

Who Does Wrong When an Organisation Does Wrong?

Stephanie Collins

This is the penultimate draft of a chapter whose final and definitive version will be published in:

Kendy Hess, Violetta Igneski, and Tracy Isaacs (eds), *Collectivity: Ontology, Ethics, and Social Justice*. Rowman and Littlefield.

Every time I open a newspaper, I read about organisations doing wrong: a police force violently prevents people from holding an outlawed vote; a national military razes the houses of an ethnic minority; a car manufacturer fiddles its emissions data. Often, it's obvious that some (if not all) members of the organisation have done wrong—and it's obvious who at least some of these wrongdoers are. But how far does this extend throughout the organisation? Have those members who remained silent—while not directly enacting the wrong—also done wrong, in the sense of having performed a blameworthy act?

Other cases are even more difficult. Think of systematic discrimination in a company's hiring processes, or in a state's policies. These wrongdoings are temporally and spatially dispersed. In such cases, it might be clear that some members did wrong in the past.¹ But it is much more difficult to pin down which *current* members—if any—are doing, or have done, wrong. Current members are my concern here.

Of course, there are sometimes current members who are doing the company's hiring, or implementing the state's policies. But it's not clear that even these individuals are doing wrong. First, if these individuals didn't implement the organisation's decisions, someone else would: the organisation-level wrong is *multiply realisable* and it's *overdetermined* that it will occur in some realisation or other. Second, the wrong is *explained* by organisation-level decisions, not by decisions of current members. Third, any harms that current members cause

(if any) are *tiny*: such microscopic harms are not obviously wrongful or blameworthy, even if those microscopic harms add up to a huge (and wrongful) macroscopic harm.

These excusing factors arise in cases where the organisation's wrongdoing is contemporary or ongoing. But historical cases give us an even clearer sense that present-day members are off the hook. In historical cases, the wrong has occurred *already*, under an entirely different membership from the present one. So there are no present-day members who are implementing the organisation's decision, even in an overdetermined, organisation-explained, and individually microscopic way.

So: what wrong, if any, is committed by present-day members when an organisation is doing wrong in the present or has done wrong in the past? This chapter tackles this question. I will assume that, if an organisation does (or did) wrong, then each current member is part of an entity that does (or did) wrong. But it does not follow that each current member does (or did) wrong. Analogously, when I cover my face with my hands, each part of me—whether present or future—is part of an entity that has covered its face with its hand. But it does not follow that each part of me has covered *its* face with *its* hands—or even that each part me (my left foot, say) has covered my face with my hands. The first section of this chapter explains why the actions (and, therefore wrongdoings) of organisations are distinct from the actions (and, therefore, wrongdoings) of members: I will outline how organisations' actions (including wrongdoings) are *multiply realisable* in members' actions, *explain* some actions of members, and are *emergent* over and above the mere conjunction of members' actions.

Because of this ontological distinctness, we cannot simply assert that all current members do wrong whenever an organisation does wrong. Indeed, the ontological distinctness raises the possibility that an organisation can do wrong, even while *none* of the members does wrong. The second section analyses such an example: the much-discussed 'discursive dilemma'. Philosophers argued that this is a case where an organisation did wrong *and*

members did wrong. I show that these attempts to implicate members will not always work: if these solutions are all we have, then there are still some cases where an organisation does wrong without any current members doing wrong. And even when these solutions work, they will often blame only a tiny portion of members.

But this result is deeply unsatisfying. As members of wrongdoing organisations, we feel ourselves on the hook. When my organisation (state, university, club) does wrong, I feel implicated—even if there is nothing that I could have done to alter its course. The distinctness of organisations' actions means that my membership alone does not vindicate this feeling. If the feeling is to be vindicated, we must appeal to something more.

The third section does just that. The general proposal is this. When an organisation does wrong—label this wrong 'W₁'—this generates obligations for members to disavow the wrong. It generates these obligations for the following reason: when a member fails to act *within and because of her organisational role* to disavow the organisation's wrongdoing, she makes it the case that the organisation itself 'pro tanto' ('to some extent' or 'in one respect') fails to disavow its own wrongdoing. That is: when a member acts within and because of her organisational role, the organisation performs a 'pro tanto' action, or acts to some extent or in one respect. Likewise, when a member *fails* to act within and because of her organisational role, the organisation fails to act, in one respect. So, if a member fails to act (within and because of her role) to disavow W₁, then *in one respect* the organisation also fails to disavow W₁. It would be wrong for the organisation to fail (in one respect) to disavow W₁; so, it is wrong for the member to perform an act that constitutes this organisational failure. If a member fails to disavow W₁, then this is a wrong done by the member (call this wrong 'W₂'), where W₂ constitutes a new pro tanto wrongdoing of the organisation itself. Call this new wrongdoing at the organisation-level 'W₃'. The member's failure to disown W₁ constitutes an organisational failure to disown

