Erkenntnis
https://doi.org/10.1007/s10670-024-00818-3

ORIGINAL RESEARCH

Blameworthiness and Causal Outcomes

Matthew Talbert[1]

Received: 27 September 2023 / Accepted: 13 April 2024

Abstract
It is widely held that whether a person is morally responsible for an outcome partly depends on whether certain causal relations obtain between that person and the outcome. This paper argues that, regardless of whether the preceding claim about moral responsibility is true, moral blameworthiness is independent of such causal considerations. This conclusion is motivated by considering cases from Carolina Sartorio and Sara Bernstein. The causal structures of these cases are complex. Sartorio and Bernstein believe that reaching conclusions about moral responsibility and blameworthiness requires wading through this complexity. But it is argued here that getting clear on these causal structures does not help us assess blameworthiness (regardless of the relevance of such considerations for drawing conclusions about moral responsibility). Moral blameworthiness simply does not turn on the considerations relevant to figuring out whether an agent causally contributed to an outcome.

Keywords Blame · Causation · Moral Affront · Moral Responsibility · Resultant Moral Luck

1 Introduction

Regardless of whether we are morally responsible for the casual outcomes of our behavior, I argue here that such outcomes do not affect moral blameworthiness. Neither the occurrence of an outcome, nor the fact that it resulted from a choice of ours, affect our openness to moral blame.

The above claim is related to a standard position in the debate about moral luck. Skeptics about resultant moral luck hold that the fortuitous occurrence of a negative outcome does not affect an agent's blameworthiness. Defenses of this position—e.g., Zimmerman (2002)—often focus on our lack of control over the consequences of our actions, or on the unfairness of blaming a negligent actor who causes harm while excusing an equally negligent actor who luckily avoids causing harm. My

✉ Matthew Talbert
matthew.talbert@mail.wvu.edu

[1] Department of Philosophy, West Virginia University, Morgantown, WV, USA

Published online: 08 May 2024

argument is motivated by different considerations. As I explain in Section 2, the central idea is that outcomes and the obtaining of causal relationships between actions and outcomes are simply not the kinds of things that make a person worthy of moral blame. In Section 3, I consider causally complex cases of agency that illustrate this claim: the lesson from these cases is that solving causal puzzles is irrelevant for assessing blameworthiness. In Sections 4 and 5, I consider responses to the view that I defend.

2 Responsibility and Blameworthiness for Outcomes

Philosophers working on moral responsibility often assume that we can be responsible[1] not just for our actions but also for the causal outcomes of our actions. A natural way to think about responsibility for outcomes is in terms of causation and derivative responsibility: responsibility for outcomes derives from responsibility for actions that are appropriately causally related to these outcomes. However, the obtaining of such causal relations is not sufficient for moral responsibility: if I act in ignorance of the likely effects of my action, I may not be responsible for these effects even if they are causally related to an action for which I am responsible. So, there seems to be an epistemic condition on responsibility for outcomes: an agent must know (or should have known) that their action would (or that it might) cause the outcome in question. Along these lines, Carolina Sartorio proposes the *Principle of Derivative Responsibility*:

> "If an agent is responsible for X, X causes Y, and the relevant epistemic conditions for responsibility obtain, then the agent is also responsible for Y."[2] (2016, 76)

I take this to be a plausible principle that many philosophers would endorse.

Perhaps many will also agree with Sartorio that agents can be blameworthy for the causal consequences of their actions (2016, 77). A prominent view in the moral responsibility literature is that blameworthiness is simply, as Sartorio puts it, "a form of responsibility" (2016, 77), and it is also widely held that being morally responsible comes down to deserving blame (or praise) (2016, 7). Given such a connection between responsibility and blameworthiness, blameworthiness for outcomes seems plausible. Here is one way to think about this: if we fulfill the causal and epistemic conditions on moral responsibility with respect to some outcome (a harm, for example), and we do so in a way that is morally objectionable (for example, by intentionally and maliciously causing the harm), then we are blameworthy for that

[1] I use "moral responsibility" interchangeably with "responsibility," and "moral blame" interchangeably with "blame."

[2] Sartorio suggests that we may need to add a condition that rules out responsibility in cases of deviant causation (2016, 76, note 32). I discuss cases of deviant causation in the next section.

outcome. The basic idea is that you are blameworthy for an outcome if you are morally responsible for it in virtue of factors that ground moral blame.[3]

It's a version of this last claim that I argue against in this paper. On my view, there may be room for saying that we can be morally responsible for outcomes, but outcomes do not affect our worthiness of moral blame. I will argue that some of the very cases that Sartorio considers, along with others from Sara Bernstein, suggest this conclusion. Through their examinations of causally puzzling scenarios, Sartorio and Bernstein have done valuable work showing how judgments about moral responsibility are informed by causal judgments. But contrary to Sartorio and Bernstein, I think that these same puzzles suggest that our judgments about blameworthiness should not be so informed. (Whether my argument also tells against responsibility for outcomes will depend on how readers take responsibility and blameworthiness to be related; I take this issue up below.)

I'll turn to Sartorio's and Bernstein's causal puzzles in the next section; in this section, I want to motivate my central claim more abstractly. First, we should examine the notion of "blameworthy for." Given what I've said so far, agent *A* is blameworthy for outcome *O* if *A* fulfilled the causal and epistemic conditions on moral responsibility (with respect to *O*) in virtue of factors that make blame appropriate. This might be so, for example, if the occurrence of *O* is explained (in the right way) by *A*'s morally objectionable motives and intentions.

But this basic account leaves room for stronger and weaker readings of "*A* is blameworthy for *O*." On the weaker reading, *A* is blameworthy for *O* if *A* caused *O* in virtue of factors that *independently and exclusively* make blame appropriate, such as *A*'s objectionable motives and intentions. I have no objection to this weak reading of "*A* is blameworthy for *O*" because, on this reading, the occurrence of *O*, and *A*'s connection to *O*, do not affect *A*'s openness to blame: A's blameworthiness is fixed just by *A*'s objectionable motives and intentions. What I object to is a stronger reading of "*A* is blameworthy for *O*" according to which the occurrence of *O,* and *A*'s connection to it, make *A more* worthy of blame (or worthy of *more* blame) than *A* would have been had *O* not occurred.[4]

Again, on my view neither the occurrence of an outcome nor an agent's causal connection to that outcome affect the agent's blameworthiness. Given the close relationship often assumed between blameworthiness and responsibility, it may seem to follow that we also are not morally responsible for outcomes. Thus, the skeptic

[3] This is similar to Sartorio's *Principle of Derivative Blameworthiness* (2016, 77). For related approaches, see Björnsson (2017) and Gunnemyr and Touborg (2023).

[4] See Graham (2014, 407), Smith (1983, 566-68), and Thomson (1989, 208-10), for discussion related to the distinction made here. Below, I argue that we can retain talk of responsibility for outcomes while dismissing blameworthiness for outcomes. A referee suggests that I could apply the weak/strong distinction just mentioned to moral responsibility, thus preserving the standard connection between responsibility and blameworthiness. I could say that we are neither responsible nor blameworthy for outcomes in a strong sense that entails openness to blaming responses, though we may be weakly responsible and blameworthy for these things in a way that does not have this entailment. Though I don't pursue this suggestion in the text, it is reasonable; indeed, the responsibility I describe below as potentially applying to outcomes can be said to be "weak" precisely because it does not entail blameworthiness in the strong sense described in the text.

about blameworthiness for outcomes may accept something like Khoury's (2018) proposal that we can be morally responsible only for our internal "willings," and never for the consequences of our actions. But I will suggest a way of accommodating responsibility for outcomes while still dismissing, with Khoury, the possibility of blameworthiness for outcomes.

