knew to be immoral is irrelevant per se to an explanation of why Marla is blameworthy for her immoral action. This latter claim, however, is difficult to swallow.

When we think about the sorts of alternatives the availability of which might plausibly contribute to an explanation of why a person is blameworthy for what she did, the ones that most naturally come to mind are those in which the agent intentionally opts to do what she correctly believes to be morally right (i.e., obligatory or permissible) rather than doing something she knows to be morally wrong. It is alternatives like that which seem like they might be most relevant to whether or why a person is blameworthy for her immoral behavior. Indeed, if any

alternatives can be relevant per se to an explanation of why an agent is blameworthy for her conduct, surely it is those in which the agent intentionally opts to do the right thing rather than doing something she knows to be immoral. When such alternatives are available to an agent who acts immorally, a plausible explanation of why the agent is blameworthy for her bad behavior is that she did something she knew to be morally wrong despite being aware that it was within her power to instead behave in a way that was morally acceptable.⁹

An example will help illustrate the point. Someone in the office has started an ugly and potentially damaging rumor about one of Lindsay's co-workers, a rumor which Lindsay knows to be unfounded. However, Lindsay stands to gain a promotion if the maligned colleague is fired. So, when asked by her employer whether she can speak to the truth of the rumor, Lindsay deliberately lies to her boss, telling him that all the nasty things he has heard about this particular employee are true. She does this, moreover, despite knowing that it is immoral, and despite being aware that it was within her power to do the right thing instead by telling her employer the truth.

It seems clear that Lindsay is blameworthy for perpetuating the false rumor about her colleague. Moreover, the fact that she could have done the right thing by telling her boss the truth hardly seems irrelevant to an explanation of why she is blameworthy for her mendacious behavior. What seems to make Lindsay especially worthy of blame in this case is not just that her behavior was harmful or that she knew better than to do such a thing. Those facts are relevant, to be sure, but by themselves they do not provide a full explanation of why Lindsay is blameworthy. It also seems relevant that she could have done the right thing, thereby preserving her colleague's reputation and preventing him from losing his job. So at least part of the reason why Lindsay is blameworthy in this case is that she did something she knew to be immoral despite being cognitively sensitive to the fact that she could have done the right thing instead.

It thus seems quite plausible to suppose that an agent had a robust alternative to what she did at t , an alternative that contributes to an explanation of why the agent is blameworthy for her behavior, if the agent was aware, at least at some level, that it was within her power to avoid wrongdoing at the time and to instead behave at that time in a way that was morally acceptable. Notice, though, that the alternative course of action available to agents in the buffer cases fits that description. Marla was cognitively sensitive to the fact that she could have avoided doing anything immoral at t by simply pausing then to reflect on the reasons for not cheating. But rather than perform that morally permissible action at t , she elected instead to do something she knew to be immoral. These facts do not seem irrelevant per se to an explanation of why Marla is blameworthy for the decision she made at t , since we can explain why she is blameworthy for that decision in part by pointing out that she was cognitively sensitive to the fact that she could have behaved at t in a way that was above reproach but elected instead to do something she knew to be morally impermissible. Arguably, then, pausing at t to reflect on the reasons for not cheating counts as a robust alternative possibility for action after all.

A slightly different argument for this conclusion emerges when we consider what, if anything, we could reasonably have expected Marla to do at t instead of what she actually did at that time.¹⁰ A plausible test for whether an agent had a robust alternative to the immoral action she performed at t is to ask whether there was a morally permissible action the agent could have performed at t such that we could reasonably have expected her to perform it at t instead of doing what she actually did at t . If there were no alternatives fitting this description, that would be evidence that the agent lacked a robust alternative to what she did at t , as it would indicate that the agent lacked a reasonable opportunity to avoid wrongdoing at t . If, on the other hand, there was a morally permissible action the agent could have performed at t such that we could

reasonably have expected her to perform it at t instead of doing what she actually did at t , this, I suggest, would be an excellent candidate for a robust alternative, for then we could plausibly explain why the agent is blameworthy for the immoral action she performed at t in part by appealing to the fact that she had a reasonable opportunity to avoid doing anything wrong at t .