W₁. Her obligation is to prevent W₃, and she herself does wrong W₂ if she does not discharge that obligation. The proposal is depicted in Box 1.² <Box 1 near here>

To be clear, then, the ‘when’ in my title question does not mean ‘contemporaneous with,’ but rather ‘as a result of.’ This is the same sense of ‘when’ that is involved when we ask ‘what happens when a ball is thrown up?’ and accept the answer ‘when a ball’s thrown up, it falls down.’ That is: W₂ is *set in motion by* or *a result of* W₁. It’s not that W₁ *entails* W₂ and W₃—as I’ll explain, a member can halt the process that has been set in motion. This is just as we can halt a thrown ball’s falling down, by catching it in a net. But—just like a ball-catcher—the member must *act* to prevent W₂ and W₃.

So much for the relation between W₁ and W₂. The relation between W₂ and W₃ is constitution. This builds on the following: when a member acts within and because of her role, then this action constitutes an action of the organisation. Thus, when a member fails to act within and because of her role to disavow a wrongdoing, the organisation *pro tanto* fails to disavow that wrongdoing. The failure ‘transfers upwards’ from the member to the organisation, to make it the case that the organisation—*pro tanto*— fails to disavow its own wrongdoing. A member does wrong when she allows such upward-transfers. The way for a member to prevent such upward-transfer is to act within-and-because-of her role in the organisation, with a view to *disavowing* W₁. The answer to my title question is therefore “all those members who don’t act within and because of their role, to disavow the organisation’s wrongdoing.”

In the final section, I apply this to citizens of liberal democratic states. The claim will be that when a citizen use their role to disavow unjust policies—for example, through protest, petition, or public condemnation—this citizen makes it the case that the state has (in one respect) disavowed its own unjust policy. Those citizens who *do not* disavow the unjust policy have made it the case that the state has (in one respect) failed to disavow its own wrongdoing. This might sound contradictory, but I will argue that it is not. This result supports the common-

sense intuition that, when such states do wrong, wrong is also done by all and only those citizens who do not disavow the wrongful policy.

1. Organisations' Distinct Wrongdoings

Organisations occupy a distinct place in social ontology; they're different from crowds or mobs or collections of people in subway cars. Following Geoffrey Hodgson, I will assume that "[i]nstitutions are systems of established and embedded social rules that structure social interactions. *Organizations* are special institutions that involve (a) criteria to establish their boundaries and to distinguish their members from non-members, (b) principles of sovereignty concerning who is in charge and (c) chains of command delineating responsibilities within the organization."³ Properties (a), (b), and (c) amount to the specification of a structure.⁴ When members realise (instantiate) that structure, the organisation can, I assume, be identified with that instantiation of the structure.⁵

Each organisation possesses, and rationally operates, its own decision-making procedure. These are processes that allow the organisation to move from (1) goal-like inputs, combined with (2) judgment-like inputs, to make (3) intention-like outputs.⁶ The judgment-, goal-, and intention-equivalents in organisations are distinct from the judgments, goals, and intentions of members. For example, an organisation might hold the profit above all else, even though no member holds the profit above all else in his or her personal life.⁷ Such divergence—and distinct ownership—of an organisation's judgments, goals, and intentions from those of its members can arise, at least in part, because the organisation's method for arriving at, and using, these features is different from the method any member uses for arriving at, and using, their analogous features. For example, an organisation might establish its judgments, goals, and intentions via a majoritarian or committee-based procedure, while no member uses such a procedure to arrive at his or her own judgments, goals, and intentions.⁸

Organisations operate their distinctive procedures from what Carol Rovane calls a “rational point of view,”⁹ and they do so in a way that seeks to maintain that rational point of view. That is: the organisation acts from a particular web of judgments, goals, and intentions, and it updates these in order to maintain both the internal coherence of the web and to achieve the goals in accordance with the judgments. (Of course, organisations, like individuals, do not always succeed in these aims.) Organisations with identical procedures can be distinguished from one another by their distinct rational points of view: two organisations might have identical procedures, but nonetheless have different judgments, goals, and intentions.

