

Luckily, We Are Only Responsible for What We Could Have Avoided

PHILIP SWENSON

The view that moral responsibility requires making a difference faces challenges from both cases of causal overdetermination and cases of causal preemption (such as Frankfurt-style Cases). In some such cases it seems that an agent is responsible for something that they could not have avoided. To deal with these cases, I enlist help from a plausible approach to moral luck. Following John Martin Fischer (1986) and Michael Zimmerman (2002), I defend the view that one's degree of responsibility is immune to moral luck but the scope of events for which one is responsible is subject to moral luck. I then argue that this view leads to a plausible error theory for our responsibility intuitions concerning cases of causal overdetermination and preemption. This error theory allows us to avoid counterexamples to the claim that responsibility requires making a difference.

1. INTRODUCTION

I aim to defend the view that being morally responsible for something requires being able to make a difference with regard to the occurrence of that thing. I endorse the following requirement:

Difference Making Requirement: An agent is morally responsible for an action, omission, outcome, etc. X, only if she could have acted in such a way that X would (or at least might) not have occurred.¹

The Difference Making Requirement is a fairly strong requirement. Some *alternative possibility* requirements on responsibility only require that, in order to be responsible for X, an agent must have been able to act differently. They do not require that her alternative action avoid X. The Difference Making Requirement holds that the agent must have a chance to avoid X. The Difference Making Requirement is *prima facie* motivated by cases like the following:

Sharks: John is walking along the beach and sees a child drowning in the water. John believes that he could rescue the child without

1. The "might" alternative comes into play if there are cases in which there is no fact about whether an agent's acting differently would have prevented X.

much effort. Due to his laziness, he decides not to attempt to rescue the child. The child drowns. Unbeknownst to John, there is a school of sharks hidden beneath the water. If John had attempted to rescue the child, the sharks would have killed him and his rescue attempt would have been unsuccessful.²

In *Sharks* it seems that John is not responsible for the omission *failing to save the child* or for the outcome *the child's death*. One natural way to account for his lack of responsibility is by appealing to something along the lines of the Difference Making Requirement. After all, if the sharks hadn't been there, John could have made a difference and would have been responsible for the child's death.³ (Also, note that merely requiring that the agent be able to do *something* different in order to be responsible for X would not explain John's lack of responsibility for the death. John could have done something different. He could have tried to rescue the child. Thus *Sharks* provides a reason to favor the Difference Making Requirement over weaker *alternative possibility* requirements.)

However, cases of overdetermination make trouble. Consider the following:

Overdetermination Case: Alissa and Emily both freely throw a rock at a window. Both rocks cause the window to break. Alissa was not aware that Emily was planning to throw a rock and could not have done anything to prevent Emily from breaking the window.

And:

Preemption Case: Alissa freely throws a rock which breaks a window. Emily does not throw a rock, but she would have broken the window if Alissa hadn't. Alissa did not know that Emily would break the window if she refrained. And Alissa would not have been able to prevent Emily from breaking the window.

In both of these cases it seems that Alissa is responsible for the destruction of the window. But she could not do anything to avoid its destruction. Consider also the following more sophisticated case of counterfactual overdetermination:

Frankfurt Case: Black wishes Jones to cast his vote for presidential candidate A. In order to ensure that Jones does this, he implants a chip in Jones's brain which allows him to control Jones's behavior in the voting booth. (Jones has no idea about any of this.) Black prefers that Jones vote for candidate A on his own. But if Jones starts to show any sign that he will vote for anyone other than A, Black will immediately use his chip to cause Jones to vote for candidate A instead.

2. This case is drawn from Fischer and Ravizza (1998: 125).

3. There are rival accounts of John's lack of responsibility. I attempt to undermine such accounts in Swenson (2016a). See also Swenson (2015) where I appeal to *Sharks* in order to defend alternative possibility requirements on responsibility.

As it turns out, though, Jones votes for candidate A on his own and Black never exerts any causal influence on Jones's behavior.⁴

In *Frankfurt Case* it seems that Jones is responsible for *voting for A*, even though he could not have avoided doing so. So it looks like we have three counterexamples to the Difference Making Requirement. But perhaps our intuitive reaction to these cases rests on a mistake.