The alternative possibility for action available to an agent in the buffer cases passes this test for robustness. To see this, consider the story about Marla again. Is there a morally permissible action she could have performed at t such that we could reasonably have expected her to perform it at t instead of the action she actually performed at t ? Indeed, there is. Marla could have voluntarily paused at t to reflect on the reasons for not cheating on her philosophy midterm, and it is clear that this is something we could reasonably have expected her to do at t instead of deciding then to cheat on the exam. So once again we arrive at the conclusion that pausing at t to reflect on the reasons for not cheating is a robust alternative possibility.

4. A Reply to Pereboom's Argument

With this background in place, we are now prepared to see where the argument of Pereboom's introduced earlier goes wrong. Recall that according to Pereboom an alternative action is robust only if the agent was cognitively sensitive to the fact that if she voluntarily performed that action she would then likely be blameless for the action she actually performed. Call this Pereboom's requirement for robustness or PRR, for short. As we have seen, PRR implies that the alternative course of action available to Marla at t was not robust. If Marla had contemplated the reasons for not cheating, she would have avoided blameworthiness, since the neuroscientist would have intervened and forced her hand, thereby rendering her blameless. The problem, though, is that Marla was completely unaware of this. She was therefore not cognitively sensitive to the fact that

if she paused at t to consider the reasons for not cheating she would then avoid blameworthiness for deciding to cheat. Marla's only alternative thus does not satisfy PRR and so, Pereboom concludes, is irrelevant to whether or why she is blameworthy for her decision.

The crucial step in Pereboom's argument is PRR. But why suppose that principle is true? Why suppose it expresses a genuine requirement for robustness? I will argue that Pereboom does not offer a compelling answer to this question. I will also argue that even if PRR is true, there is an acceptable version of the principle that does not support Pereboom's conclusion.

Pereboom tries to motivate PRR by appealing to the following example. Joe decides at t to take an illegal deduction on his tax form, and the only way he could have avoided deciding to do so was by voluntarily taking a sip of his coffee, which, unbeknownst to him, has been spiked with a drug that induces compliance with the tax code. Pereboom claims that whether Joe could have taken a sip of coffee is irrelevant per se to whether or why Joe is blameworthy for the decision he made at t to take the illegal deduction on his taxes. Pereboom goes on to suggest that this is because "Joe has no inkling that taking the sip would render him blameless" (2014, p. 11).

Consider Pereboom's claim that whether Joe could have voluntarily taken a sip of coffee is irrelevant to explaining whether or why Joe is blameworthy for his decision to cheat on his taxes. Whether this claim is true, I shall argue, turns on the details of the case. Let t be the time at which Joe decided to take the illegal deduction. Evidently we are to suppose that Joe could have refrained from deciding at t to take the illegal deduction and taken a sip of his coffee at t instead. For ease of expression I shall henceforth say that Joe could have *simply* taken a sip of coffee at t , where "simply" indicates that Joe was not also making a decision then to cheat on his taxes. I contend that whether this alternative course of action is robust depends on whether Joe was cognitively sensitive to the fact that it was a *morally relevant alternative* to what he did at t ,

something morally acceptable he could do at the time instead of the immoral action he actually performed at that time, and not on whether Joe was cognitively sensitive to the fact that by availing himself of this alternative he would then be blameless for cheating on this year's taxes.

Suppose Joe was cognitively sensitive to the fact that simply taking a sip of coffee was a morally relevant alternative to deciding to take the illegal deduction on his tax form. Suppose, in other words, that Joe was aware, at least at some level, that simply taking a sip of coffee was something morally acceptable he could do at t instead of deciding then to cheat on his taxes. In that case, that Joe could have simply taken a sip of his coffee at t does not seem obviously irrelevant to whether or why Joe is blameworthy for the decision he made at t to cheat on his taxes, for we could plausibly explain why Joe is blameworthy for that decision in part by pointing out that he was cognitively sensitive to the fact that he could have behaved at t in a way that was morally acceptable but elected instead to do something at t that he knew to be immoral.

