

# Responses to Blomberg, Payton, O'Brien, Bernstein, and Sebeq

Saba Bazargan-Forward

## 1. INTRODUCTION

The *Journal of Social Ontology* (JSO) has granted me this opportunity to address a panoply of insightful comments and criticisms which an exceptional group of symposiasts have provided in response to my book, *Authority, Cooperation, and Accountability*. I am grateful to the editors at JSO for selecting my work among the many on offer, and to the symposiasts for their generosity in dedicating their time and attention to addressing the many challenges my account faces, of which I can respond only to a handful here. It is convention—but no less true—to say at this point that I only wish I had the benefit of these comments and criticisms *before* the book was published. As is, this meagre, rearguard attempt to defend the account I have developed will have to suffice.

## 2. OLLE BLOMBERG

I develop authority-based accountability as a distinct way in which one agent can be accountable for the purpose with which another agent acts, even when the former is not a cause of what the latter does. In such cases, the executor's purpose is grounded in his protected reasons to do as the deliberator

10.25365/jso-2024-9151

Published online November 30, 2024

Saba Bazargan-Forward, University of California, San Diego, E-mail: sbazargan@ucsd.edu

Open Access. © 2024 Author(s) published by the Journal of Social Ontology. This work is licensed under the Creative Commons Attribution 4.0 International License.

says. Olle Blomberg deftly argues that ordinary causal influence can ground accountability both for what others do and for the purpose with which they do it.

He imagines a case in which I call your mobile phone at 5:30 am to awake you. I thereby furnish your phone's buzzing with the use-based agentic function of waking you up. Similarly, Blomberg suggests, Contender causally assigns to Goon—through the agreement with him—the use-based agentic function of facilitating her victory by maiming Victim. On this view, ordinary intentional action can be enough to confer use-based agentic functions upon events in the world. If so, why believe that use-based agentic functions must be assigned via protected reasons?

Let's grant that Blomberg is correct in his surmise that ordinary causal influence can ground accountability both for what others do *and* why they do it. But whether it *can* do so is not enough; an aim of the account I develop throughout the book is to demonstrate that there is an important way in which we can be accountable for what others do even when we do *not* causally contribute to their conduct, such as in large-scale cooperative action. This is why I develop an alternative to the causal account. According to this alternative, protected reasons can license attributing to the executor's actions the purpose of enacting the deliberator's reasons (Bazargan-Forward 2022, 5–7). Still, I take Blomberg's point that it is possible to confer a purpose in the minimal, causal way he suggests.

But Blomberg goes further, by arguing that causation isn't just sufficient but necessary in that a deliberator cannot confer a purpose on an executor without it. On this view, the deliberator's accountability requires a non-deviant causal connection between her instructions and the executor's actions. To show this, Blomberg considers cases in which Goon happens to do what he promised to do, but for reasons completely unrelated to his prior agreement with Contender. Suppose Goon maims Victim because she happened to be the woman who insulted him on the street; he is unaware that she is the one he had earlier promised Contender to maim. Suppose now Victim wants to determine why she was maimed. Blomberg maintains that Deliberator's motivating reasons are irrelevant here. "Absent a non-deviant causal connection between the deliberator's instructions and the executor's action," Blomberg says, "the deliberator arguably bears no authority-based accountability for the (alleged) fact that the executor's action has a wrongful purpose furnished by the deliberator" (Blomberg 2024, 15).

This is a compelling argument, but I think the conclusion is premature. Generally speaking, if I make a promise to do  $\varphi$ , I thereby fulfil it by doing

$\varphi$ , even if I didn't do  $\varphi$  because of the promise, and even if I didn't know I was fulfilling the promise at the time. The conditions of promise-fulfilment are "externalist" in this way. Similarly, if we agree to do  $\varphi$  for others, and they thereby confer a particular purpose on us, and we subsequently do  $\varphi$ , we thereby act in accordance with that purpose, even if we didn't do  $\varphi$  because of the agreement, and even if we didn't know we were fulfilling the agreement at the time. The conditions of purpose-fulfilment are "externalist" in this way. This applies to Blomberg's version of *Olympic Sabotage*; Goon has fulfilled the terms of his promise to Contender. Given that Goon was promised payment in return for fulfilling the terms of his promise, Contender owes Goon that payment. We can imagine Contender saying to Goon as she hands him the cash: "Well, you did what you agreed to do, though not quite in the way either of us expected. Either way, you fulfilled the purpose I had for you." This is a perfectly apt claim. Now imagine that Victim learns of Goon's relationship to Contender; moreover, she learns that Contender conceded that Goon fulfilled the purpose she conferred upon him. There is good reason to think that Victim would now be unsatisfied with the explanation Goon had previously given her as to why he assaulted her. In this respect, *Goon's motivating reasons under-describe the purpose of the assault*. It seems that if Victim wants to know what purpose her assault had, she has to consult not just Goon but Contender. For this reason, I think that Contender does indeed bear authority-based accountability for the wrongful purpose she furnished upon Goon. The same point applies to Blomberg's discussion of *Double Promise* (Blomberg 2024, 17).