2. THE HYBRID APPROACH TO MORAL LUCK

Zimmerman offers us this pair of scenarios:

Successful Murder: "George shot at Henry and killed him."

Attempted Murder: "Georg shot at Henrik in circumstances which were, to the extent possible, exactly like those of George (by which I mean to include what went on 'inside' the protagonists' heads as well as what happened in the 'outside' world), except for the fact that Georg's bullet was intercepted by a passing bird (a rather large and solid bird) and Henrik escaped injury." (Zimmerman 2002: 560)

Here are two takes one might have on these cases:

Pro-luck take: George is blameworthy for Henry's death. But Georg is not blameworthy for Henrik's death (since he didn't die). So George is more blameworthy than Georg.

Anti-luck take: Georg and George are equally blameworthy. So George is not blameworthy for anything beyond what Georg is blameworthy for. Thus George is not blameworthy for Henry's death.

Both of these takes start out with a plausible claim and derive an implausible conclusion. It seems clear to many of us that "Georg and George are equally blameworthy." But it also seems clear that "George is blameworthy for Henry's death." Fortunately, there is a way to accommodate both claims.

Fischer (1986) notes a distinction between the *degree* of blameworthiness and the *content* of blameworthiness. Zimmerman (2002) notes a very similar distinction between the *degree* of responsibility and the *scope* of responsibility.⁵ Both apply the distinction to the issue of moral luck. The degree of blameworthiness reflects how much blame it is appropriate for the agent to receive. The scope/content of blameworthiness reflects which events the agent is blameworthy for. The Fischer and Zimmerman view is that the scope of blameworthiness is affected by luck in how things turn out (e.g., by whether or not a bird interferes with your attempt to kill). But the degree of blameworthiness is immune to luck. This leads to the following view:

Hybrid take: Georg and George are equally blameworthy. That is to say that they have the same degree of blameworthiness. But George

4. The origin of such cases is Frankfurt (1969).

5. Frankfurt (1982) also points out this sort of distinction.

is blameworthy for more events than Georg, because he is blameworthy for killing Henry. So the scope of George's blameworthiness exceeds the scope of Georg's blameworthiness.

The three takes on Georg and George correspond to three views one might have concerning resultant moral luck (i.e., luck in how things turn out downstream from one's choices and attempts).

Pro-luck approach: Both the degree and scope of blameworthiness are subject to resultant luck.

Anti-luck approach: Neither the degree nor the scope of blameworthiness is subject to resultant luck.

Hybrid approach to luck: The scope of blameworthiness is subject to resultant luck, but the degree of blameworthiness is not.

I accept the hybrid approach. (I formulate these approaches in terms of blameworthiness. But they could also be put in terms of moral responsibility more generally. I also accept the hybrid approach when it comes to moral responsibility in general.)

The hybrid approach nicely explains why many of us feel that it would be unfair to blame George more than Georg. It's because they deserve the same degree of blame! And unlike the anti-luck approach, it also makes sense of the fact that we would blame George *for killing Henry*, not just for *trying to kill Henry*. The anti-luck approach must hold that we are often blaming people for the wrong things. While the pro-luck approach leads to apparent unfairness. The hybrid approach thus appears to have an edge over both of its rivals.

The key insight of the hybrid approach is that agents can differ in the scope of their blameworthiness without differing in their degree of blameworthiness. This allows the hybrid approach to generate attractive results in cases like *Sharks* as well. Compare *Sharks* with the following:

Absent Sharks: Jane is walking along the beach and sees a child drowning in the water. Jane believes that she could rescue the child without much effort. Due to her laziness, she decides not to attempt to rescue the child. The child drowns. There were no sharks in the area. If Jane had tried to rescue the child, she would have succeeded.