Moral responsibility is certainly connected to praiseworthiness and blameworthiness, but we need not conceive of responsibility as entailing openness to either praise or blame.[5] Rather, we can think of moral responsibility as fundamentally involving a sort of control—characterized by satisfaction of causal and epistemic conditions—such that an agent *would be* praiseworthy or blameworthy *if* the thing over which they exercised control reflected morally well or poorly on them. On this view, it wouldn't follow that when we are morally responsible we are also either praiseworthy or blameworthy, for we can exercise responsibility-relevant control over behaviors that are morally neutral, grounding neither praise nor blame.[6]

If we allow for moral responsibility for morally neutral actions (as I think we should), then we should also regard morally neutral consequences as falling within the scope of moral responsibility when agents exercise the right sort of control over these consequences—e.g., by intentionally or knowingly contributing to the occurrence of those consequences. But on the view I offer below, all outcomes are morally neutral in the relevant sense because they never in themselves reflect well or poorly on anyone. So, we would never be praiseworthy or blameworthy for outcomes even when we exercise the sort of control sufficient for moral responsibility over them.

Of course, on this picture, attributions of responsibility for outcomes will tend to be of little practical significance (and the same is typically true of responsibility for morally neutral behaviors).[7] No doubt there is something counterintuitive about speaking of "moral responsibility" in such contexts. But I think the move suggested here has some appeal as a way of regimenting our conception of moral responsibility, at least if the core of that conception has to do with exercising a certain sort of control, which is a separate matter from whether the thing over which we

[5] I take Khoury to accept such an entailment. See his suggestion that responsibility does not entail blameworthiness only because it may instead entail praiseworthiness (2018, 1358).

[6] The view that morally neutral acts can fall within the scope of our moral responsibility has proponents: Fischer and Ravizza (1998, 8 note 11), Hieronymi (2008, 363 note 13), Talbert (2016, 1–2). McKenna suggests that we can be responsible for morally neutral actions, but only if they occur in morally salient contexts (2012, 17). Zimmerman defends a related view (1988, 61–2).

[7] But such attributions may not be *entirely* without significance. If I cause a bad outcome and satisfy relevant epistemic conditions with respect to it, then the outcome is my doing in a way that—while not making blame appropriate—could give rise to certain obligations. For example, perhaps I have a particular duty to compensate a person if I have knowingly caused them harm. (See note 20 for more on this.) Khoury allows that outcome-causing agents may acquire such obligations, but he notes that we need not posit responsibility for outcomes to explain this (2018, 1375-76). Indeed, Khoury can't posit responsibility here and preserve the (correct) conclusion that we are not blameworthy for outcomes, since he understands responsibility "as the extent to which an agent is blameworthy or praiseworthy" (2018, 1375). But if we understand responsibility as a function of control, the exercise of which need not entail praiseworthiness or blameworthiness, then there is room for conceiving of holding a person to the obligations in question—if they arise—as a way of holding them responsible for what they have caused.

exercise control redounds to our moral credit or discredit.[8] However, if readers disagree, they may simply set aside what I have said here. I am much more concerned with establishing that we are not blameworthy for outcomes than that we are morally responsible for outcomes, and my arguments below regarding blameworthiness do not depend on accepting the proposal that I have made about the use of "moral responsibility."

I'll turn, then, to explaining why I find it implausible that the obtaining of causal relationships between actions and outcomes affects blameworthiness. Three related thoughts motivate my perspective. The first is that to be blameworthy is to be worthy of blame: it is to be an appropriate target of the morally angry and offended attitudes—particularly resentment—involved in blame. The second thought is that these attitudes are made appropriate by instances of moral affront, and the third is that an unwelcome outcome is never in itself a moral affront, and neither is the fact that an agent caused such an outcome.[9]

Clearly, bad things can happen without anyone having grounds for being morally offended. It's less obvious that an agent's having caused a bad outcome is not itself direct grounds for moral offense.[10] Still, there's good reason to accept this. The central thing to note is that even if an agent's action causes a bad outcome, and even if the agent foresaw this outcome, we generally do not take ourselves to have grounds for moral offense unless we believe certain things about *why* an agent

[8] I take this to be compatible with the attributionist claim that agents can be open to blame on account of things over which they lack control, such as their beliefs, desires, and moral character. However, what I have said may not fit well with the attributionist claim that agents can be morally responsible for things over which they lack control. For an account of attributionism, see Talbert (2022). A referee suggests that I note the connection between what I say in this paper and attributionist accounts. I am happy to do so. I have not presented my argument here as an attributionist one only because I don't want to give the impression that it depends on accepting controversial features of that view; again, see Talbert (2022).

[9] The "moral affront" that I take to ground blaming responses like resentment is, roughly, a failure to show others due moral regard. These failures make resentment appropriate because they make it fitting; they do so because, I shall assume, resentment in some way represents actions as conveying (something like) lack of regard. Here I follow Graham (2014), Rosen (2015), and Strabbing (2019). They take resentment to be partly constituted by the thought that another "has violated a moral requirement of respect" (Graham 2014, 408), "showed an objectionable pattern of concern" (Rosen 2015, 77), or "expressed insufficient good will" (Strabbing 2019, 3127). On these accounts, resentment is fitting only if one or another of the thoughts just mentioned is true. (Also see D'Arms and Jacobson's claim that resentment is a "cognitive sharpening" of anger—which itself has no cognitive content—constituted by the thought that one was wronged [2003, 143].) For my disagreement with Rosen and Strabbing about the nature of the cognitive content of resentment, see Talbert (2023). A referee suggests that any action that causes anger or outrage may be a moral affront. In this case, actions that give rise to unwelcome consequences, even in the absence of moral disregard, may qualify as moral affronts. But I do not think of simple anger as a moral blaming response; at any rate, it is not the sort of blaming response that I am arguing is inappropriate in the context of outcomes. Rather, my focus is on the moralized anger involved in resentment, which, for the reasons given, I take to be made fitting only by moral disregard. If someone is angry because of an unwelcome outcome even though they judge that they have not been shown disregard, then I do not think it is resentment and moral blame that they are experiencing (cf. D'Arms and Jacobson [2003, 143])—or if it is, then it is experienced inappropriately. See note 30 for more on this last point.

[10] The occurrence of a bad outcome may give us reason to suspect a moral affront, but I am arguing that neither the bad outcome, nor an agent's connection to it, is itself a moral affront or in itself grounds for moral offense.

acted as they did. Were they moved by considerations that we ourselves take to justify their action? Do we therefore find the agent's motives and intentions, and their patterns of moral regard, unobjectionable? If so, then by our own lights there is no moral affront, and we have no grounds for moral offense and offended blaming responses like resentment.[11] Or do we find the agent's motives and intentions to be in themselves morally offensive: for example, do we find them to signal a morally offensive disregard for others and the moral objections that they might be raise to certain forms of treatment? If so, then, we have grounds for moral offense and for the morally offended responses involved in blame.