More generally, Blomberg raises the following doubt. How can the deliberator's reasons have the function of normatively "guiding" the executor's conduct, and how can the executor successfully enact those reasons, unless they causally influence the executor's conduct? When I say that the deliberator's reasons have the function of normatively "guiding" the executor's conduct, I do not mean that the executor can or should consult the purpose conferred upon him in deciding what to do. Blomberg is correct in thinking that this is impossible in cases of the sort he has in mind. Rather, what I mean is this: the executor's conduct is *teleologically* directed toward ends which the deliberator fixes. It is in this sense that the deliberator's reasons normatively "guides" the executor's conduct. I confess, it was misleading for me to use the word "guiding" here; I should have spoken in terms of teleology instead.<sup>1</sup>

Near the end of his article Blomberg considers a collective-action case, *Assassination Fund*; there, he suggests my account doesn't go far enough.

<sup>1</sup> Indeed, I do at Bazargan-Forward (2022, 8, 33, 39).

We want to say not only that each donor is accountable for conferring a purpose upon the villain, but for the act of hiring the hitman. To this end, we might advert to more sophisticated causal accounts, or to membership-based accounts; in either case authority-based accountability might prove otiose. Regarding causal accounts: consider a case where a Villain accepts a particular donor's money, but then loses it. Since the funding overdetermined, Villain is able to hire the assassin anyway. It is difficult to see how a causal account, no matter how sophisticated, can accommodate the donor's accountability for anything other than an attempt, whereas according to authority-based accountability the causally ineffective donor's attempt succeeds in conferring a particular purpose upon the Villain. Regarding membership-based accounts: to evaluate them, we of course need to hear more about how and why mere membership is inculpatory. Ultimately, I agree with Blomberg when he suggests that authority-based accountability might understate accountability in cases like *Assassination Fund*. For this reason, I think authority-based accountability is only part of the picture when it comes to individual accountability for cooperatively committed harms. Still, I maintain that it is an indispensable part.

### 3. JONATHAN D. PAYTON

Payton suggests that the analysis I provide of intrapersonal diachronic decision-making—which itself is drawn from Luca Ferrero's work—isn't quite right. My decision at  $t_1$  to  $\varphi$  at  $t_2$  settles the matter for me in that at  $t_2$  I needn't redeliberate about what to do. The earlier decision does not add to my stock of reasons when it comes time to act at  $t_2$ . Rather, the reasons I took there to be at  $t_1$  for  $\varphi$ , are applicable now at  $t_2$ . So, if I decided to  $\varphi$  for foolish reasons, then it's not only my previous decision but my current action that can be criticized as such. Ferrero suggests that this is because my decision to  $\varphi$  at  $t_2$  acts as a kind of anaphoric device, referring back to my original reason at  $t_1$ . I borrowed from Ferrero's account, not only by suggesting that *intrapersonal* decision-making functions anaphorically in this way, but that *inter-personal* decision-making does so as well.

Payton casts doubt on Ferrero's claim by pointing out that the proposition "I decided to  $\varphi$ " contains no anaphor or other mechanism by which it might literally refer to the reasons for which I made my decision. Rather, the proposition's content is that I decided to  $\varphi$ . Strictly speaking, Payton is correct; the content of the proposition in question contains no anaphoric device. But it might be better to understand Ferrero's point in this

way: the content of my purported *justification* for doing  $\varphi$  at  $t_2$  will include an anaphoric device referring back to my earlier decision from  $t_1$ .

Moreover, though I borrow from Ferrero's account, I weaken it in an important way when applying it to inter-personal divisions of agential labor. Instead of claiming that the executor acts for the deliberator's *reasons*, I claim that the executor acts according to a *purpose* conferred upon him, where the purpose derives its content from the deliberator's motivating reasons. Payton notes this, and points out that by making this move, I sidestep his criticism of Ferrero's account.