These definitional remarks serve to limit the scope of my enquiry. Entities like ‘the wealthy’ or ‘humanity’ are not organisations, so my argument will not apply to them.¹⁰ I’m also not concerned with small, temporally limited, and teleologically focused groups, such as two people moving a tree,¹¹ painting a house,¹² or going for a walk.¹³ But my characterisation is neutral on many organisational quirks and characteristics: organisations might be more or less bottom-up or top-down; democratic or authoritarian; large or small; diverse or homogenous; long-term or short-term; and so on. I will assume that an organisation’s sharp boundaries, distinctive decision-making procedure, and unique rational point of view justify the claim that an organisation is an agent.

My focus is not just organisations’ *agency*, but their propensity for *wrongdoing*. If they can perpetrate wrongdoing, then organisations must be *moral* agents. A moral agent is one whose goals can be determined by moral reasons. Organisations can be moral agents. After all, if an organisation’s members are human moral agents, and so can recognise moral reasons, then it is natural to assume they can design an organisational procedure to which they can present moral reasons. And it is natural to assume they can design the procedure so it processes those moral reasons such that the organisation forms the goal of responding to those reasons. So some organisations are moral agents.¹⁴

There are (at least) three reasons to view organisations' actions—including their wrongdoings—as irreducible to (that is, non-identical to) the combination of members' actions.¹⁵ The first reason is multiple realisability. When a state wages a war, when a for-profit pollutes the environment, or when a non-profit does poor research, there are multiple ways in which the organisation-level action can be realised by individual actions. Consider the US's action of invading Iraq. There are numerous ways in which the various individual actions that make up this organisational action could go differently. A different commander, soldier, or logistician could have had this or that idea. This or that idea could have been replaced by a slightly different idea, though in the mind of the same individual. Individuals could have worked more or less callously, carefully, obediently, and so on.

Yet the invasion of Iraq exists across each of these realisations. So, we can't identify the implementation with any one realisation.¹⁶ Each of these realisations is 'The US invading Iraq,' just so long as two things are true: first, the relevant individual actions were performed by members within and because of their role in the US; second, Iraq was indeed invaded. There is something that unites all the possible individual-level realisations: in each case, the individuals were acting within, and because of, their role in an organisation that was governed rationally with a particular decision-making procedure that operated from a particular point of view; in each case, a significant macro-level event occurred.

The second reason for irreducibility is explanatory power. To adequately explain the members' actions that make up the organisation's action—to capture their motivations, constraints, goals, and cause—we need to refer to the organisation's decision-making procedure. If a soldier were to engage in weapons training by herself, for her own reasons, while off-duty, without using the US's organisational or material resources, then we would not be inclined to describe this as the US's action. This is because the US's point of view and decision-making procedures—with its distinctive judgment, goal, and intention-forming

methods, and ways of updating those—would not play any role in motivating, constraining, guiding, or causing the soldier’s actions. But if an adequate explanation or description of the soldier’s actions does require referring to the US’s point of view and decision-making procedure—for example, she did it because she was instructed, or because the organisation gave her access to the relevant resources—then we face pressure to view the soldier’s action as not merely her own (though it is that too), but also the action of the US. We might naturally say: ‘The US is engaging in weapons training.’ The more dependent the action was on the organisation’s decisions, procedure, and resources, the more pressure we face to endorse such a statement.¹⁷

The final reason for irreducibility is that organisations can do things their members taken in aggregate cannot do. The invasion of Iraq could not be performed by even a huge collection of individuals, if they were not coordinated by an organisation-level decision-making procedure. A number of legislators, commanders, soldiers, and so on—working *without* a clear division of roles and method for setting and dividing sub-aims—couldn’t control for the invasion of Iraq. Now, such a collection might cause Iraq to be invaded, by sheer fluke (say, if they each privately aimed at the outcome, without coordination). However, if it resulted from sheer fluke, then we would not describe this as the action of any one of them—or even as the action of the collection. It would be outside of their individual or collective control. By contrast, if the individuals do it under an organisation-level decision-making procedure, it’s not just a fluke. It’s organised and intentionally controlled.¹⁸ But it’s still not true that any one of the individuals performs the invasion on their own. Instead, the action of invading Iraq—if it is to be an action at all, which it seems it should, because it’s intentionally controlled—must be attributed to the organisation itself. Because the action could not come about from a conjunction of actions by unorganised individuals, the action is *emergent*: it arises

only when the individuals are structured just so, such that the complex relations between the individuals are crucial to grounding the action qua action.¹⁹

These three reasons together give us sufficient reason to think that at least some organisations can act irreducibly, at least sometimes. And organisations can perform their own, organisation-level, wrongdoings—for the same three reasons. But what does this imply for members? Which members, if any, do wrong when an organisation does wrong?