In *Absent Sharks* Jane is blameworthy for the child's death. This means that the scope of her blameworthiness is larger than the scope of John's in *Sharks* (since the presence of the sharks prevented John from being responsible for the death). However, does it seem that Jane deserves more blame than John? I think not. Given that both of them thought they could rescue the child, it seems that they deserve the same amount of blame. The hybrid approach allows us to say both that (1) Jane is blameworthy for more events than John and (2) Jane and John deserve the same amount of blame.⁶

6. Fischer (1986) applies the hybrid approach to similar cases.

Another nice feature of the hybrid approach is that it provides (at least a partial) explanation of why resultant moral luck is so puzzling. We feel pulled toward both the pro-luck and anti-luck positions. The hybrid approach helps make sense of this because, given the hybrid approach, both the pro-luck and anti-luck positions have some of the truth. The pro-luck position is right about the scope of blameworthiness and the anti-luck position is right about the degree of blameworthiness.

The problem is that we have not sufficiently separated questions regarding scope of blame from questions regarding degree of blame. We are pulled in two directions because we are failing to clearly distinguish between two different phenomena. We are conflating the scope of blameworthiness with the degree of blameworthiness.

I am unsure if this purported conflation can fully explain away the problem of resultant luck all by itself. It might just seem to some that George deserves more blame than Georg. One way to handle such intuitions is to claim that they are based on a bad (by the lights of the hybrid approach), implicit inference from the claim that George is blameworthy for killing and Georg is not. But suppose that is not always the case. We could then appeal to additional error theories to help handle such cases. For example, we could appeal to the thought that we often take successfully causing harm to be evidence of a more culpable mental state. And this leads us to overgeneralize and end up with pro-luck intuitions even when we stipulate the mental states of the two agents we are comparing are the same.⁷

3. QUESTIONS ABOUT THE HYBRID APPROACH

In this section I will attempt to answer two questions that naturally arise when one considers the hybrid approach. First:

How should we think about one's degree of blameworthiness for a particular act?

The hybrid approach says that George and Georg deserve the same total amount of blame, even though George is blameworthy for more events. How does this work? Surely George's success in killing Henry does not make him less blameworthy for trying to kill Henry. So how can his total amount of blameworthiness be no greater than Georg's?

One answer would be that degree of blame attaches only to attempts. So perhaps George deserves ten units of blame for attempting to kill Henry. But, though he is also blameworthy for killing Henry, he does not deserve any amount of blame for it. The obvious problem with this view is that it seems to make any amount of blame for killing Henry excessive. And the hybrid view is supposed to capture the intuitive thought that we can blame George *for killing Henry*.

7. Lewis (1989) considers a similar account.

A better answer, I think, is to say that in many cases there is no precise degree of blame (or punishment) which is uniquely fitting for each act. Rather, there are multiple permissible ways of divvying up the blame. Suppose George and Georg each deserve ten units of blame. In Georg's case, we cannot blame him for killing, so we can only distribute the blame among entities such as the attempt, the decision to try to kill, and so on. In George's case, however, I think it would be permissible to blame him only for the killing (and spend all ten units worth of blame there). But it might also be permissible to blame George some for the attempt, and blame him some more for the killing. Suppose, for example, that you blame him for the attempt before learning whether he succeeded, you should not conclude that you have gone wrong in blaming him for the attempt once you learn that he succeeded in killing. (Though it might be disrespectful to Henry to blame only for the attempt in cases where you *know* that George succeeded in killing Henry.) So there is no unique amount of blame George deserves for the killing; rather, he deserves some amount of blame for his entire course of action (the same amount of blame Georg deserves), and we can divide it up in different ways.

Now for the second question about the hybrid approach. It is clear that one's degree of blameworthiness is morally significant. It affects how much one can be permissibly blamed or punished. But one might wonder whether the scope of blameworthiness is of any importance. Harry Frankfurt (1982) appears to suggest that the scope of blameworthiness is a matter of "no *moral* significance". So my second question is the following:

Is the scope of blameworthiness morally relevant?

Would it be a problem for the hybrid approach if the scope of blameworthiness had no significance? Perhaps so. Perhaps it is implausible that our moral intuitions would track a property that does not matter. The hybrid approach is more plausible if it can account for the moral significance of scope. I will argue that there are (at least) two ways in which scope is morally relevant even if the hybrid approach is correct.

First, scope might affect an agent's degree of liability, the extent to which it is fitting to harm the agent in order to prevent or compensate for some harm. Two agents who have the same degree of blameworthiness might, nonetheless, differ with regard to their degree of liability. Consider the following:

House Destruction 1: Neal and Chris both fire rockets into the air. Neal's rocket lands on your house, causing it to burn down. Chris's rocket causes no harm.