But Payton suggests that switching from reasons to purposes in this way undercuts some of the moral claim I wish to make. In particular, I claim that conferring a wrongful purpose on an executor can turn an otherwise permissible act into an impermissible one, or a virtuous act into a vicious one. How can this be if, by hypothesis, the executor is *not* acting for the deliberator's reasons?

On my account, the functional relationship between the deliberator's motivating reasons and the executor's conduct is what licenses adverting to the those motivating reason in evaluating the executor's conduct. Payton worries that this isn't enough. He considers a case I describe in which a sexist business owner orders one of his managers to promote a male subordinate, not because the subordinate is deserving of a promotion (which he is) but because he is male. Here, the executor's motivations are pure, and his conduct, considered apart from the deliberator's reasons, is permissible. "How, then," Payton writes, "could he have acted impermissibly?" (Payton 2024, 24)

It's true that the manager's actions when "considered apart from [the deliberator's] reasons" are morally unproblematic (Payton 2024, 24). But I deny that we should evaluate the manager's actions in isolation. The thought here is that it can be a pro tanto wrong for you to fulfill the sexist intentions that others have—specifically, in cases where you are purposed with doing so—even if you don't share in those sexist intentions. Payton seems to suggest otherwise. Of course, Payton might be right. But we cannot settle this matter by just *assuming* that we should evaluate the manager's conduct in isolation, as *that* is precisely what's at issue.

Continuing this line of thought, Payton says of the executor that if he "does something permissible for permissible reasons, how could his *character* possibly be impugned?" (Payton 2024, 24) But as I suggested, what he does is a pro tanto wrong, and his character is impugned to the extent that this he blameworthy for that wrong.

Taking a step back, Payton argues that Ferrero's analysis of diachronic rational agency allows you under certain conditions to disavow at  $t_2$  the decision you made at  $t_1$ . The basis for acting on a previous decision is that, were you to deliberate anew under idealized conditions, you would reach the same conclusion now as you did before when it comes to deciding what to do. When this is not the case—as in the example Payton describes derived from *Total Recall*—there is little rational basis for enacting the decision you make in the past. The suggestion here seems to be that the same goes for cases in which a deliberator confers a purpose upon an executor that the latter wouldn't accept were he to know what that purpose is. In such cases the purpose conferred doesn't yield a protected reason for the executor.

I agree with Payton's analysis of the *Total Recall* case and its implications for intrapersonal divisions of agential labor but deny that it has much bearing on cases of interpersonal divisions of agential labor. In intrapersonal cases, what grounds the protected reasons to "obey" your past self—and what explains when and why you are permitted to "disobey"—is the importance of a commitment to a diachronically unified set of aims. But this isn't necessarily so when it comes to interpersonal cases (Bazargan-Forward 2022, 32). There might be a presumption in favor of a Bratmanian *instrumental* rationality among the cooperants, but not necessarily any deeper commitment to ensure that we're all acting for the same motivating reasons (Bratman, 1993). As a result, the conditions under which an executor can unilaterally disavow the purpose that a deliberator confers upon him will be quite different from the conditions under which you can unilaterally disavow a decision your past self made for you. Either way, Payton has effectively demonstrated that I must say more to elucidate the analogy between intrapersonal and interpersonal decision making—a task which I cannot undertake here.

On Payton's preferred analysis, when I do as I previously decided, or what you decided for me, "I'm volunteering to let those reasons serve as the standard against which my behaviour is to be judged" (Payton 2024, 26).

In this way, the reasons behind my decision at  $t_1$  serves as a basis for evaluating my conduct at  $t_2$  even if that conduct wasn't motivated by those reasons. Payton adds an important caveat though. Typically, I will *conditionally* volunteer to let your reasons or my previous reasons serve as a basis for evaluating what I do now. Put in terms of interpersonal divisions of agential labor: the executor might vest authority in the deliberator on the condition that that the purpose the deliberator thereby bestows is one which the executor would avow, given the choice. This provides the executor with what Payton describes as an "escape clause" (Payton 2024, 27). But consider this case,

entitled “Ignorant Executor” which I discuss in chapter 4:

**Ignorant Executor.** Deliberator asks Executor to drive a car to a junkyard—with the added instruction that he refrain from looking in the trunk. Executor isn’t sure why Deliberator wants this done. Neither is he sure what is in the car’s trunk. Nor is he sure who Deliberator is exactly. He suspects though, that Deliberator is a person of some influence, engaged in some sort of illicit activity. Executor agrees to do as Deliberator says, since Executor wants to ingratiate himself to Deliberator. (Bazargan-Forward 2022, 116–7)

