

Précis of *Authority, Cooperation, and Accountability*

Saba Bazargan-Forward

Imagine a group of individuals cooperating in a way resulting in a collectively committed wrong. They might be employees working for a corporation, soldiers fighting in a war, mobsters robbing a bank. Some of the cooperants contribute a lot. Others contribute little. Some fail to contribute anything at all. How do we make sense of individual accountability in such cases?

We might think that each cooperant's accountability is limited by her causal reach. But this sort of account is infamous for its counterintuitive implications. The aim of *Authority, Cooperation, and Accountability* (2022) is to develop an alternative account. I make the case for thinking that distinct aspects of human agency, normally "wrapped up" in a single person, can be "distributed" practically across different people. We "distribute" agency routinely, by forming promises, by making requests, by issuing demands, and by undertaking shared action. This resulting division of agential labor makes possible a distinctive way in which one person can be accountable for the actions of another. I call this phenomenon "authority-based accountability" and argue that it helps make sense of individual accountability in cooperative contexts. In the book's first half, I develop the argument for this view. In the

10.25365/jso-2024-9145

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.

book's second half, I apply it to just war theory, accomplice liability in the law, the legal doctrine of respondeat superior, and institutional racism. Here, though, I focus on summarizing the argument itself.

When evaluating a wrongful act, we often must repair to the reasons for which it was committed. It is generally presumed that the *wrongdoer's* reasons are the ones to which we must repair. But I argue that such a presumption is mistaken. To evaluate a wrongdoer's act, we will often need to repair to the reasons that *someone else* took there to be in favor of that act. This happens when the wrongdoer's role is to execute the decisions of someone whom the wrongdoer is tasked with helping. A dyadic example of this sort, where one individual helps another achieve a goal, will serve as a "building block" for more complex cases.

Olympic Sabotage. Contender is vying with Victim for first place in an Olympic competition. Contender consequently hires Goon to maim Victim in a way that will put her out of the running. Goon promises to do so and Contender accepts that promise. Goon subsequently maims Victim.

Imagine that after the assault, Victim demands that Goon explain why he attacked her. Goon might give *his* reasons for attacking her, of which there might be a variety: he wanted the money, he enjoyed the work, he liked Contender, he disliked Victim, and so on. But at some point, he might say, "Look, I attacked you because *I was supposed to*. If you really want to know why you were attacked, you're going to have to ask Contender." If Victim wants a full account of the purpose for which Goon acted, there is a sense in which she has to "look through" Goon's reasons to Contender's reasons.

Put less metaphorically, Contender *constitutively* determines at least one purpose of Goon's conduct. Suppose Goon does not know that Contender and Victim are competing for first place. Rather, he is just given a face and a name, and instructed to maim the target. Still, Contender has made it so that Goon has the purpose of facilitating Contender's victory in the upcoming competition. If Contender had some other reason for contracting the assault, Goon's purpose in committing it would *ipso facto* change accordingly.

Contender is thereby accountable for a morally relevant feature of what Goon does: a purpose for which the assault is committed. Of course, Contender is also accountable for what Goon does in a more familiar sense: by contracting the assault, Contender *causally* influences what Goon subsequently does. Contender presumably bears accountability for doing so. But Contender is also accountable for *constitutively* determining the purpose of Goon's actions.

This example demonstrates that morally evaluating one person's conduct will sometime require advertent to someone else's reasons. I call an individual who has the function of constitutively determining a purpose for which another person acts a "deliberator." And I call an individual who has the function of acting for a purpose constitutively determined by someone else an "executor." I analyze these dovetailing functions in terms of Razian protected reasons: the purpose the deliberator furnishes for the executor yields a Razian protected reason to act accordingly (Raz, 1979). Put differently, the executor has practical reasons to enact the deliberator's purpose—but at least some of those practical reasons derive not from the merits of the purpose itself, but rather from the fact that the executor is supposed to enact that purpose. In elucidating this relationship between the deliberator and the executor, I liken them functionally to the relationship you bear to your past self in cases of diachronic decision-making.