The reasonable expectations test for robustness introduced earlier supports this conclusion. Is there a morally permissible course of action Joe could have undertaken at t such that we could reasonably have expected him to undertake it at t rather than doing what he actually did at t ? It seems there is. Given our assumption that Joe was cognitively sensitive to the fact that simply taking a sip of coffee at t was a morally relevant alternative to what he actually did at t , it seems we could reasonably have expected him to do that at t instead of deciding then to take the illegal deduction. But if so, then the reasonable expectations test implies that simply taking a sip of coffee at t is a robust alternative to deciding at t to take the illegal deduction.

Now suppose that Joe was *not* cognitively sensitive to the fact that simply taking a sip of his coffee was a morally relevant alternative to deciding to cheat on his taxes. Suppose that, through no fault of his own, Joe had no clue that simply taking a sip of his coffee was something

morally acceptable he could do at t instead of deciding then to cheat on his taxes. In that case, the fact that Joe could have simply taken a sip of coffee at t may indeed be irrelevant to whether or why Joe is blameworthy for deciding as he did at t . What would explain its irrelevance, however, is not, as Pereboom suggests, that Joe had no idea that simply sipping his coffee at t would render him forever blameless for choosing to take the illegal deduction. Instead the explanation is that Joe was not cognitively sensitive to the fact that simply taking a sip of coffee at t was a morally relevant alternative to deciding then to take the illegal deduction.

That this is a better explanation than the one Pereboom offers is supported by the observation I made a moment ago that simply taking a sip of coffee would be a robust alternative to deciding to take the illegal deduction if only Joe were cognitively sensitive to the fact that simply taking a sip of coffee was something morally acceptable he could do at t instead of deciding then to cheat on his taxes. Joe need not also have been aware that simply taking a sip of coffee at t would render him blameless for deciding to cheat on this year's taxes. What makes the difference to whether an alternative is robust, it would seem, is whether the agent was cognitively sensitive to the fact that the alternative was a morally relevant alternative to the action the agent actually performed, and not whether the agent realized that availing himself of the alternative would render him forever blameless for doing what he actually did.

Pereboom's attempt to motivate PRR thus faces a dilemma. Either Joe was cognitively sensitive to the fact that simply taking a sip of coffee at t was a morally relevant alternative to deciding at t to cheat on his taxes, something morally acceptable he could do at the time instead of the immoral action he actually performed at that time, or he was not. If he was cognitively sensitive to that fact, then contrary to what Pereboom claims, whether Joe could have simply taken a sip of coffee at t is not obviously irrelevant to whether or why Joe is blameworthy for

what he actually did at t . If, on the other hand, Joe was not cognitively sensitive to the fact that simply taking a sip of his coffee was a morally relevant alternative to what he actually did at t , then the fact that he could have avoided deciding to cheat on his taxes by simply taking a sip of coffee at t may indeed be irrelevant to whether or why Joe is blameworthy for deciding to take the illegal deduction. Contrary to what Pereboom suggests, however, we can plausibly explain the irrelevance of that fact without appealing to the likes of PRR. We could explain it instead by appealing to the fact that, through no fault of his own, Joe was unaware that simply taking a sip of coffee at t was a morally relevant alternative to the immoral action he actually performed at that time. Either way, reflection on stories like the one about Joe does not provide us with a compelling reason to accept PRR, and I see no other reason to think that that principle is true.