According to the hybrid approach, Neal and Chris have the same degree of blameworthiness. Nevertheless it seems that it would be more fitting for Neal to compensate you for the damage to your house. After all, Neal is responsible for destroying your house! One might be tempted to claim that, ideally, Neal and Chris (and anyone else who performed an equally reckless act) should equally share in the costs of compensating you. That might be correct.

However, suppose that, for some reason, only one of them can be made to compensate you. It seems that it is more fitting if Neal is the one chosen. So there is still an asymmetry with regard to liability.

One might object that Neal is more liable because he caused the damage, not because he is responsible for the damage. But consider this case:

House Destruction 2: Neal culpably fires a rocket in the air. Chris nonculpably (accidentally) fires a rocket in the air. Both rockets land on your house, jointly causing it to burn down.

Here they both cause the damage, but only Neal is responsible for causing the damage. And it still seems that Neal is more liable. Thus, scope of responsibility provides a better explanation of Neal's liability.

But now one might object that in *House Destruction 2* it is Neal's degree of blameworthiness that explains his increased liability. Chris is not blameworthy at all, and Neal is blameworthy. So we can account for the difference between them by appeal to degree of blameworthiness, rather than scope of blameworthiness. But consider the following:

House Destruction 3: Neal culpably fires a rocket in the air. Chris nonculpably (accidentally) fires a rocket in the air. Both rockets land on your house, jointly causing it to burn down. Chris has previously committed fraud against someone else, so (as it happens) his total degree of blameworthiness is the same as Neal's.

In this case their total degree of blameworthiness is the same and they both cause the destruction of your house. But it still seems that Neal is more liable. Appealing to the fact that Neal is responsible for causing the damage allows us to explain this. So it looks like there is a role for scope of blameworthiness in accounting for liability.⁸

The claim that Neal is more liable than Chris in *House Destruction 1* commits me to the claim that one's degree of liability is sensitive to luck. One might wonder why the degree of liability is sensitive to luck if, as the

8. A determined critic might object that Neal is more liable than Chris in *House Destruction 3* because he is responsible for a choice that caused your house to burn down, not because he is responsible for burning your house down. (Thanks to Peter Sutton for this worry.) I reply first that this view still demonstrates the relevance of scope. Neal and Chris have the same degree of blameworthiness. The particular choice Neal is responsible for is what explains his increased liability. Second, consider the following:

House Destruction 4: Neal culpably fires a rocket in the air. The rocket lands on your house, causing it to burn down. Chris commits fraud against someone else, so (as it happens) his total degree of blameworthiness is the same as Neal's. Due to a completely unforeseeable chain of events, Chris decision to commit fraud also causes your house to burn down.

Here both Chris and Neal are responsible for a choice that caused your house to burn down. But only Neal is responsible for burning your house down. Thus it looks as though it's his responsibility for burning down your house that explains his increased liability.

hybrid approach has it, the degree of blameworthiness is immune to luck?⁹ Here is a tentative answer to this question. In cases of liability it is fixed that some harm will occur. Either you will suffer the uncompensated loss of your house or someone will suffer the harm of compensating you. So luck is playing a role in determining who should suffer harm (given that harm is unavoidable). The fact that Neal is liable to be harmed does not entail that there would be any positive reason to harm Neal, if all harm could be avoided.¹⁰

Blameworthiness, on the other hand, involves the introduction of a reason to create harm even if no harm need occur. (Here I am appealing to a retributive view of blameworthiness.) Blameworthiness entails that the agent deserves some degree of harm. Since no one need suffer harm, there is not the same pressure (as there is in cases of liability) to allow luck to determine who should suffer the harm.¹¹

Now for the second way in which scope is morally relevant. Scope is relevant to the value of relationships. (In defending this claim I will switch the focus to the scope of moral responsibility more generally.) Compare the following two cases:

Rescue 1: Justin appears to see his friend Taylor drowning. He goes to great effort to throw a raft out to him. But Taylor was not actually drowning. Justin's act does not accomplish anything.