So, unwelcome outcomes—even when knowingly brought about—do not ground the moral offense involved in blame because they are not in themselves instances of moral disregard (though such disregard often gives rise to unwelcome outcomes). Neither are such outcomes necessary for constituting a moral offense. An agent's motives and intentions can be morally objectionable even if they don't give rise to bad outcomes.[12] Finally, the occurrence of a bad outcome, or an agent's being causally connected to it, does not plausibly add to the moral offensiveness of an agent's objectionable motives and intentions.[13] The offensiveness of an agent's motives and intentions is fully explained by what they indicate about an agent's regard for others or for important moral values, and the occurrence of a bad outcome doesn't affect this. A bad outcome may make a person's objectionable motives and intentions—the things that legitimate our blaming responses—more salient than they would have been in the absence of the outcome, but this does not affect an agent's blameworthiness, only our readiness to blame them.[14]

Blameworthiness, then, is not affected by outcomes, or by bearing a causal relation to an outcome, because these things are not in themselves moral affronts. But motives and intentions can be in themselves moral affronts. And we have a reasonably straightforward story to tell here. Suppose that a person is motivated by unjustified malice to injure me. The injury is unwelcome, but it is the malice and the aim to injure that morally offends (and that grounds blame[15]) because these things involve a failure to show me the moral regard to which I am entitled. Failures of regard constitute moral affronts that make the moral offense expressed in blame appropriate. What is the analogous story that explains how the addition of a bad outcome (to

[11] We have no such grounds because, by our own lights, there is no disregard for moral standing, so no moral affront, in the vicinity. I find this reasoning very plausible, but it can be questioned: see Sher (2009), Clarke (2014), and Ayars (2021).

[12] Those who think that outcomes matter for blameworthiness typically allow that intentions have independent moral significance. Bernstein thinks that a lucky agent who injures someone may be "more morally responsible" than an agent who injures no one, but she allows that the second agent may still be blameworthy for their intentions (2019, 152). For a similar admission, see Sartorio (2012, 22).

[13] See Lang (2021) for disagreement. I discuss Lang's view in Section 5.

[14] Skeptics about resultant moral luck often make this sort of point. See, e.g., Richards (1986). For development of the idea, from Scanlon (2008), that the significance of an outcome can affect the appropriateness of blame even though it does not affect an agent's blameworthiness, see Riedener (2021). For criticism of Scanlon on this point, see Talbert (2019), which defends an attributionist approach to moral luck; for objections to the attributionist approach, see Miller (2022).

[15] Thus, as Strawson (1962) notes, a central form of moral excuse involves showing that a harm was not maliciously caused.

a malicious will) properly elicits morally offended blaming responses (beyond the responses already made proper by the attribution of malice)? How are outcomes in themselves instances of disregard or something similar, and if they are not, how do they in themselves constitute moral affronts? I have not seen compelling answers to these or similar questions from those who defend the proposition that outcomes in themselves make a difference to the appropriateness of blaming responses (though I consider contenders in Section 5).

In the next section, I introduce cases from Sartorio and Bernstein that illustrate the points made above. The causal structures of these cases are complex and Sartorio and Bernstein believe that reaching conclusions about moral responsibility requires wading through this complexity. As I have suggested above, perhaps this is true (though in a way that makes an attribution of moral responsibility less morally interesting than we might have expected). But unlike Sartorio and Bernstein, I don't think that getting clear on the causal structures of these cases helps us draw conclusions about moral blameworthiness. An agent's openness to blame simply doesn't turn on the considerations that go into assessing whether they causally contributed to an outcome.

3 Causation Cases

3.1 How Morally Important is Moral Luck?

I'll start with a relatively simple example, Carolina Sartorio's version of *Sharks* (borrowed from Fischer & Ravizza, 1998):

> … Bad Samaritan sees a child drowning. He thinks he could easily jump into the water and save him but decides not to do so. The child dies. Unbeknownst to Bad Samaritan, the water is infested by sharks. Had he jumped in, the sharks would have attacked him and prevented him from saving the child. (2012, 70)

Sharks illustrates a form of resultant moral luck.[16] Sartorio says that "Bad Samaritan is lucky that the water was infested by sharks" since this fact "seems to relieve him of at least some moral responsibility"; in particular, Bad Samaritan "is not responsible for his failure to save … [the child], or for the child's death" (2012, 70–1). In addition, agents in contexts like Bad Samaritan's "are less blameworthy because their behavior didn't have harm as a consequence" (2012, 72).

I agree that Bad Samaritan is not morally responsible for the child's death since the death is not explained by Bad Samaritan's failure to act. But if this is a case of moral luck, it is a relatively uninteresting variety. This is because Bad Samaritan's being relieved of moral responsibility seems, *pace* Sartorio, irrelevant to our moral assessment of him and of his moral blameworthiness. In discovering the fact that relieves Bad Samaritan of his moral responsibility—the presence of the sharks—we do not discover anything that casts his failure to act in a morally better light. Thus,

[16] Sartorio is agnostic about resultant luck in this paper. Her aim is to get "a grip on the kind of thing to which we would be committed if we were to be committed to resultant luck" (2012, 64).

we have no reason to revise the morally offended blaming responses that we might have had toward Bad Samaritan if we had not known about the sharks.[17] Learning about the sharks may tell us something important about the causal structure of the case, but it tells us nothing about Bad Samaritan's openness to blame.

3.2 Degrees of Causation and Responsibility

Consider a pair of cases from Sara Bernstein. *Victim* is a case of overdetermination: two assassins shoot Victim, and he dies, but one shot would have killed Victim; *Hardy Victim* is a case of joint causation: two assassins shoot Hardy Victim, and he dies, but Hardy Victim is hard to kill and both shots were necessary to kill him (Bernstein, 2017, 165).

Bernstein asks whether these cases "differ morally": "does each assassin's proportion of moral responsibility for Victim's death quantitatively differ between Victim and Hardy Victim?" (2017, 165). As part of the answer to this question, Bernstein proposes a principle that she calls Proportionality: "An agent's moral responsibility for an outcome is proportionate to her actual causal contribution to the outcome" (2017, 167).[18] If we accept Proportionality, we should think that the assassins who make a greater causal contribution to Victim's death bear greater moral responsibility than those who make a smaller causal contribution.

But as Bernstein observes, determining which set of assassins makes a greater causal contribution will depend on how we conceive of causation. Suppose we employ a "'productive' … theory of causation" according to which "causation is a matter of energy transference between a cause and an effect" (Bernstein, 2017, 170). It takes both assassins in *Hardy Victim* to transfer the energy required to kill the victim, but in *Victim*, both assassins individually transfer the energy required. We might conclude, then, that each assassin in *Hardy Victim* is only "partially responsible for the death," but "in the overdetermination case [*Victim*], both assassins are, in a sense, 'fully' responsible for Victim's death" (Bernstein, 2017, 170–71). So, we should conclude that each assassin in *Victim* bears greater moral responsibility than each assassin in *Hardy Victim*.

Alternatively, we might have "counterfactual causal intuitions" and focus on the fact that in *Hardy Victim* the contribution of each assassin is "more essential to the victim's death" than in *Victim* (Bernstein, 2017, 171). If either assassin in *Hardy Victim* had refrained, the victim would have survived, but in *Victim*, either assassin could have refrained, and the victim would still have been killed. Since in *Hardy Victim* the death depends on both assassins' contributions, these contributions may seem more causally significant than those of the assassins in *Victim*. This could lead us to conclude "that each assassin in (Hardy Victim) is *more*, rather than *less*, causally responsible than each assassin in (Victim)" (Bernstein, 2017, 171, emphasis in original).

If different conceptions of causation yield different conclusions about comparative causal contributions in *Victim* and *Hardy Victim*, perhaps this should be

[17] This conclusion depends, of course, on the account of appropriate moral offense in the preceding section.