Where the purpose that the deliberator furnishes for the executor is morally bad, the deliberator can be accountable for that wrong-making feature of what the executor does (without thereby indemnifying the executor). This is authority-based accountability. So, in morally evaluating conduct, the agency that matters belongs not only to those we metaphysically attribute the conduct, but also to those whose ends the conduct has the function of enacting.

There is a lot of machinery in this gloss: various kinds of reasons (motivating, practical, moral, protected), as well as various telic concepts (intentions, functions, purposes), and still more besides. But the machinery is in service of what I hope is an intuitively compelling conceit: when I'm under your authority, and you have me harm someone, my victim might need to consult not just me but *you* for a comprehensive explanation of why *I* acted the way I did. In this way you can be on the hook for elements of what I did.

Authority-based accountability might seem otiose in dyadic cases like *Olympic Sabotage*. After all, ordinary accountability is enough to get us what we want: both Competitor and Goon are accountable for intentionally causing harm. But authority-based accountability helps us locate accountability in more complex examples, such as this one:

Toxic Dump. A couple dozen employees working for a contractor, Environmental Solutions, are tasked with safely disposing industrial waste. But they conspire to dump it in a river instead. No single cooperant's actions make a significant difference to the river's toxicity level. But in combination, their actions cause an environmental crisis for the local community.

By hypothesis, the cooperants implicitly agree to help one another in furtherance of dumping the waste in the river. The result is tantamount to a case in which each cooperant makes pairwise agreements with each and every other cooperant. Because of these agreements, each cooperant has authority over every other cooperant, in that each confers upon every other a protected reason to do her part. Should any cooperant fail to “pull her weight,” the other cooperants would have a basis for complaint.

So, in virtue of the agreement they made, each cooperant qualifies as both a deliberator and executor in $n-1$ pairwise agreements, where n is the total number of cooperants. The point can be put this way: in *Olympic Sabotage*, Goon acted at Competitor’s behest, whereas in *Toxic Dump* the cooperants are acting at one another’s behest.

The purposes which the cooperants confer upon one another depend on each cooperant’s motivating reasons. Suppose I’m one of the dozen cooperants working for Environmental Solutions. I thus serve as both a deliberator and an executor. I want everyone else on the team to help me dump the waste because it will line my pockets. I thereby furnish that purpose for every other teammate in my capacity as a deliberator. The result is that each cooperant has—wittingly or not—the following particular purpose, among others: to help me profit by illicitly dumping toxic waste into the river. Likewise, the other teammates severally furnish for me, in my capacity as executor, a concomitant purpose derived from each of *their* motivating reasons which might differ from mine. Some might cooperate in the scheme because it avoids travel across the state during the holidays. Some might cooperate out of peer pressure. But one way or another, each cooperant, qua deliberator, furnishes a purpose for me. Some of these ends might seem morally benign, such as avoiding harassment, or even morally commendable, such as spending time with family during the holidays. But each cooperant’s purpose is construed broadly, to include the means of achieving these ends: i.e., by dumping toxic waste.

Recall that a deliberator who furnishes a wrongful purpose for an executor is accountable for a wrong-making feature of what the executor does: a purpose for which she acts. The same goes for the employees of Environmental Solutions. Each cooperant, qua deliberator, is accountable for the purpose she furnishes for every other cooperant. If it is wrongful to dump toxic waste as a means of lining my pockets, then I am accountable for conferring that wrongful purpose upon the eleven other cooperants. So, even if I have little or no *causal* influence over what the others do, and even if I fail to contribute much to what we together do, I still have *normative* influence over what they are supposed to do and why they are supposed to do it.

How bad is it, though, to furnish for others a wrongful purpose to act in a way that they would have acted anyway? Suppose you are among the cooperants for whom I furnish a purpose in *Toxic Dump*. I'm not a but-for cause of the fact that what you do is wrongful. But the purpose I furnish is alone enough to make your conduct wrongful. The result is that if we *severally* rather than *jointly* evaluate what each cooperant does when she furnishes a purpose, it turns out that the purpose I furnish is wrong-making after all; *mutatis mutandis* for each of the other cooperants who furnishes a purpose for you. Proceeding in this way, every individual cooperant is accountable for making it so that every other cooperant's conduct counts as wrongful. The result is that even if I causally contribute nothing to the scheme (maybe the barrel I help shove into the river is empty) I *still* end up on the hook for a wrong-making feature of what the others do, by furnishing a purpose for which its done.