Suppose further that Executor agrees to do as Deliberator asks, but only if the Deliberator possesses no illicit intentions or aims. It’s absurd to think that Executor can so easily get out from under Deliberator’s practical authority. For this reason, I argue in chapter 4 that escape clauses generally fail to function as such *if* the executor has reason to suspect that the deliberator won’t fulfill its conditions. And even if the executor has no reason to suspect that the deliberator possesses illicit intentions or aims, it’s *still* unclear why we should need an escape clause. In such cases, the inimical purpose conferred upon the executor’s actions will not impugn her character; she is not to blame for that wrong-making feature of her conduct. At worst, the executor should feel something akin to agent regret. Indeed, Payton suggests this sort of analysis, and I agree with it—a point I discuss in further detail in the book’s conclusion.

#### 4. LILIAN O’BRIEN

What is the relationship between motivating reasons (as I understand them) and purposes in action? I claim that motivating reasons find expression in the content of an action’s purpose; Lilian O’Brien casts doubt on this claim. She imagines a case in which you are deciding between purchasing two cars, A and B. Car A has lower mileage and is cheaper. These facts serve as motivating reasons for you to purchase car A. But these facts do not seem to “constitutively determine” your purpose in purchasing that car. Your purpose is, simply, to purchase a car, *sans phrase*. It seems, then, that it is mistaken to think that motivating reasons for an action feature in the content of that action’s purpose.

Imagine, though, an obscenely rich person who really doesn’t care about mileage or cost in purchasing a car. At a suitably general level of description, both of you have the same purpose: to buy a car. But my claim is that if we unpack your purpose and his purpose, we will see that they differ in important ways. On O’Brien’s view, they don’t. On her view, the content of a purpose does

*not* include all the satisfaction-conditions specified in the motivating reason behind it. On my view, it does.

On the picture O'Brien adopts, when we deliberate about what to do, the candidate options include actions, each of which are characterized independently in terms of a particular purpose. On this picture, the motivating reasons in favor of a particular action do not determine that action's purpose. Instead, the motivating reasons serve as the basis for choosing among candidate purposes. This picture, if correct, suggests that a deliberator's motivating reasons do *not* feature in the description of the purpose the deliberator confers upon an executor. Such a conclusion would undermine much of the argument I develop in favor of authority-based accountability.

Let's call a "narrow purpose" a purpose the content of which excludes a description of the actor's motivating reasons. And let's call a "broad purpose" a purpose the content of which includes a description of the actor's motivating reasons. At one remove, it might be misguided to ask which concept is "correct" in the abstract. The "right" concept depends on the role it plays in a theory of mind. I take it that the role of attributing a purpose to an action is to *explain* that action. Purposes construed broadly do a better job of this than purposes narrowly. This is because narrow purposes fail to be modally robust.

To see why, consider the canonical example of the pilot deciding whether to drop bombs on a munitions factory. Suppose the pilot's motivating reason in doing so is to kill the villagers nearby. If O'Brien is right, the pilot's purpose should be described in terms free from reference to motivating reasons. The pilot's purpose, then, is just to drop the bomb on the munitions factory. But this narrow purpose fails to explain why the pilot *wouldn't* drop the bomb if there were no villagers nearby. For a purpose to explain action, its content needs to be derived at last in part from a description of the actor's motivating reasons.

This goes also for ordinary cases in which double-effect isn't an issue. Suppose you take your dog for a walk. You do so to maintain your dog's health. Your purpose, construed narrowly, is simply to take your dog for a walk; this purpose includes no reference to your dog's health. But this narrow purpose, as described, would fail to explain why you wouldn't take your dog for a walk if it were cold and raining. Again, this narrow purpose fails to be modally robust and hence fails to explain action.

Of course, O'Brien might deny that the role of the concept "purpose" is to explain action. Instead, narrow purposes *in combination with* motivating reasons might be what explains action. What explains the pilot's conduct isn't his purpose, simply, but his purpose in combination with the motivating reasons behind it; ditto for the dog-walker. In keeping with this, O'Brien points



out that regardless of how we explain action, it is important to keep motivating reasons separate from purposes. Suppose I adopt the purpose of driving myself to work; it would be misleading to say that this purpose gives me *added* reasons to get into a car. Instead, that purpose rationally *necessitates* getting into a car. So, O'Brien suggests, if motivating reasons *just are* considerations favoring a course of action, then they cannot serve as purposes.