Suppose, however, that PRR is true. Even so, it is unclear whether the principle supports Pereboom's conclusion that the alternative action available to an agent in Marla's position is not robust. There are at least two possible ways of reading PRR.¹¹ According to the first reading, an alternative action is robust only if the agent was cognitively sensitive to the fact that if she voluntarily performed that action she would then likely be *permanently* blameless for the action she actually performed. According to the second reading, an alternative action is robust only if the agent was cognitively sensitive to the fact that if she voluntarily performed that action she would then likely be *at least temporarily* blameless for the action she actually performed. On the first reading, the alternative action available to Marla is not robust, since, as I explained earlier, Marla was not cognitively sensitive to the fact that if she were to perform that alternative action at t she would then be permanently blameless for deciding to cheat on her exam. On the second reading, though, the alternative action available to an agent in Marla's position may very well be robust, for the agent presumably was aware that if she were to perform that alternative action at t

rather than deciding at t to cheat on her midterm she would then be at least temporarily blameless for deciding to cheat on the exam. She was presumably aware of this for the simple reason that if she had performed the alternative action at t instead of deciding then to cheat on her midterm, she would not yet have made a decision to cheat, and we may safely assume that she is smart enough to know that an agent cannot be blameworthy for an offense she has yet to commit.

So we have at least two possible readings of PRR, one that supports Pereboom's conclusion that agents in the buffer cases have no robust alternatives, and a second that does not support that conclusion. Pereboom seems to have the first reading in mind, clearly. However, on the assumption that a version of PRR is true, I contend that we have reason to prefer the second reading. Earlier I argued that there are independent reasons for thinking that the alternative action available to an agent in the buffer cases is robust. The agent was cognitively sensitive to the fact that, by performing that alternative action, she would be doing the right thing at t rather than doing something at t that she knew to be morally wrong. Moreover, we could reasonably have expected the agent to perform that alternative action at t rather doing what she actually did at t . The agent thus had a reasonable opportunity to avoid wrongdoing. These facts, I argued, support the conclusion that the alternative action available to the agent in buffer cases is, in fact, a robust alternative. So, assuming that PRR expresses a genuine requirement for robustness, these facts also provide us with reason to think that that requirement is satisfied. But, as we have just seen, the requirement is only satisfied on the second reading of PRR. Hence, my earlier arguments provide us with some reason to prefer that second reading.

Pereboom attempts to show that the alternative possibility for action available to an agent in Marla's position is not robust. His argument for that conclusion is unpersuasive for two

reasons. First, PRR, the argument's central premise, is insufficiently motivated. Second, even if PRR is true, there is a plausible reading of it that does not support Pereboom's conclusion.

5. The Timing Objection

I have argued that the buffer cases are not counterexamples to PAP; they do not provide us with scenarios in which a person is blameworthy for what she did despite lacking a robust alternative possibility for action. Central to my argument for this conclusion is the observation that the agent in buffer cases had it within her power at the time to do the right thing at that time. This observation has also been used as the basis for another objection to the buffer cases known as the timing objection.¹² As I hope to make clear, the objection to the buffer cases defended here is importantly different from the timing objection. Nevertheless, the two arguments are similar enough to warrant comparison.¹³ While the discussion that follows will bring to light some potential difficulties with the timing objection, I should like to make clear that my principal aim is not to show that the timing objection fails. Rather, my aim is to highlight what I think are some important advantages of the objection to the buffer cases developed in this paper.

According to proponents of the timing objection, Marla is not blameworthy for deciding to cheat on her exam *simpliciter*. This is because, through no fault of her own, she could not have avoided deciding to cheat on it. What Marla is blameworthy for, however, is deciding *precisely at t* to cheat on the exam. But, as proponents of the timing objection are quick to point out, Marla did have a robust alternative to deciding precisely at *t* to cheat, even by Pereboom's lights. Marla could have paused at *t* to reflect on the reasons for not cheating, and she was presumably cognitively sensitive to the fact that if she voluntarily performed this alternative action at *t*, she would then be *forever* blameless for deciding precisely at *t* to cheat. So, by paying close attention

to the exact time at which Marla acted, proponents of the timing objection insist that we find nothing for which Marla is blameworthy but for which she lacked a robust alternative.

The main difference between my objection to the buffer cases and the timing objection has to do with what we are willing to deem Marla blameworthy for. According to the view defended here, Marla may be blameworthy for deciding to cheat on her exam. Contrary to what defenders of the buffer cases contend, however, this fact poses no threat to PAP, since, as we have seen, there was a robust alternative available to her. By contrast, timing objectors insist that, because Marla could not have avoided deciding to cheat on the exam, she is not blameworthy for deciding to cheat on it. What she is blameworthy for, though, is the temporally more specific fact that she decided precisely at t to cheat on the exam. Attention to some worries about the timing objection will illustrate why this difference between the two objections matters.