[18] This suggests a form of resultant moral luck—"proportionality luck" (Bernstein 2017, 168)—that occurs when an agent's proportion of moral responsibility for an outcome is a matter of luck because their degree of causal contribution is out of their control.

reflected in our judgments about moral responsibility. But I would reject a similar conclusion about blameworthiness.[19] Suppose that I come to find the productive account of causation more compelling than I once did. In coming to this conclusion, I may have reason to adjust my causal judgments about *Victim* and *Hardy Victim*, but I won't have discovered anything that should alter my judgment about the comparative blameworthiness—the openness to blaming responses—of the pairs of assassins in these cases. And this is because I won't have discovered anything that should alter my sense of the degree to which these examples involve moral affronts.

Or perhaps my causal intuitions about *Victim* and *Hardy Victim* are simply unclear because I waiver between the two pictures of causation that Bernstein considers. But suppose that I am certain about everything else that could be morally relevant in the two cases: I know about the assassins' attitudes toward their occupation and toward the consequences of engaging it, I know how these attitudes informed the assassins' intentions and decisions, and so on. *If* I knew all this, isn't it strange to suppose that I must decide between theories of causation before I can conclude whether blaming attitudes are fittingly held toward the assassins? Getting the causal facts straight doesn't seem relevant here. Even without being certain about the causal picture, I may already know everything required to make a judgment about the moral responses appropriate in *Victim* and *Hardy Victim*.

I've just suggested that adjustments to, or ambiguity in, our causal intuitions won't necessarily yield corresponding adjustments to, or confusion in, our judgments about blameworthiness.[20] But suppose that we have stable and clear views about causation. Suppose, for example, that we have settled counterfactual casual

[19] I assume that Bernstein sees differences in causal contributions and moral responsibility as entailing differences in blameworthiness because she takes responsibility to track blameworthiness (2017, 166), and because she says that proportionality luck affects blameworthiness (2017, 168). Also, Bernstein suggests that the view that there is no moral difference between *Victim* and *Hardy Victim* is motivated by the idea "that the way things turn out is irrelevant to the blameworthiness of agents" (2017, 167). Bernstein rejects the *no difference* claim partly, I assume, because she thinks that the way things turn out *does* affect blameworthiness.

[20] At several points in this section, referees press the following sort of objection. I say that shifts in our causal judgments shouldn't lead to changes in our judgments about blameworthiness. But I have also suggested that such shifts might lead to changes in our judgments about moral responsibility. Is this plausible? I suggested in Section 2 that bearing the right causal and epistemic relations to an outcome might mean that one is morally responsible for that outcome (without this entailing anything about blameworthiness). This way of defining the scope of moral responsibility has the virtue of accommodating the central role that fulfillment of causal and epistemic conditions plays in accounts of moral responsibility. But what of the suggestion (in note 7) that responsibility for outcomes may give rise to obligations? In an earlier draft of this paper, I suggested that bearing causal relations to outcomes might generate duties of apology. The referees' comments convince me that this is not plausible: if shifts in causal judgments shouldn't affect blameworthiness, why should they affect who should apologize or how much they should apologize? I have amended the presentation in Section 2 accordingly. But I have left in place the suggestion that a person who causes an outcome may acquire duties of compensation. I still think it reasonable that if accepting a theory of causation leads us to conclude that it was *A*'s malice rather than *B*'s that led to *C*'s harm, then, in choosing where to place the burden of making *C* whole, we should place it on *A* rather than *B* (even if it is only a matter of luck that *B*'s malice did not play the causal role that *A*'s malice did). We don't have to attribute responsibility to *A* to explain why they should bear this burden, but putting this burden on them is at least compatible with attributing the sort of moral responsibility for outcomes described in Section 2.

intuitions. Recall that in this case we should, on Bernstein's view, see each assassin in *Victim* as bearing less moral responsibility than each assassin in *Hardy Victim*. This is because—on the counterfactual view—each assassin in *Hardy Victim* makes a more significant causal contribution to Victim's death because each assassin's contribution was necessary for that outcome.

Some might "resist the intuition that agents in a case of overdetermination [like *Victim*] are less responsible than agents in a case of joint causation [like *Hardy Victim*]" (Bernstein, 2017, 172 note 6). To motivate this intuition, Bernstein proposes a case in which pushing any of 100 buttons is sufficient to kill Victim: "100 assassins push their buttons, and Victim dies…. Victim's death is overdetermined 99 times over" (Bernstein, 2017, 172 note 6). Bernstein expects readers to share "the intuition that the sheer number of assassins involved lessens the moral responsibility of each one" (2017, 172 note 6).

Call this new case *100 Assassins*. Assuming the truth of the counterfactual view, if there were 101 assassins, then each would be slightly less morally responsible, and if there were 99 assassins, then each would be slightly more responsible. The general lesson is that if degree of moral responsibility tracks degree of causation, then moral responsibility may increase or decrease depending on how many other causes are in play.

We should see a corresponding effect if we adopt a productive view of causation. In that case, adding extra overdetermining causal agents, like the assassins in *Victim* or in *100 Assassins*, won't affect the moral responsibility of each individual agent: each agent's contribution would still be just as causally significant (in the productive sense) as any other agent's because each agent's contribution would be sufficient for the relevant outcome. But adding additional jointly causing agents, like the assassins in *Hardy Victim*, would reduce the moral responsibility of each agent because each agent would contribute less (from the productive standpoint) to the outcome.

We see this effect in Carolina Sartorio's treatment of a set of cases. Similar to Bernstein's Proportionality principle, Sartorio says, "it seems natural to expect that … to the extent that we can make a *more significant* or *less significant* contribution to an outcome's occurrence, this will also result in different degrees of responsibility for that outcome" (2015a, 140, emphasis in original).[21] Here are Sartorio's cases:

> **CASE 1**: I want an explosion E to occur. I have good reason to believe that pressing button A will trigger an explosive that will result in E. I press A, and E occurs. (Sartorio, 2015a, 140)
> **CASE 2**: *Three* buttons (A, B, and C) need to be pressed for E to occur. Two other agents independently press B and C while I press A …. E occurs. (Sartorio, 2015a, 140, emphasis in original)
> **CASE 3**: Again, there are three buttons and three fully informed and responsible agents. But this time pressing *any* of the buttons would be sufficient to bring about E. I press A, and the two other agents press B and C. E occurs. (Sartorio, 2015a, 141, emphasis in original)

[21] At the end of the paper, Sartorio expresses reservations about the proposal that causation can come in degrees; also see Sartorio (2020). For defense of the proposal that contributions to the causing of an outcome can come in degrees (and application in the context of tort law), see Kaiserman (2017); for an overview of the debate about degrees of causation, see Kaiserman (2018).

Sartorio says, "given that I seem to make a significantly smaller contribution in CASE 2 than in CASE 1, this supports the idea that I'm responsible to a *significantly lesser* degree in that case" (2015a, 141, emphasis in original). She adds: "if it helps pump the relevant intuitions, we can always increase the number of agents and buttons to, say, one thousand" (2015a, 141). The thought is that in a case with 1000 button pushers—1000 joint causal actors—it should be obvious that Sartorio's own contribution is vastly smaller than in *CASE 1* where she is the sole button pusher. This is the same move that Bernstein makes with *100 Assassins*, but it is made in the opposite causal context.

When Sartorio turns to *CASE 3*, an overdetermination case, she suggests that the causal and moral responsibility of each button pusher is the same as that of the individual button pusher in *CASE 1* (2015a, 141). So, Sartorio sees the overdetermination case (and not the joint causation case) as the one involving comparatively greater individual moral responsibility. Perhaps this illustrates Bernstein's claim that different causal intuitions will lead to different comparative assessments of moral responsibility when assessing cases like *Victim* and *Hardy Victim*. But for my purposes it doesn't matter if it is Bernstein's *100 Assassins* or Sartorio's *CASE 3* in which adding causal agents reduces moral responsibility. What's important is that they both accept that adjusting the number of causal contributors to an outcome can affect the moral responsibility of each individual contributor.