This summary raises a host of questions. What if one of the cooperants plays his role but denies that any of the others have authority over him, or denies that he has any authority over the others? What if some are ignorant of what it is exactly they're doing? What if some were coerced into participating in the scheme? I address these complicating issues, and others, in the book.

Authority-based accountability makes sense of accountability for what others do in the context of cooperation without invoking group-agents (List and Pettit, 2013), novel propositional attitudes (Searle, 1990), propositional attitudes with alternative "modes" (Tuomela, 2006; Schmitz, 2017), or propositional attitudes with irreducibly collective content (Kutz, 2000). This doesn't mean accounts invoking such concepts are mistaken. Rather, I believe that such concepts are not necessary in understanding, at the most fundamental level, why individuals can be accountable for what others do in the context of cooperative action.

Authority-based accountability has many practical implications, which I explore in the book's second half. I begin with war ethics. Suppose you're a soldier fighting an unjust war. But you occupy a highly subsidiary, marginal role, contributing almost nothing to the war effort. The "revisionist" turn in war ethics has difficulty grounding your liability, since it's determined solely by reference to your individual contributions to the war's aims. I explain how authority-based accountability helps address this challenge.

Turning to the legal doctrine of accomplice liability, suppose you're a participant in a bank robbery, but you end up contributing virtually nothing to its success. I argue that your complicity in such a case is best construed as a species of vicarious liability; you are vicariously liable for the wrong your partners commit because you bear a special, formal relationship to them: you

serve as both deliberator and executor within a division of agential labor.

I then address the legal doctrine of *respondeat superior*. Suppose you're an employee who abuses his position in a way harming a client. Why believe that your boss is accountable? I argue that *respondeat superior* is best construed in terms of authority-based accountability. Your boss, serving as a deliberator, implements duties of care owed toward her clients by conferring such duties upon you. Should you as an executor violate these duties of care, your boss thereby does so as well.

Authority-based accountability also helps us make sense of individual accountability in the context of institutional racism. Suppose your boss adopts policies perpetuating racial injustice. You, however, keep your hands clean: your actions and attitudes, as well as their effects, are morally unproblematic. But this doesn't mean you're off the hook. Given authority-based accountability, the moral status of your actions depends in part on your boss's motivating reasons.

I end the book by considering the implications of authority-based accountability for Bernard Williams's notion of integrity. In discussing his famous example of George the chemist, I argue that the attack on his integrity lies not simply in causally contributing to an end he finds morally abhorrent. In addition, George accepts an alien purpose, antithetical to his values. The upshot is that the divisions of agential labor in which we participate are crucial not just in assessing accountability for cooperatively committed harms, but in reckoning with the demands of integrity.

REFERENCES

- Bazargan-Forward, S. 2022. *Authority, Cooperation, and Accountability*. Oxford: Oxford University Press, URL <https://doi.org/10.1093/oso/9780192862419.001.0001>.
- Kutz, C. 2000. *Complicity: Ethics and Law for a Collective Age*. Cambridge: Cambridge University Press.
- List, C., and P. Pettit. 2013. *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>.
- Raz, J. 1979. *The Authority of Law*. Oxford: Clarendon Press, URL <https://doi.org/10.1093/acprof:oso/9780198253457.001.0001>.
- Schmitz, M. 2017. "What is a mode account of collective intentionality?" *Social Ontology and Collective Intentionality: Critical Essays on the Philosophy of Raimo Tuomela with his Responses*, edited by G. Preyer, and G. Peters, Springer, 37–70, URL <https://doi.org/10.1007/978-3-319->

33236-9_3.

Searle, J. 1990. "Collective Intentions and Actions." *Intentions in Communication*, edited by P. R. Cohen, J. Morgan, and M. E. Pollack, Cambridge, MA: MIT Press, 401–415.

Tuomela, R. 2006. "Joint Intention, We-Mode and I-Mode." *Midwest Studies in Philosophy* 30: 35–58, URL <https://doi.org/10.1111/j.1475-4975.2006.00127.x>.