Some critics of the timing objection are skeptical that ascriptions of blameworthiness are time sensitive in the way the objection requires.¹⁴ I am not convinced that their skepticism is warranted, but let us set that aside. The point I should like to make at present is that even if they are right, even if ascriptions of blameworthiness are not time-sensitive in the way required by the timing objection, that would pose no threat to the argument of this paper. Unlike the timing objection, the objection to the buffer cases developed here does not require us to time index ascriptions of blameworthiness and so avoids these sorts of worries altogether.

A more pressing worry for proponents of the timing objection has to do with their claim that Marla is not blameworthy for deciding to cheat on her exam.¹⁵ That claim is controversial, at best. After all, no one forced Marla to decide as she did, though the neuroscientist would have, had she not decided on her own at t to cheat. We may suppose that Marla satisfied a robust set of conditions that contemporary compatibilists and libertarians would deem sufficient for free

action. Furthermore, it was within her power to avoid deciding at t to cheat and to behave at that time in a way that was entirely above reproach. She thus had some control at or immediately prior to the time of action over whether she behaved immorally at that time. So not only did Marla do something she knew to be immoral, it appears that by most accounts she did it freely. But if she freely decided to cheat on the exam, knowing full well that it was the morally wrong choice, it is difficult to see why she is not worthy of at least some blame for deciding as she did.

A number of possible responses to this worry are available to proponents of the timing objection.¹⁶ For instance, they could adopt an account of action-individuation according to which the timing of a token action is essential to it.¹⁷ On such accounts, even though it was inevitable that Marla would make a decision to cheat, she could have avoided making the particular decision to cheat that she actually made, since she could have avoided deciding precisely at t to cheat. But if she could have avoided making that particular decision to cheat, then proponents of the timing objection could grant that the decision Marla actually made to cheat was a free action of hers, and, crucially, that she is blameworthy for it, without having to abandon PAP.

The key to this defense of the timing objection is the assumption that actions are individuated in part by the times at which they occur. But that assumption is hardly uncontroversial. Indeed, I am inclined to think it is false. I will not argue for that conclusion here, however. Instead, I should simply like to point out that the position developed in this paper again enables us to bypass these issues entirely. It allows us to accept the claim that Marla freely decided to cheat, and that she is blameworthy for the decision she made, without having to take a stand on the vexed question of action individuation, and without having to jettison PAP.

It seems to me, then, that the objection to the buffer cases on offer here has at least two important advantages over the timing objection. First, because it does not restrict what Marla is

blameworthy for to the temporally specific fact that she decided precisely at t to cheat on the exam, the objection to the buffer cases defended here does not require us to engage in the controversial practice of time indexing ascriptions of blameworthiness. Second, it allows us to accept the plausible claim that the featured agent in buffer cases is blameworthy for the decision she made, without forcing us to take a stand on the difficult question of action-individuation. To be clear, none of this is to say that the timing objection is indefensible. For all I have said here, that objection may succeed, and if it does, I certainly would not begrudge that fact. There may be more than one way to skin the cat in this case. Some ways, though, are easier than others.¹⁸

Justin A. Capes

East Tennessee State University

Reference

- Fischer, John Martin. 1994. *The Metaphysics of Free Will* (Cambridge, Mass.: Blackwell Publishers).
- Frankfurt, Harry. 1969. "Alternate Possibilities and Moral Responsibility," *Journal of Philosophy*, vol. 66, no. 3, pp. 829-839.
- Franklin, Christopher E. 2011. "Neo-Frankfurtians and Buffer Cases: The New Challenge to the Principle of Alternative Possibilities," *Philosophical Studies*, vol. 152, n. 2, pp. 189-207.
- Ginet, Carl. 1996. "In Defense of the Principle of Alternative Possibilities: Why I Don't Find Frankfurt's Argument Convincing," *Philosophical Perspectives*, vol. 10, pp. 403-417.
- , 2002. "Review of *Living Without Free Will*," *Journal of Ethics*, vol. 6, no. 3, pp. 305-309.