This point is well taken. A purpose cannot be something that *adds reasons* favoring actions contributing (instrumentally or constitutively) to the achievement of that purpose. But the reasons you take there to be in favor of a purpose can be folded into the description of that purpose without thereby suggesting that the purpose is something *that adds reasons* favoring actions contributing to the achievement of that purpose. Suppose I adopt the purpose (narrowly construed) of driving myself to work. And suppose I adopt that purpose because I *enjoy* the drive. That's my motivating reason. We might therefore say that my purpose (broadly construed) was to *enjoy driving to work*. Here, the motivating reason features in the description of my purpose. This broad purpose doesn't *add reasons* to get into the car. Rather, it rationally necessitates getting into the car, so long as doing so instrumentally or constitutively contributes to the achievement of that purpose. Likewise, Contender's adopted purpose does not serve as an added reason for Goon to maim Victim. Rather, the maiming is rationally necessitated by dint of the adopted purpose. The account I develop is friendly to this analysis, since it regards Goon's reasons as *protected* which means they are supposed to be decisive rather than function merely by adding to the stock of Goon's reasons.

I'm reluctant, then, to abandon the claim that motivating reasons partly determine the content of our purposes. Still, O'Brien helpfully develops an alternative picture of how my account might go. On her alternative reconstruction, the purpose of Goon's attack is determined not by Contender's motivating reasons, but by Contender's intention to *make it the case* that Victim is put out of the running. But if a *deliberator* must form an intention for it to be the case that an executor fulfills her overall purpose, then the two of them seem function to more like co-executors rather than a deliberator and an executor. As O'Brien suggests, this alternative formulation will have difficulty accommodating the range of cases to which I put my own.

In addition, what warrants conceptualizing Contender and Goon as a deliberator and an executor (and not as co-executors) is in part that Contender has authority over Goon but not vice-versa, where that authority yields protected reasons. It's less clear that a co-executor retains this kind of unilateral authority. But perhaps we can formalize this asymmetry in O'Brien alternative

reconstruction insofar as one co-executor but not the other is in charge of determining the overall purpose for which they are acting. Either way, O'Brien alternative reconstruction is clearly worth developing especially considering that my account is heterodox in its treatment of the distinction between purposes and reasons.

## 5. SARA BERNSTEIN

Bernstein makes two principal claims about the account I develop. First, she suggests that there are many more sources of reasons and purposes than those I allude to in the account I develop. Second, she suggests my view undergenerates results in some cases and overgenerates in others. In particular, my view has difficulty accommodating various forms of moral luck especially as they pertain to the sources of reasons.

Bernstein points out that a variety of social phenomena can provide motivating reasons. For example, one's upbringing, social categories, conventions, and algorithms on social media, can each yield motivating reasons to behave in certain ways. Bernstein suggests that in such cases, the executor is acting on the basis of purposes that are not furnished by any identifiable deliberator.

It's undoubtedly true that the phenomena of the sort Bernstein describes yield motivating reasons for us. Meena has a motivating reason to remain a vegetarian; Jamal has a motivating reason to be cautious around police; Cordelia has a motivating reason to behave in ways stereotypical of professors; Alexia has a motivating reason to dislike Balkaners; and so on. But on the account I develop, they have the purpose or function to behave in these ways only if they are under the authority of others (Bazargan-Forward 2022, 35-53). Bernstein suggests that this is indeed what's going on; in each case the protagonist is responding to "authoritative" reasons. As such, the protagonist in each case qualifies as an *executor*. Who, then, is the deliberator? Excepting Meena, each protagonist's purpose isn't conferred by any discernible or identifiable individuals, but rather by social structures or institutions. Bernstein suggests that my account of authority-based accountability misses something important in overlooking these sorts of cases. (Though I briefly consider the possibility that *group agents* might confer purposes (Bazargan-Forward 2022, 33, 95-6), these comprise only a subset of the relevant cases; in most instances the conferring social entity will not qualify as an agent).

I take her point, though there is a reason why I focus on cases where "there is a clear chain of transmission from the deliberator to the executor"

(Bernstein 2024, 45). Normally, to fix the content of an executor's purpose, we defer to the deliberator's motivating reasons. But the "deliberator" in the majority of Bernstein's cases are social structures and institutions. Do these entities have anything analogous to motivating reasons serving as a basis for fixing the content of the purposes they confer upon executors? We might deny that in such cases there is any division of agential labor to speak of, since by hypothesis there is no agent serving as a deliberator.