Rescue 2: Justin appears to see his friend Taylor drowning. He goes to great effort to throw a raft out to him. Taylor uses the raft to avoid drowning. Taylor would have drowned if Justin had not acted.

It seems that the value of Justin and Taylor's relationship has increased more in *Rescue 2* than in *Rescue 1*. In *Rescue 1*, Justin is responsible for *trying to aid Taylor*. And that is a good thing. But in *Rescue 2* Justin is responsible for rescuing Taylor. And this seems more valuable. Taylor can view Justin as responsible for the fact that he is alive. And Justin knows that he is responsible for making a major difference in Taylor's life. These facts seem to add value to their relationship. Thus the increased scope of Taylor's responsibility increases the value of their relationship.

4. IN DEFENSE OF THE DIFFERENCE MAKING REQUIREMENT

I now return to the Difference Making Requirement. Recall this potential counterexample:

9. Daniel Statman (2015) argues that accepting that luck can result in an innocent threat being liable to defensive harm commits one to accepting moral luck.

10. Thanks to Jimmy Goodrich for helpful discussion of this feature of liability.

11. For a defense of a retributive account of blameworthiness similar to the one I am appealing to here see Rosen (2015).

Overdetermination Case: Alissa and Emily both freely throw a rock at a window. Both rocks cause the window to break. Alissa was not aware that Emily was planning to throw a rock and could not have done anything to prevent Emily from breaking the window.

We are naturally inclined to say that Alissa is responsible for the destruction of the window. But I want to suggest that this inclination is based on a conflation of the degree of responsibility with the scope of responsibility. We correctly intuit that Alissa bears the full degree of blame regarding her act (despite the overdetermination). But we end up falsely judging that she possesses the normal scope of blameworthiness.

Recall that on the hybrid account one's scope of blameworthiness can change without one's degree of blameworthiness changing. I think the core intuition that we have about Alissa is that *she does not get off the hook (even to a small degree)* because of the presence of Emily. This core intuition is correct. But it is an intuition about her degree of blameworthiness. Because we fail to hold the distinction between degree and scope firmly in mind and attend to the fact that a change in Alissa's scope of blameworthiness need not render her deserving of less blame, we end up inclined to assert that she is blameworthy for the destruction of the window.

When a student is inclined to say that some act that caused unforeseeable harm is not wrong, I often remind them of the distinction between wrongness and blameworthiness. (You've probably done the same.) Oftentimes, with that distinction in mind, the student will be more inclined to say that the act was wrong. I am suggesting something similar is going on when we consider cases of overdetermination. At least in my own case, when I consider Alissa's status with the distinction between scope and degree firmly in mind, and note that she is blameworthy to the same degree either way, I am much less inclined to insist that she is blameworthy for the destruction of the window.¹²

Now consider the following:

Preemption Case: Alissa freely throws a rock which breaks a window. Emily does not throw a rock, but she would have broken the window if Alissa hadn't. Alissa did not know that Emily would break the window if she refrained. And Alissa would not have been able to prevent Emily from breaking the window.

12. Another analogy to the strategy I am pursuing is a response some externalists about epistemic justification have given to the New Evil Demon problem. They explain away the intuition that our demon world counterparts have justified beliefs by claiming that the demon worlders do have some positive status (which is driving the intuition). But that positive status is really excused belief, not justified belief. See Littlejohn (Forthcoming) and Williamson (Forthcoming). For a response, see Madison (2018).

Once again, the driving intuition is that Alissa does not get off the hook (even a little) because of Emily's presence. But if this is an intuition about degree of blameworthiness, not scope of blameworthiness, it is not a threat to the Difference Making Requirement.

One might object as follows. I have granted that the scope of responsibility is relevant to one's degree of liability. In *Preemption Case* it seems that Alissa is fully liable for the damage to the window. We have just as much justification for requiring her to compensate for the damage as we would have if Emily had not been present. Can we make sense of her full liability if we deny that she is responsible for the destruction of the window?

I think we can make sense of Alissa's full liability. Note that Alissa could have avoided causing the destruction of the window. The window would still have been destroyed, but not by her. Thus it is plausible that she is responsible for *causing the destruction of the window*. I claim that being responsible for causing a particular harm is what is relevant (so far as scope is concerned) regarding liability for that harm.¹³ So scope is relevant to liability, but what matters is responsibility for causing the harm, not responsibility for the occurrence of the harm.