Like Bernstein, Sartorio recognizes that there's a kind of moral luck in play here.[22] If adding causes reduces responsibility, and you lack control over how many causal factors are present, then your moral responsibility depends on something over which you lack control. This form of moral luck arises insofar as (i) causation is the sort of thing that can be divided up such that you may contribute more or less of it, and (ii) your moral responsibility for an outcome depends on your causal contribution to it.

Regardless of what we say about moral responsibility, it seems clear to me that this sort of moral luck does not affect blameworthiness. (So, as in Section 3.1, this is a less morally interesting form of moral luck than we might have thought.) Perhaps causal contributions can be divided up amongst individuals, but grounds for blame are not divisible in this way.[23] In particular, the moral offensiveness of an agent's contributing to an unwelcome outcome is not a direct function of the number of other agents contributing to an outcome. Rather, the offensiveness of an individual's action, and the blame they merit, depends on their individual motives and intentions. Once these are fixed, the blameworthiness of the individual assassin is also fixed even if their actual causal contribution depends on the size and causal structure of the group of which they are a member.[24]

[22] This form of moral luck is "illustrated by cases where our responsibility appears to depend exclusively on whether *other responsible agents* are present and what their contributions are" (Sartorio 2015a, 134, emphasis in original).

[23] See Kaiserman (2021) for defense of the related claim that responsibility for an outcome is not something that is divided up among contributing individuals.

[24] For additional concerns about the difficulties that proportional responsibility raises for resultant moral luck, see Demirtas (2022a, b).

3.3 Deviant Causation

In Bernstein's *Angry Cassowary* case, "Assassin shoots at Victim, intending to kill him, but the shot misses. However, the shot startles a sleeping cassowary who then angrily mauls Victim to death" (2019, 151). This case features deviant causation: Assassin's general aim is achieved through Assassin's efforts, but the route to Victim's death deviates significantly from what Assassin intended. According to Bernstein, "the prevailing intuition is that Assassin is not fully morally culpable for Victim's death, given that it happened in such a strange way" (2019, 152).[25]

Perhaps the strangeness of the causal route in *Angry Cassowary* means that Assassin does not cause Victim's death in the right way to support moral responsibility. But I reject the idea that this causal deviance would affect Assassin's moral blameworthiness in any meaningful way. Even if Assassin had killed Victim in the way he intended, he wouldn't be open to a fundamentally different sort or degree of moral blame than he is in Bernstein's version of the case.

Now consider *Ricochet*: "Jayantha shoots at Jordan, intending to kill him. The bullet ricochets off of a nearby rock and kills Jordan" (Bernstein, 2019, 155). Is the causal deviance here sufficient to undermine responsibility? Bernstein suggests that it is not: "the natural intuition is that Jayantha is still morally responsible for Jordan's death, since the causal process is not so different than the one Jayantha intended" (2019 155).

But as Bernstein notes, it's not clear how deviant a causal process must be to call responsibility into question (2019, 155). I take it that reasonable people can disagree about whether the fit between Jayantha's intentions and Jordan's death are of the right sort to support moral responsibility. I suggest, though, that parties to this disagreement need not also disagree about blameworthiness: they need not see themselves as having grounds for disagreeing about the morally offended attitudes that might be appropriately directed toward Jayantha.

Similarly, if we revise our judgment about the presence of causal deviance in a case, this won't necessarily give us grounds for revising our blaming responses. In Sartorio's *Choking with Water*, a husband aims to kill his wife: he offers her a glass containing only water, but he believes that it contains a lethal poison. As it happens, the wife chokes on the water and dies. Sartorio says that even though "the evil husband does cause his wife's death," because of deviance in the causal sequence, "we still wouldn't hold him morally responsible for her death (only for his evil intentions, or for trying to kill his wife)" (2012, 82). But what if we learned that the husband intended for his wife to die in the unlikely way that she did? Or suppose we learn, contrary to what we initially thought, that Assassin in *Angry Cassowary* intended to startle the cassowary and to bring about Victim's death via the cassowary's attack. (Perhaps Assassin wanted to conceal their involvement, and struck on the plan involving the cassowary.) In these revised versions of the cases, perhaps we should

[25] Bernstein mentions *moral culpability* here; elsewhere in the paper, she puts the point in terms of moral responsibility. I take these phrases to come to the same thing and to entail, in cases of objectionable behavior, conclusions about blameworthiness since Bernstein assumes that moral responsibility "amounts to blameworthiness and praiseworthiness" (2019, 152).

revise our initial responsibility assessments. Perhaps the husband and the assassin are no longer responsible just for evil intentions and attempts, but also for killings. But it is not plausible that we have learned something that shows these agents to be more fitting targets of moral blame than in the original versions of the cases. This is because the new information about these cases does nothing to suggest that they involve greater instances of moral disregard than we had initially supposed.

Now consider Sara Bernstein's version of *Thirsty Traveler* (2019, 158).[26] Victim is about to travel into the desert with a canteen of water. Aiming to kill Victim, Billy empties the canteen and fills it with a poison that kills via dehydration. Later, when Victim is in the desert, Suzy steals the canteen with the aim that Victim die from dehydration. Victim dies from dehydration. Who caused Victim's death? It's difficult to answer this question because Billy's and Suzy's efforts to kill Victim interfere with one another. Suzy's theft of the canteen prevents Billy's poison from having its effect. But since Suzy merely deprived Victim of something that would have killed him, how could she have caused his death?

Thirsty Traveler involves deviant causation since Victim's death "does not occur via the causal process planned by Billy or Suzy" (Bernstein, 2019, 159). This partly explains our reluctance to assign full responsibility for Victim's death to either Billy or Suzy. Complicating things further is the fact that Suzy acts as both a *preemptor* and a *delayer*. Suzy is a preemptor in that "stealing the canteen preempts Billy's poison from killing Victim" (Bernstein, 2019, 159). Preemption often signals causation. To use Bernstein's example, if Billy and Suzy throw rocks at a window and Billy's rock reaches the window first, then Billy's throw "is the preempting cause" of the window's shattering (2019, 159). This might incline us to think that Suzy caused Victim's death. But Suzy also delays Victim's death since "if she had not stolen the canteen, Billy's poison would have killed Victim sooner" (2019, 160). Delayers, Bernstein says, are typically not taken to be causes (2019, 157). This leads us to the contrary conclusion that Suzy does not cause Victim's death but "merely changed a death that was already going to occur by making it later" (Bernstein, 2019, 160–61).

But now imagine a version of the case in which Suzy slightly hastens rather than delays Victim's death. That is, imagine that because Suzy stole Victim's canteen, Victim dies sooner of dehydration than if Suzy had left Billy's poison to do its work. Perhaps the poison that Billy used, though ultimately causing death by dehydration, also has an initial hydrating property such that Victim would have survived longer—perhaps only a few seconds longer—if Suzy had not stolen his canteen.

Is the causal picture clearer in this version of the case? I think so. Suzy preempts Billy's effort to kill Victim, and her theft of the canteen leads to Victim dying (by dehydration, as Suzy intended) sooner than would have otherwise been the case. Perhaps this should affect our judgement about what Suzy caused (and what she is responsible for) but it is not plausible to expect a corresponding effect on Suzy's blameworthiness. It is not plausible that Suzy's openness to the morally offended responses involved in blame partly depends on whether Billy used a slower- or

[26] See Sartorio (2015a, b) for her take on this puzzle. For a reply to Sartorio related to the one in the text, see Talbert (2015).

faster-acting poison. Learning what poison Billy used might give us insight into the causal structure of the case, but it doesn't give us any additional insight into the propriety of moral blame.