But this move would impoverish my account by excluding from it social structures and institutions as grounds for conferring purposes. For this reason, I am inclined to extend my account by invoking a functionalist analysis of motivating reasons thin enough to accommodate the relevant social structures and institutions as bearers of such reasons. The account, extended in this way, can help make sense of purposes conferred in cases of the sort Bernstein has in mind. The relevant social structures and institutions, as deliberators, will then bear authority-based responsibility for a wrong-making feature of what their executors do in accordance with the purposes conferred upon them. But what are the implications of attributing responsibility to groups qua groups? At this point, I would invoke the work of others to help answer this question.<sup>2</sup> The upshot is that I am largely sympathetic to Bernstein's arguments in favor of expanding my account to accommodate social structures and institutions.

Bernstein then turns to the role that moral luck plays in my account. The account over-generates cases, she suggests, "in which a deliberator furnishes a purpose for an executor, but the executor turns out not to be causally responsible for an outcome." She has in mind cases where "[i]t is bad moral luck" that these outcomes occur. The problem, she says, is that the account I develop "divorces causal contribution from moral responsibility in a worrisome way" (Bernstein 2024, 47).

Crucially, though, the account I develop is *supposed* to be non-causal, in that its purpose is to reveal one way in which some can be non-causally accountable for elements of what others do. It isn't enough, then, to cite the account's conclusion as a basis for rejecting it wholesale. Bernstein says "[i]ntuitively, causal contribution to an outcome is a condition of being held accountable for that outcome" (Bernstein 2024, 47). But there are familiar cases in which individuals seem accountable for harms despite the absence of any causal connection to it. The book's introduction presents a panoply of such cases, which the account I develop helps accommodate.

---

2 See for example Isaacs (2011); List and Pettit (2011); Collins (2023), among others.

Still, Bernstein worries that my commitment to non-causal accountability forces me to deny the moral relevance of resultant moral luck. The same goes for the moral relevance of deviant causation; Bernstein worries that my commitment to non-causal accountability forces me to deny the moral relevance of deviant causation. It's true that, on the account I develop, resultant luck and deviant causation are not morally relevant when the basis of accountability is non-causal. The account I develop makes no claims about the moral relevance of resultant moral luck where the basis of accountability is causal. The upshot is that, contrary to what Bernstein suggests, my account doesn't force me to deny wholesale either the moral relevance of resultant moral luck or the moral relevance of deviant causation. If *all* instances of accountability were non-causal, Bernstein's worry would be valid. But I do not make that claim.

How then do we draw a principled distinction between Mastermind 1 (M1) and Mastermind 2 (M2), "given that they each confer a purpose onto a causally sufficient executor"? (Bernstein 2024, 47) The thought here is this: if in both cases the basis of accountability is at least partly non-causal, why is M1 on the hook but not M2? Though the *basis* of a deliberator's accountability is indeed non-causal, the *object* of her accountability is what her executor does qua executor. M1's assassin committed murder; M2's assassin merely attempted murder. Thus, M1's accountability is accountable for murder whereas M2 is not.

Bernstein then turns from cases in which the deliberator is accountable for what the executor does, to cases in which the executor is accountable for enabling a deliberator who harbors malicious aims. She considers cases in which executors are attempting to "change the system from within" (Bernstein 2024, 50). She worries that my account will yield the result that such individuals are acting wrongly. But authority-based accountability constitutes only part of morality in that it yields only *pro tanto* reasons for executors to refrain from enabling deliberators harboring malicious aims. These reasons can be outweighed by reasons to promote good outcomes. And when they do, my account captures the moral complexity of the situation, in that an executor attempting to change the system from within might nonetheless be said to have "dirty hands." One way of making sense of this common reaction is by noting that what the executor does, though all-things-considered permissible, includes this wrong-making feature: her actions retain an inimical purpose conferred by her deliberator. In this way, the executor might feel that her integrity is compromised even though what she did was praiseworthy. I discuss this phenomenon in the book's concluding chapter.

Bernstein ends by noting that motivating reasons can be shaped by constitutive and circumstantial luck—but this is of course true in cases of ordinary responsibility. That being said, Bernstein is right that my account, as much as any other, must struggle with the relevance of moral luck.