- Hunt, David. 2000. "Moral Responsibility and Unavoidable Action," *Philosophical Studies*, vol. 97, no. 2, pp. 195-227.
- . 2005. "Moral Responsibility and Buffered Alternatives," *Midwest Studies in Philosophy*, vol. 29, no. 1, pp. 126-145.
- Hunt, David and Shabo, Seth. 2013. "Frankfurt Cases and the (In)Significance of Timing," *Philosophical Studies*, vol. 164, no. 3, pp. 599-622.
- Kane, Robert. 1985. *Free Will and Values* (Albany, NY: State University of New York Press).
- Mele, Alfred. 2006. *Free Will and Luck* (New York: Oxford University Press).
- Otsuka, Michael. 1998. "Incompatibilism and the Avoidability of Blame," *Ethics*, vol. 108, no. 4, pp. 685-701.
- Palmer, David. 2011. "Pereboom on the Frankfurt Cases," *Philosophical Studies*, vol. 153, no. 2, pp. 261-272.
- . 2013. "The Timing Objection to the Frankfurt Cases," *Erkenntnis*, vol. 78, no. 5: 1011-1023.
- Pereboom, Derk. 2000. "Alternate Possibilities and Causal Histories," *Philosophical Perspectives*, vol. 14, pp. 119-138.
- . 2001. *Living Without Free Will* (New York: Cambridge University Press).
- . 2005. "Defending Hard Incompatibilism," *Midwest Studies in Philosophy*, vol. 29, no. 1, pp. 228-247.
- . 2009. "Further Thoughts About a Frankfurt-style Argument," *Philosophical Explorations*, vol. 12, no. 2, pp. 109-118.
- . 2012. "Frankfurt Examples, Derivative Responsibility, and the Timing Objection," *Philosophical Issues*, vol. 22, no. 1, pp. 298-315.

- . 2014. *Free Will, Agency, and Meaning in Life* (New York: Oxford University Press).
- Shabo, Seth. 2007. "Flickers of Freedom and Modes of Action: A Reply to Timpe," *Philosophia*, vol. 35, no. 1, pp. 63-74.
- Widerker, David. 1995. "Libertarianism and Frankfurt's Attack on the Principle of Alternative Possibilities," *The Philosophical Review*, vol. 104, no. 2, pp. 247-261.
- . 2000. "Frankfurt's Attack on Alternative Possibilities: A Further Look. *Philosophical Perspectives*, vol. 14, pp. 181-201.
- . 2003. "Blameworthiness, and Frankfurt's Argument Against the Principle of Alternative Possibilities," in *Moral Responsibility and Alternative Possibilities: Essays on the Importance of Alternative Possibilities*, ed. David Widerker and Michael McKenna (Aldershot, UK: Ashgate), pp. 53-73.

Notes

For helpful comments on earlier versions of this paper I'm grateful to Carl Ginet, David Hunt, Al Mele, Michael McKenna, David Palmer, and two anonymous referees for this journal.

¹ The idea that PAP requires only "robust" alternatives is originally due to Fischer (1994) pp. 140-147. The definition of a robust alternative given in the text is the standard one. For a slightly different use of the term "robust alternative," see Mele (2006), p. 92.