I turn now to the Frankfurt cases. Recall:

Frankfurt Case: Black wishes Jones to cast his vote for presidential candidate A. In order to ensure that Jones does this, he implants a chip in Jones's brain which allows him to control Jones's behavior in the voting booth. (Jones has no idea about any of this.) Black prefers that Jones vote for candidate A on his own. But if Jones starts to show any sign that he will vote for anyone other than A, Black will immediately use his chip to cause Jones to vote for candidate A instead. As it turns out, though, Jones votes for candidate A on his own and Black never exerts any causal influence on Jones's behavior.

Once again, I claim that the core intuition is that Jones *does not get off the hook (even to a small degree)* because of the presence of Black. But this is an intuition about degree, not scope.

Things are somewhat more difficult here. In Alissa's case it is clear that she still has *some* scope of responsibility even if she is not responsible for the destruction of the window. She is responsible for throwing the rock. And she is responsible for causing the destruction of the window. But in *Frankfurt Case* one might worry that there is nothing that Jones is responsible for which he could have avoided. In that case, in order to say that he still has a high degree of responsibility, we would have to say that he is responsible to a high degree even though he is responsible for nothing. But it seems highly

13. Or perhaps, if what I say about the Frankfurt cases below is correct, what really matters for liability is responsibility for *causing the harm on your own*.

plausible that in order to be responsible one must be responsible for something!¹⁴

Fortunately, I think we can find something for which Jones is responsible. The *fine-grained response* to the Frankfurt cases holds that Jones is responsible for *voting for A on his own* (which he could have avoided), even though he is not responsible for *voting for A* (which he could not have avoided).¹⁵ Here “on his own” means “not as a result of outside force or coercion” (Capes and Swenson 2017).

One worry for the *fine-grained response* is that it allows for unacceptable moral luck. After all, Jones would have been responsible for *voting for A* if Black had not been present.¹⁶ So Black’s presence (which is just a matter of luck) changes what Jones is responsible for. Justin Capes and I have responded to this worry by appealing to the hybrid account. Luck does impact the scope of Jones’s responsibility. But it does not affect his degree of responsibility (Capes and Swenson 2017).

Though she ultimately rejects our view, Carolina Sartorio points out a significant advantage of our approach:

The underlying thought is that luck can affect what you are responsible for without affecting the degree of your responsibility, which would remain the same to the extent that you made the decision on your own. This view has the potential advantage that it allows a proponent of [views along the lines of the *fine-grained response*] to capture more of Frankfurt’s original intuitions about Frankfurt-style cases, which many people have found convincing. For it allows us to say that there is a sense in which agents in Frankfurt-style cases are just as blameworthy as if they had been in ordinary circumstances, although they are responsible for fewer things. (Sartorio 2019: 7)

I think this is exactly right. I add the claim that this view captures *the core intuition* behind the Frankfurt cases. On the view I am putting forward, we are correctly intuiting that Jones is “just as blameworthy” as he would have been in normal circumstances. But we mistakenly slide to the thought that he must be blameworthy for just as many events. As with Alissa above, when

14. Although Zimmerman’s (2002) view seems to allow for degrees of responsibility without any responsibility. The possibility of degrees of responsibility without responsibility for anything makes room for an interesting potential strategy for compatibilists about responsibility and determinism. They could grant the Difference Making Requirement and grant that, as a result, determinism rules out responsibility for anything at all. But they could claim that, nonetheless, determinism is compatible with possessing nonzero degrees of responsibility. Of course, compatibilists could also pursue the more standard strategy of accepting the Difference Making Requirement but claiming that agents have the ability to avoid certain events even if determinism is true. I develop an incompatibilist account of ability in Swenson (2016b). (See also Section 5 of Cyr and Swenson 2019).

15. Michael Robinson (2012, 2014, Forthcoming) defends the similar view that Jones is *basically* responsible only for acting on his own. But, on Robinson’s view, Jones is also derivatively responsible for voting for A. Thus Robinson must reject the Difference Making Requirement.