4 What are We Blameworthy For?

Does the approach to thinking about blameworthiness outlined above leave out something important? Consider the following example (*Battlefield*) from Bernstein:

> Jane is at the battlefield and sees that four of her soldiers are about to be slaughtered by the enemy. She could save any one of them, but only one of them. (She only has one bullet left, and each one of her soldiers is being attacked by one enemy soldier). She cannot get herself to choose which one to save so they all die. (Bernstein, 2016, 436)

Bernstein says that this is a case of causal indeterminacy since "it is indeterminate *which* death … [Jane] caused" (2016, 438, emphasis in original). For this reason, Bernstein suggests that Jane's moral responsibility (2016, 442) and blameworthiness (2016, 447) are also indeterminate. I agree that we don't know everything needed to assess Jane's blameworthiness in *Battlefield*, but I suggest that this gap in our knowledge won't be filled by a more determinate picture of what Jane caused.

Bernstein considers the possibility of divorcing Jane's blameworthiness from her causal and moral responsibility, but she says that this would leave the crucial thing out of the story. It is the outcome, "the deaths of the soldiers," Bernstein says, "that are morally important to the case": the "[f]amilies of the deceased soldiers testifying in court against Jane" are mainly concerned "that her failure to act resulted in at least one death that could have been prevented" (2016, 443). "Attributing moral blameworthiness" to Jane independently of this outcome, Bernstein concludes, "leaves out a large swath of morally important information" (2016, 443).

While the deaths in *Battlefield* are obviously an important feature of the case, I don't believe that they are morally important in a way that's relevant to assessing Jane's blameworthiness. What's morally important in this way is the element of the case that is such that if we change it, we get a different conclusion about whether Jane is an apt target of blame. This element in *Battlefield* is not a soldier's death or even Jane's failure to prevent that death. Rather, the potentially blame-relevant factor in the case is whatever explains *why* Jane failed to prevent the death. Suppose that Jane failed to act because she suffers from an undiagnosed obsessive–compulsive disorder—perhaps akin to moral scrupulosity—that seriously impaired her ability to put an end to deliberation about which soldier's life to save. In this case, it would still be indeterminate which death Jane caused, but the issue of blameworthiness might be clearer: we might think that Jane, given her impairment, is not very blameworthy at all. Things would be different if Jane didn't act because she wasn't concerned about the fate of her comrades or because she bore them ill will. In this case, the puzzle about what Jane caused would remain, but blame directed at her would have a foundation given the explanation of her inaction.

But what if we asked a family member in *Battlefield* how they feel about Jane? To make things straightforward, let's assume that there's a determinate casual relation between Jane's inaction and the death of this person's child, and that Jane's inaction is explained in a way that makes it a clear case of moral affront. Suppose the parent says: "I blame Jane *for the death of my child*." This is a natural response, but perhaps not entirely clearheaded. We might note, as I did in the last paragraph, that Jane's being causally related to the soldier's death isn't enough to make Jane blameworthy. For blameworthiness, Jane needs to have brought about the death in a way that grounds blame. So, following Graham, (2014, 2017) we might say that what Jane is *really* blameworthy for is not the death of the soldier but rather the blame-grounding motives and intentions that led Jane to not try to save the soldier.[27] More generally, we might maintain with Khoury that "the only things that one can be blameworthy for are those things that make one blameworthy," which, for Khoury, can only be our objectionable willings (2018, 1363).

I largely agree with Graham and Khoury, but I think it also helps to employ here the distinction from Section 2 between weak and strong senses of "blameworthy for." Applied to outcomes, the weak sense of "blameworthy for" presupposes an appropriate causal relation between an outcome and a blameworthy agent. We can say that Jane is blameworthy for the soldier's death in this weak sense if the death is appropriately causally related to factors that (independently and exclusively) make Jane blameworthy. But the presence of this causal relationship would not make Jane blameworthy for the death in a stronger sense. The stronger sense of "blameworthy for" specifies the factors in virtue of which a person is open to blaming responses, and I claim that while Jane is open to such responses, this is not because of the soldier's death or her connection to it. Rather, Jane is a fitting target of blame because of her independently objectionable motives and intentions. So, Jane is blameworthy in the strong sense for these motives and intentions, and since these blame-grounding factors are appropriately related to the soldier's death, we can also say, if we like, that Jane is to blame for the death, but only in the weak sense of "blameworthy for."[28]

Does this still leave something important out of the story? Two points can be added. The first is that, in the version of *Battlefield* where Jane clearly caused a particular soldier's death, the family's awareness of the death is the occasion for their blame. The death is *why* they blame Jane (even if it is not why Jane is blameworthy); it is what alerts family members to the possibility that there are grounds for blame, and it largely accounts for the family's interest in blaming Jane. The second, closely related point is that the soldier's death, and the fact that Jane brought it about, is the lens through which Jane's blameworthiness comes into focus for the family members. Their experience of blaming Jane, and their morally aggrieved emotions, are filtered through the experience of losing a family member and the emotional and cognitive states—grief, bewilderment, remembrance, etc.—that accompany such a

[27] Graham (2014) argues that we are blameworthy only for our bad motivations. In Graham (2017), the claim is put in terms of intentions.

[28] See Smith (1983, 567) for such a use of "to blame for."

loss. So, the soldier's death will play an important role in the overall account of the blame that the family members direct toward Jane, even if the death does not materially affect Jane's worthiness of blame.

And what happens if we remove the soldier's death from the equation? Suppose the soldier's family believes that Jane's objectionably-motivated intention to not save the soldier led to the soldier's death, but then the family learns that the soldier actually survived. The grounds for blaming Jane—her objectionably-motivated intention—are still present, but the outcome for which Jane was thought to be (weakly) to blame is absent. This will affect the experience of blaming Jane since the moral offense felt will no longer be accompanied by dismay and grief over the soldier's death. And because there is no attention-focusing outcome for which Jane is (weakly) to blame, the morally offended blaming responses may themselves be muted and may more quickly recede even though the grounds for these responses are still present. In this case, Jane may simply be worthy of more blame than she receives.[29] The soldier's survival makes Jane lucky insofar as it inhibits blame, but this is not *moral* luck since, on my view, it does not actually free Jane from blameworthiness.[30]

5 Objections from the Resultant Luck Literature

Zimmerman (2002) distinguishes between the *scope* of an agent's moral responsibility and that agent's *degree* of moral responsibility. On Zimmerman's view, what falls within the scope of one's responsibility may be subject to luck, but one's degree of responsibility is unaffected by chance outcomes.[31] Suppose that George shoots at Henry and kills him, and that Georg shoots at Henrik but fails to kill him (Zimmerman, 2002, 560–61). In both cases, the result is not entirely within the agents' control, but Zimmerman says that Henry's death—given the way that it was brought about—falls within the scope of George's moral responsibility, whereas Georg has nothing similar within the scope of his responsibility (since no action of his resulted in a death). However, according to Zimmerman, George and Georg may be morally responsible to the same degree. (I take this to mean that they are equivalently blameworthy.) What makes George and Georg responsible to the same degree is that they both "would have freely killed someone," though of course only George does kill

[29] As Graham notes on the basis of similar reflections, "[w]e often don't blame people as much as it would be altogether appropriate for us to" (2014, 406). See also, Talbert (2019, 33).