## 6. EVRENSEL SEBEP

The book's second half applies the argument for authority-based accountability to a range of issues and topics, including war. In his article, Evrensel Sebepe addresses this attempt.

Sebepe notes that, given, authority-based accountability some combatants will be liable to greater harms than others, depending on the roles they occupy and the situations they're in. Yet it seems that combatant liability should "maintain a binary nature" in that "a combatant's right not to be defensively killed is either forfeited or not." Such forfeiture, Sebepe says, "cannot be a matter of degree" (Sebepe 2024, 57).

It is true that on my account, combatant liability will vary; some will be liable only to sub-lethal harms. But even such a result helps resolve the "responsibility dilemma" which Sebepe helpfully articulates. This is because liability to substantial but ultimately sub-lethal harms makes it easier to meet a lesser-evil justification for killing. Suppose you can save the lives of three innocent persons but only by killing one other person—call him "J." Normally it is impermissible to kill one non-liable individual in order to save three. But suppose J bears *some* accountability for the threat the five face. As a result, he is morally liable to lose his legs, but *not* morally liable to be killed. J's liability is relevant to the lesser-evil calculation; as a result, killing him might be permissible after all. I call this a "hybrid" justification for killing.<sup>3</sup>

Suppose, then, authority-based accountability fails to ground lethal liability for all combatants. Instead, many are morally liable only to substantial sub-lethal harms. Even so, this helps explain why killing them will typically be permissible after all. Moreover, it does so in a way avoiding the unpalatable conclusion that killing typical civilians is similarly permissible (since they do not bear authority-based accountability for what combatants do in war).

This also helps resolve a related worry Sebepe raises. It would be prohibitively difficult, he says, to calculate and assign varying degrees of liability in the midst of a complex and dynamic environment such as warfare. But on my account, we generally needn't distinguish between (a) the combatants who are liable to be killed and (b) the combatants who are liable to substantial but sub-

<sup>3</sup> I make this argument in much greater detail in Bazargan (2014).

lethal harms. This is because *both* can be permissibly killed, albeit for different reasons. Still, a liability-based justification for killing is importantly different from a hybrid justification for killing. Combatants in category “b,” but not those in category “a,” are wronged by being killed. We can imagine cases where the number of combatants liable only to sub-lethal harms is so large, and the good done by killing them is so small, that such killings are impermissible after all, on the grounds that they don’t do enough good to warrant wronging so many. Hence, I believe Sebepe is correct in suggesting that my account faces a practical challenge, though not one quite so dire as what he envisions. For any given war, we generally won’t need to distinguish on a case-by-case basis those liable to lethal violence from those liable to sub-lethal violence. But we will need to ascertain the ratio between the two, as well as the point at which a hybrid-justification for killing fails. The epistemic difficulties endemic to this kind of moral calculus are par for the course in an activity as complex and fraught as warfare, especially within an overall revisionist framework.

Because liability will depend on the group in which one is participating, moral luck will play a role in determining who is liable for what. Sebepe suggests this is a problem: he envisions a scenario in which some combatants are assigned to combat operations while other are assigned to support roles. The severity of the harms to which each combatant is liable might depend on the group to which each is assigned, as a matter of luck. Sebepe is describing circumstantial moral luck, which undercuts authority-based accountability no more than it does “ordinary” accountability. Setting aside authority-based accountability, imagine a mob boss who arbitrarily chooses one of his men to steal a pizza and another to commit a murder. The fact that it might have been the other way around doesn’t obviously affect either’s accountability for what he does. Here, circumstantial moral luck seems to matter; ditto for my account and its application to war.

Sebepe goes on to note that on my account combatants occupying support aren’t be liable for the war *in toto*. (This is also true for combatants who are indeed engaged in combat operations). But the overall effectiveness of the military, Sebepe points out, relies on the contributions of support personnel. It seems, then, that the account I develop would eliminate or substantially mitigate the liability of support personnel who contribute substantially to an unjust war.

But this doesn’t mean that on my account such combatants cannot be permissibly killed. They furnish wrongful purposes for one another: the purpose of contributing to a war for which there is insufficient moral rationale to fight. As a result, each is accountable for a wrong-making feature of what

each of their cohorts do; thus each is liable for more than what he or she causes. This account raises a host of questions, of course. How can equally ranked combatants serve as co-deliberators for one another? Why believe that alienated combatants are conferring war-related purposes upon one another? I address these issues in the book (Bazargan 2014, 151–9). The point here, though, is that the account I develop presents reasons for thinking that support personnel are individually liable (for more than the difference each makes) by dint of their participation in a cooperatively committed harm. That being said, Sebepe is right, I believe, in suggesting that I must do more to show that these purposes which the combatants confer upon one another make *enough* of a difference to liability—enough to ground a hybrid-based justification for killing them.