² PAP is typically construed as a principle about moral responsibility, in which case it would apply to both praiseworthiness and blameworthiness. I have elected to focus on a version of the principle that is restricted to blameworthiness for two reasons. First, as I just mentioned, PAP is

often said to provide a principled basis for the exculpatory force of pleas like “I couldn’t help it” and “I had no choice.” Because such pleas are typically offered as excuses—considerations designed to show that a person is not to blame for some offense—it is natural to focus on a version of PAP restricted to blameworthiness. Second, and more importantly, praiseworthiness raises its own distinct set of issues that I think are best addressed separately. For example, some people equate moral responsibility with blameworthiness, while others insist that a person can be morally responsible without being blameworthy (e.g., by being praiseworthy instead). And even among those who agree that moral responsibility encompasses both praiseworthiness and blameworthiness, there is disagreement about whether the conditions for praiseworthiness and blameworthiness are the same. For instance, some philosophers insist that there is an asymmetry between praiseworthiness and blameworthiness in that blameworthiness requires having alternative possibilities for action, whereas praiseworthiness does not. By focusing on a version of PAP restricted to blameworthiness, my aim is to stay neutral on these disputes.

³ See Frankfurt (1969).

⁴ Hunt (2000, 2005) and Pereboom (2000, 2001) independently developed the buffer cases. Others who have defended them include Hunt and Shabo (2013) and Shabo (2007).

⁵ This “dilemma defense” was first suggested by Kane (1985), p. 51. Versions of it have subsequently been developed and defended by Widerker (1995) and Ginet (1996), among others.

⁶ See, for example, Hunt (2000, 2005), Pereboom (2000, 2001, 2012, 2014), and Hunt and Shabo (2013).

⁷ I will focus on the most recent version of Pereboom’s argument, which is found in Pereboom (2014). Earlier versions can be found in Pereboom (2000, 2001, 2005, 2009, and 2012). Hunt and Shabo (2013) also defend a version of this argument against some recent objections to it. I

focus on Pereboom's version in what follows, but what I have to say in response to it can be applied, *mutatis mutandis*, to the version defended by Hunt and Shabo.

⁸ An earlier version of this requirement said that an alternative action is robust only if the agent "understood" that if she voluntarily performed that action instead she would then likely be blameless for the action she actually performed. But Pereboom has since come to believe that an agent can have a robust alternative even if the agent did not fully understand that availing herself of the relevant alternative would render her blameless. What is required, Pereboom now says, is not full understanding, but rather some lower level of "cognitive sensitivity" to the fact that availing herself of the alternative in question would render her blameless. See Pereboom (2014), pp. 11-14 for his discussion of these points.

⁹ For a similar position, see Otsuka (1998), pp. 694-698.

¹⁰ Cf. Widerker (2000) and (2003).

¹¹ A third reading is suggested by Hunt and Shabo (2013). As far as I am aware, Franklin (2011) was the first to distinguish the two readings of PRR discussed in the text. The argument that follows builds on his critique of Pereboom's argument.

¹² The timing objection was first developed by Carl Ginet (1996, 2002). Versions of it have recently been defended by Franklin (2011) and Palmer (2011, 2013).

¹³ It could perhaps be said that the objection on offer here is a version of the timing objection, insofar as it insists that the featured agent in buffer cases had a robust alternative available to her *at the time of action*. I have no objection to this, as long as we keep the differences between the two arguments clearly in view.

¹⁴ See Hunt and Shabo (2013) for an objection along these lines.

¹⁵ Hunt and Shabo (2013) also press this objection.

¹⁶ See Palmer (2013), p. 1018 for two such responses.

¹⁷ Note that this might also help proponents of the timing objection respond to the earlier worry about ascriptions of blameworthiness being time indexed. As Palmer (2013), p. 1018 points out, if token actions are individuated in part by the times at which they occur, it would not make sense to speak of an action without at least an implicit reference to the time at which the action is performed, since the time at which the action occurs is part of what makes it the particular action it is. But if so, would it make any more sense to talk of being blameworthy for an action without at least implicit reference to the time at which the action was performed?

¹⁸ Pereboom (2014), pp. 22-29 has produced a new Frankfurt-style example that he claims avoids the timing objection. In response, Palmer (2013) argues convincingly that this new example is no less vulnerable to the timing objection than are the older buffer cases. In any event, because the agent in the new example made a morally bad decision at t , and because the agent was cognitively sensitive to the fact that he could have done something morally permissible at t instead, the same sorts of arguments given in section 3 would also show that the agent in Pereboom's new example has a robust alternative possibility available to him.