[30] I've given some reasons why an agent's causing an unwelcome outcome might give rise to more blame than if the outcome had not occurred. But a referee rightly emphasizes that our actual tendency to respond with blame to agents who cause bad outcomes goes beyond anything my approach can accommodate. Why do our blaming practices go astray? Here, I would appeal to the idea, defended by several authors (with some empirical support)—Royzman and Kumar (2004), Domsky (2004), Levy (2016), Talbert (2017)—that there is a human tendency to *mis*attribute to others the sort of moral disregard that would (if it were rightly attributed) ground blaming responses, particularly when their actions give rise to unwelcome outcomes. Aspects of our blaming practices are explained, but not justified, by such misattributions.

[31] I take Zimmerman's view here to be related to my suggestion that we may be morally responsible for outcomes (if we stand in the right causal and epistemic relations to them) but are not blameworthy on account of such outcomes.

someone (Zimmerman, 2002, 565). On Zimmerman's view, facts about what agents *would have done*—facts that do not depend on actual outcomes—fix their degree of moral responsibility.

Hartman (2017) argues that this counterfactual component of Zimmerman's view leads to a serious problem.[32] On a view like Zimmerman's, "agents may be praiseworthy or blameworthy in virtue of events that are radically different from the kind for which they are praiseworthy and blameworthy in the actual world" (Hartman, 2017, 65). A person's actual moral character, and the motives and intentions upon which they act, may be entirely unobjectionable, but it might be true that under suitable counterfactual conditions they would act in morally terrible ways. Plausibly, such counterfactual claims are true of all of us, so we may all be blameworthy to a much higher degree than we ordinarily suppose. Zimmerman is aware of this feature of his view. He says that "since an indefinite number of counterfactuals about what one would do … can be true at once, one can be morally responsible … to an indefinite number of degrees at once"; this "opens up the floodgates, as it were, when it comes to ascriptions of responsibility…" (2002, 570).

I agree with Hartman that this is an unappealing feature of Zimmerman's view, but it is not a feature of the view that I defend. For Zimmerman, degree of responsibility is not affected by outcomes because it is fixed by the truth of counterfactual claims, which is unaffected by actual outcomes. The view I defend is that degree of blameworthiness is unaffected by outcomes because blameworthiness depends on such things as an agent's *actual* motives and intentions, which are unaffected by outcomes. The fact that a morally decent person might have objectionable motives and intentions under counterfactual conditions does not, according to my view, entail that the agent is open to moral blame because it does not necessarily entail anything about the agent's actual motives and intentions.

Hartman also argues that the tendency to deny that outcomes affect blameworthiness stems from a conflation of two different things: an agent's moral worth and their moral record. Hartman follows Greco (1995) here. Hartman and Greco allow that a drunk driver who unluckily kills someone is no worse than an identical driver who luckily avoids killing anyone. The problem, Hartman says, is that skeptics about resultant moral luck "mistakenly infer *from* the claim that each driver is no worse of a person than the other *to* the claim that each driver is no more blameworthy than the other" (2017, 118, emphasis in original). Hartman and Greco argue that this inference is faulty: while the two drivers may be equivalent in terms of moral worth, the unlucky driver is more blameworthy because they have a worse moral record insofar as they are morally responsible for killing someone.

We can, of course, distinguish between a person's moral record and their moral worth. The problem is to explain *why* a person's being responsible for a bad outcome—having a worse moral record—makes it appropriate to direct more blame at them. Hartman is clear both that he construes blame as targeting a person with morally offended emotions like resentment (2017, 32), and that an agent's being

[32] For a similar argument, see Hanna (2014). For more detailed consideration of Hartman's (and Hanna's) objection, along the lines suggested below, see Talbert (2019).

blameworthy for an outcome means that they deserve more blame—more resentment—than if the outcome had not occurred (2017, 34). But as I argued in Section 2, it is obscure why the occurrence of a bad outcome should ground more intense morally offended responses since outcomes do not, in themselves, make a person's actions more morally offensive. I don't see that Hartman's or Greco's account addresses this issue.

Lang (2021), however, takes this issue head on. Lang's defense of resultant moral luck depends on fusing moral record and moral worth: the quality of the outcomes that contribute to an agent's moral record should, Lang argues, inform our judgment about the moral quality of the agent's intentions. To illustrate, Lang introduces *Two Buttons* (2021, 52). In this case, Akira and Dylan are both disposed to unjustifiably harm strangers, and they are both inclined to press one of two buttons, knowing that pressing either button will injure some innocent person or persons. What Akira and Dylan do not know is that pressing Button 1 will cause the deaths of ten innocent strangers and that pressing Button 2 will cause one innocent stranger's arm to be broken.

Selecting at random, Dylan presses Button 1 and Akira presses Button 2. Lang argues that the difference in outcomes caused by Dylan's and Akira's actions affects their blameworthiness. Specifically, "[*b*]*ecause* he caused more harm … Dylan is more blameworthy than Akira" (2021, 53, emphasis in original). And Dylan is more blameworthy because the result of his action reveals him to be a morally worse agent, with morally worse intentions, than Akira: "because what Dylan caused to happen is worse than what Akira caused to happen, we should conclude that Dylan's internal states are morally worse than hers. And that, in turn, makes Dylan morally worse than Akira" (2021, 61). For Lang, the "badness of [Dylan's and Akira's intentions] … is fixed by actual outcomes…. the badness of the internal states which are embodied in these agents' acts is fixed retrospectively, through importing information about how these acts turn out" (2021, 63).

Lang is correct, I think, in seeing a move along these lines as crucial for defending resultant luck's supposed effect on blameworthiness (2021, 60–61). If outcomes affect the moral status of internal states, then it is clear how they can affect blameworthiness. Moreover, there's certainly a sense in which the badness of Dylan's intentions is fixed by the outcome of his action. Given the results, it is a worse thing than it might have been that Dylan intended to cause harm, and the presence of that intention in Dylan is, given the results, a more regrettable thing than the presence of the same intention in Akira. It's also true, as Lang says, that Dylan's intentions "have a significance that is not fully disclosed until we discover what they cause to happen" (2012, 61). As discussed in the last section, the salience of a person's intentions for us often depends on the salience of the outcomes to which those intentions give rise.

But none of these considerations entail that Dylan's intentions are morally more objectionable than Akira's in a way that's relevant to the fittingness of morally offended blaming responses. Independently of the need to make room for resultant moral luck, I don't think that Lang gives us much reason to move from the modest claims in the last paragraph to a substantial conclusion about Dylan's and Akira's blameworthiness. It may be a more regrettable thing that Dylan had his intention, but

that doesn't make his intention a greater offense against morality in a way that matters for blameworthiness. As I argued in Section 2, intentions matter for blameworthiness because they manifest our degree of regard for others. For outcomes to affect our intentions in a way that matters for blameworthiness, the occurrence of a bad outcome would have to show that an agent's intention embodies a more significant disregard for others than would have been the case had the outcome not occurred. And that, I suggest, is not something that the occurrence of an outcome can do.

6 Conclusion

I have argued that regardless of whether we can be morally responsible for outcomes, we cannot be substantively blameworthy for them: the occurrence of an unwelcome outcome does not affect an agent's openness to the responses involved in moral blame. The basis for this argument is the thought that the morally offended reactions that characterize moral blame are responses to moral affronts, and that unwelcome outcomes are not in themselves such affronts. I have attempted to motivate this central idea by considering cases of causally complex agency. Determining the causal structures of these cases may help us decide what an agent has caused, but the factors relevant to making such determinations do not help us decide whether a given case involves a moral affront, so such determinations are not relevant for our judgments about an agent's openness to blame. Still, as I have argued, outcomes do play an important role in explaining *why* we blame in specific instances. So, reference to outcomes need not disappear from our description of a particular instance of blame, it is just that such references are not part of the account of whether, and why, an agent is open to blame.