Turning from combatants to noncombatants, Sebepe points out that authority-based accountability assigns liability for unjust wars to a state's leadership, which might include civilian leadership. *If* a state is waging a war that lacks a just cause, and *if* that war is authorized by the civilian leadership, and *if* the targeting the civilian leadership is effective in bringing the war to a close, the leadership can, in principle, be liable to such an attack. This is because individual civilian leaders authorizing such a war will bear much more accountability for that war than individual combatants do. Sebepe suggests that by admitting as much, I am “bit[ing] the bullet” (Sebepe 2024, 59). But it seems to me plainly obvious that, *ceteris paribus*, those who authorize an unjust war are morally preferable targets to those who fight such a war.

But suppose the war is authorized by a parliament (or other decision-making body) some members of which spoke out and voted *against* the war. Much has been written on the accountability of the minority in a majoritarian decision-procedure. Suffice it to say that even given authority-based accountability, it's not obvious that the dissenting members of government would be liable for the unjust war which they voted against. The upshot is that the account I develop does not entail the view that the civilian leadership in an unjust war is necessarily liable; much will depend on how the leadership is constituted.

Sebepe worries, though, that on my account, even ordinary citizens will end up liable for unjust wars. This is because, he suggests, individual citizens function as deliberators and their government function as executors. So, those who vote for a party or candidate promising to wage an unjust war would bear authority-based accountability for what that party or candidate does.

But this gets things backward; by electing a government, the citizenry *divests* itself of authority over its armed services. There is a crucial difference between standing atop a chain of command, and deciding who stands atop

that chain of command. The electorate has no authority over the military, even though it decides who does. This doesn't mean that the electorate bears little or no accountability for the government it puts in power; but it does mean they bear little or no *authority-based* accountability, in cases of the sort Sebeq considers. The same goes for lobbyists—another class of civilians Sebeq considers. Though lobbyists who convince politicians to wage an unjust war might be accountable for the foreseeable difference they make, they bear no authority-based accountability for that war. The moral here is that though authority-based accountability might in principle yield civilian liability for unjust wars, it does so only in those rare cases where such liability seems apt in any case.

#### REFERENCES

- Bazargan, S. 2014. "Killing Morally Responsible Threats." *Ethics* 125 (1): 114–136, URL <https://doi.org/10.1086/677023>.
- Bazargan-Forward, S. 2022. *Authority, Cooperation, and Accountability*. New York: Oxford University Press, URL <https://doi.org/10.1093/oso/9780192862419.001.0001>.
- Bernstein, S. 2024. "Reasons and Luck: Comments on Bazargan-Forward's Authority, Cooperation, and Accountability." *Journal of Social Ontology* 10 (4): 41–51, URL <https://doi.org/10.25365/jso-2024-9149>.
- Blomberg, O. 2024. "Authority-Based Accountability is Causally Extended Accountability: On Saba Bazargan-Forward's Authority, Cooperation, and Accountability." *Journal of Social Ontology* 10 (4): 8–19, URL <https://doi.org/10.25365/jso-2024-9146>.
- Bratman, M. 1993. "Shared Intention." *Ethics* 104 (1): 97–113, URL <https://doi.org/10.1086/293577>.
- Collins, S. 2023. *Organizations as Wrongdoers*. Oxford: Oxford University Press, URL <https://doi.org/10.1093/oso/9780192870438.001.0001>.
- Isaacs, T. 2011. *Moral Responsibility in Collective Contexts*. New York: Oxford University Press, URL <https://doi.org/10.1093/acprof:oso/9780199782963.001.0001>.
- List, C., and P. Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. New York: Oxford University Press, URL <https://doi.org/10.1093/acprof:oso/9780199591565.001.0001>.
- Payton, J. D. 2024. "Divided Agency, Manipulation, and Regret." *Journal of Social Ontology* 10 (4): 20–29, URL <https://doi.org/10.25365/jso-2024-9147>.



Sebep, E. 2024. “Welfare and Authority-Based Accountability.” *Journal of Social Ontology* 10 (4): 52–62, URL <https://doi.org/10.25365/jso-2024-9150>.