16. Linda Zagzebski (2000) raises this worry.

I consider Jones's status with the distinction between scope and degree firmly in mind, and note that he is blameworthy to the same degree either way, I am much less inclined to insist that she is blameworthy for *voting for A*.

It looks as though, so long as the *fine-grained response* is plausible, I can explain away our intuitions about Frankfurt cases in the same way I explain away our intuitions about standard overdetermination and preemption cases. The view that we are only responsible for what we could have avoided is intuitive in its own right and is motivated by cases like *Sharks*. Now that we have an error theory for intuitions that initially seem to point in the other direction, the claim that responsibility requires making a difference appears quite plausible.¹⁷

REFERENCES

- Capes, Justin, and Philip Swenson. 2017. "Frankfurt Cases: The Fine-Grained Response Revisited." *Philosophical Studies* 174(4): 967–81.
- Cyr, Taylor W., and Philip Swenson. 2019. "Moral Responsibility without General Ability." *Philosophical Quarterly* 69(274): 22–40.
- Fischer, John Martin. 1986. "Responsibility and Failure." *Proceedings of the Aristotelian Society* 86: 251–70.
- Fischer, John Martin, and Mark Ravizza. 1998. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Frankfurt, Harry. 1969. "Alternate Possibilities and Moral Responsibility." *Journal of Philosophy* 66: 829–39.
- . 1982. "What We Are Morally Responsible For." In *How Many Questions?*, ed. L. Cauman et al., 321–35., Indianapolis, IN: Hackett.
- Lewis, David. 1989. "The Punishment That Leaves Something to Chance." *Philosophy and Public Affairs* 18(1): 53–67.
- Littlejohn, Clayton. Forthcoming. "A Plea for Epistemic Excuses." In *The New Evil Demon Problem*, ed. Fabian Dorsch and Julien Dutant. Oxford: Oxford University Press.
- Madison, B. J. C. 2018. "On Justifications and Excuse." *Synthese* 195(10): 4551–62.
- Robinson, Michael. 2012. "Modified Frankfurt-Type Examples and the Flickers of Freedom." *Philosophical Studies* 157: 177–94.
- . 2014. "The Limits of Limited-Blockage Frankfurt-Style Cases." *Philosophical Studies* 169: 429–46.
- . Forthcoming. "Robust Flickers of Freedom." *Social Philosophy and Policy*.
- Rosen, Gideon. 2015. "The Alethic Conception of Moral Responsibility." In *The Nature of Moral Responsibility: New Essays*, ed. Randolph Clarke, Michael McKenna and Angela M. Smith. Oxford: Oxford University Press.
- Sartorio, Carolina. 2019 (this issue). "Flickers of Freedom and Moral Luck." *Midwest Studies in Philosophy*.
- Statman, Daniel. 2015. "Moral Luck and the Problem of the Innocent Attacker." *Ratio* 28(1): 97–111.
- Swenson, Philip. 2015. "A Challenge for Frankfurt-Style Compatibilists." *Philosophical Studies* 172: 1279–85.

17. Thanks to audiences at VCU and FLAW. Thanks also to Justin Coates, Neal Tognazzini, and the Board Certified Epistemologists for helpful discussion of some of these issues. Thanks to Andrew Khoury for helpful comments. And thanks to Nevin Climenhaga and Jonathan Trerise for (between the two of them) coming up with the title. Work on this paper was supported by summer funding from the College of William & Mary.

- Swenson, Philip. 2016a. "The Frankfurt Cases and Responsibility for Omissions." *Philosophical Quarterly* 66(264): 579–95.
- . 2016b. "Ability, Foreknowledge, and Explanatory Dependence." *Australasian Journal of Philosophy* 94(4): 658–71.
- Williamson, Timothy. Forthcoming. "Justifications Excuses, and Sceptical Scenarios." In *The New Evil Demon Problem*, ed. Fabian Dorsch and Julien Dutant. Oxford: Oxford University Press.
- Zagzebski, Linda. 2000. "Does Libertarian Freedom Require Alternative possibilities?" *Philosophical Perspectives* 14: 231–48.
- Zimmerman, Michael. 2002. "Taking Luck Seriously." *Journal of Philosophy* 99: 553–76.