Acknowledgements I would like to thank several people with whom I have had fruitful conversations about the topics in this paper: Gunnar Björnsson, Devin Curry, Thor Grünbaum, Mattias Gunnemyr, Robert Hartman, Daniel Miller, Björn Petersson, Toni Rønnow-Rasmussen, András Szigeti, Daniel Telech. I'd particularly like to thank Olle Blomberg, Huzeyfe Demirtas, and Caroline Touborg, as well as two anonymous referees for Erkenntnis, for their very helpful written comments.

Funding Open access funding provided by Lund University.

Declarations

Competing Interests The author has no relevant financial or non-financial interests to disclose.

The author has no competing interests to declare that are relevant to the content of this article.
The author certifies that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.
The author has no financial or proprietary interests in any material discussed in this article.

References

Ayars, A. (2021). Blaming for unreasonableness: Accountability without Ill will. *Journal of Ethics and Social Philosophy, 19*, 56–79.
Bernstein, S. (2016). Causal and moral indeterminacy. *Ratio, 29*, 434–447.
Bernstein, S. (2017). Causal proportions and moral responsibility. In D. Shoemaker (Ed.), *Oxford studies in agency and responsibility* (Vol. 4, pp. 164–182). Oxford University Press.
Bernstein, S. (2019). Moral luck and deviant causation. *Midwest Studies in Philosophy, 43*, 151–161.
Björnsson, G. (2017). Explaining (Away) the epistemic condition on moral responsibility. In P. Robichaud & J. W. Wieland (Eds.), *Responsibility: The epistemic dimension* (pp. 146–162). Oxford University Press.
Clarke, R. (2014). *Omissions: Agency, metaphysics, and responsibility*. Oxford University Press.
D'Arms, J., & Jacobson, D. (2003). The Significance of recalcitrant emotion (or, Anti-quasijudgmentalism). In A. Hatzemoysis (Ed.), *Philosophy and the emotions* (pp. 127–145). Cambridge University Press.
Demirtas, H. (2022a). Moral responsibility is not proportionate to causal responsibility. *The Southern Journal of Philosophy, 60*, 570–591.
Demirtas, H. (2022b). Against resultant moral luck. *Ratio, 35*, 225–235.
Domsky, D. (2004). There is no door: Finally solving the problem of moral luck. *The Journal of Philosophy, 101*, 445–464.
Fischer, J. M., & Ravizza, M. (1998). *Responsibility and control: A theory of moral responsibility*. Cambridge University Press.
Graham, P. (2014). A sketch of a theory of moral blameworthiness. *Philosophy and Phenomenological Research, 88*, 388–409.
Graham, P. (2017). The epistemic condition on moral blameworthiness: A theoretical epiphenomenon. In P. Robichaud & J. W. Wieland (Eds.), *Responsibility: The epistemic dimension* (pp. 163–179). Oxford University Press.
Greco, J. (1995). A second paradox concerning responsibility and luck. *Metaphilosophy, 26*, 81–96.
Gunnemyr, M., & Touborg, C. T. (2023). You just didn't care enough: Quality of will, causation, and blameworthiness for actions, omissions, and outcomes. *Journal of Ethics and Social Philosophy, 24*, 1–35.
Hanna, N. (2014). Moral luck defended. *Noûs, 48*, 683–698.
Hartman, R. J. (2017). *In defense of moral luck: Why luck often affects praiseworthiness and blameworthiness*. Routledge.
Hieronymi, P. (2008). Responsibility for believing. *Synthese, 161*, 357–373.
Kaiserman, A. (2017). Partial liability. *Legal Theory, 23*, 1–26.
Kaiserman, A. (2018). 'More of a cause': Recent work on degrees of causation and responsibility. *Philosophy Compass, 13*, 1–10.
Kaiserman, A. (2021). Responsibility and the 'Pie Fallacy.' *Philosophical Studies, 178*, 3597–3616.
Khoury, A. (2018). The objects of moral responsibility. *Philosophical Studies, 175*, 1357–1381.
Lang, G. (2021). *Strokes of luck: A study in moral and political philosophy*. Oxford University Press.
Levy, N. (2016). Dissolving the puzzle of resultant moral luck. *Review of Philosophy and Psychology, 7*, 127–139.
McKenna, M. (2012). *Conversation and responsibility*. Oxford University Press.
Miller, D. J. (2022). Attributionism and degrees of praiseworthiness. *Philosophical Studies, 179*, 3071–3087.
Richards, N. (1986). Luck and desert. *Mind, 95*, 198–209.
Riedener, S. (2021). Don't make a fetish of faults: A vindication of moral luck. *Philosophical Studies, 178*, 693–711.
Rosen, G. (2015). The alethic conception of moral responsibility. In R. Clarke, M. McKenna, & A. M. Smith (Eds.), *The nature of moral responsibility* (pp. 65–87). Oxford University Press.
Royzman, E., & Kumar, R. (2004). Is consequential luck morally inconsequential? Empirical psychology and the reassessment of moral luck. *Ratio, 17*, 329–344.
Sartorio, C. (2012). Resultant luck. *Philosophy and Phenomenological Research, 84*, 63–86.

Sartorio, C. (2015a). A new form of moral luck. In A. Buckareff, C. Moya, & S. Rossell (Eds.), *Agency, freedom, and moral responsibility* (pp. 134–149). Palgrave Macmillan.
Sartorio, C. (2015b). Resultant luck and the thirsty traveler. *Methode, 4*, 153–171.
Sartorio, C. (2016). *Causation and free will*. Oxford University Press.
Sartorio, C. (2020). More of a cause? *Journal of Applied Philosophy, 37*, 346–363.
Scanlon, T. M. (2008). *Moral dimensions: Permissibility, meaning, and blame*. Harvard University Press.
Sher, G. (2009). *Who knew?: Responsibility without awareness*. Oxford University Press.
Smith, H. (1983). Culpable ignorance. *The Philosophical Review, 92*, 543–571.
Strabbing, J. T. (2019). Accountability and the thoughts in reactive attitudes. *Philosophical Studies, 176*, 3121–3140.
Strawson, P. F. (1962). Freedom and resentment. *Proceedings of the British Academy, 48*, 1–25.
Talbert, M. (2015). Responsibility without causation, luck, and dying of thirst: A reply to sartorio. *Methode, 4*, 173–184.
Talbert, M. (2016). *Moral responsibility*. Polity Press.
Talbert, M. (2017).Omission and attribution error. In D. Nelkin and S. C. Rickless (Eds.), *The ethics and law of omissions*. Oxford University Press, pp. 17-35.
Talbert, M. (2019). The attributionist approach to moral luck. *Midwest Studies in Philosophy, 43*, 24–41.
Talbert, M. (2022). Attributionist theories of moral responsibility. In D. K. Nelkin & D. Pereboom (Eds.), *The Oxford handbook of moral responsibility* (pp. 53–70). Oxford University Press.
Talbert, M. (2023). Blaming reasonable wrongdoers. *Res Publica*. https://doi.org/10.1007/s11158-023-09644-w.
Thomson, J. J. (1989). Morality and bad luck. *Metaphilosophy, 20*, 203–221.
Zimmerman, M. (1988). *An essay on moral responsibility*. Roman and Littlefield.
Zimmerman, M. (2002). Taking luck seriously. *Journal of Philosophy, 99*, 553–576.