2. Limited Routes to Members' Wrongdoing

I will assume that, when an organisation does wrong, each current member is part of the entity that did the wrong. This is so even if the wrong was done in the past. This is just like how, when I grow a new hair, this hair is part of the entity that was once enrolled in my primary school. This doesn't mean the hair was there at the primary school—still less that the hair was enrolled in my primary school. But it is part of the entity that was so enrolled. Likewise, I am part of an entity—the state of New Zealand—that denied its native Māori population political representation until 1867. This doesn't mean I was there in (or before) 1867—still less that I personally denied Māori political representation before that date. But I am part of the entity that did so. This is also true of contemporary wrongdoing. Take New Zealand's ongoing failure to adequately address inequalities between its Māori and Pākehā ('of European descent') populations. This failure is committed by the state. The state is upheld through the tax contributions and the law-abidingness of citizens, of which I am one. But the state's failure is not my personal doing. The arguments of the previous section demonstrate that the state is an ontologically distinct entity. So we cannot automatically impute its wrongs to me.

Given all this: which members do wrong when an organisation does wrong? There are straightforward ways in which philosophers have tried to get at least some members on the hook. But even in temporally limited and small-scale cases, these straightforward methods will

sometimes capture no members at all. To demonstrate this, it will be cleanest to consider a toy example. Take the example depicted in Table 1.²⁰ <Table 1 near here>

There is a three-person tenure committee. Mr Borderline performs adequately at teaching, research, and service. He needs all three to deserve tenure. The committee votes on each in turn, and decides that he meets each component, because a majority of members vote that he does. But, once the committee votes on the question of tenure, Mr Borderline is denied tenure. This is because each committee member believes that Mr Borderline fail one of the three components, meaning no committee member believes he deserves tenure. The decisive vote is the vote on tenure, rather than the votes on the components of tenure. So, Mr Borderline is denied tenure. Assume this is an injustice to Mr Borderline: the tenure committee has done wrong (that is, performed a blameworthy action). Assume, further, that this wrongdoing could have been avoided if the members had been required (by the committee's constitutional rules) to discuss the components before voting. But the committee's decision-making procedure did not require members to do this. It seems here that no members did wrong, since they all voted in good faith and none had the power to amend the constitution to require pre-vote discussion. Yet the committee itself has wronged Borderline. So, which members—if any—do wrong?

Perhaps some members reasoned in a wrongful way when deciding how to vote. For example, perhaps they were insufficiently attuned to the evidence, where they owed it to Mr Borderline to be better-attuned.²¹ Or perhaps the members ought to have voted strategically, but failed to do so.²² If so, the relevant individuals have done wrong. However, we can imagine that the judgments were reasonable given the evidence (and gathering alternative evidence would have been impossible or unduly demanding),²³ or that the tenure committee's procedures forbid strategic voting. If so, no individual member did wrong.

Or perhaps the committee had "steering members" (i.e., members who design, direct, and can change the procedure), such that, for example, the group's duty to require pre-voting

deliberation translates into a duty for the steering member(s) to change the procedure so that it requires pre-voting deliberation. If so, steering members are straightforwardly on the hook for the group's harm.²⁴ Plainly, there will not always be such steering members, for example if the constitution of the tenure committee was decided decades ago, or by university managers who are not committee members.

Third, perhaps members control for the fact that group-level decisions and reasoning can depart from individual-level decisions and reasoning.²⁵ If this control was wrongfully exercised, then those who exercised that control did wrong when the committee did wrong. However, the discussion in the previous section gives us reasons to think the organisation itself (not the individual members) controls for the group's separation, decisions, and reasons. Even if members do exercise some influence over these things, it will often be that this influence is not exercised by *current* members.

All of this suggests that that there are some contexts in which the answer to my title question is 'no current members.' And there are numerous contexts in which the answer is 'only the leaders.'

3. A More Expansive Route

The above result does not vindicate the complicity felt by current members. When my state does wrong—even when I think of wrongs it has done wrong in the past, long before I was born—I feel implicated. Perhaps our philosophical accounts of wrongdoing will ultimately show that this feeling is mistaken. (Indeed, I will argue that, if I have performed certain disavowing actions, the feeling would be mistaken.) But I want to see if we can develop an independently plausible account of the member-organisation relation that vindicates the feeling, at least for many members and in many contexts—even when the organisation's wrongdoing has already happened, is multiply realisable, explains the members' contributions,

and is emergent; and even when the members reasoned well, were not steerers, and did not control for organisational distinctness. The relevant member-wrong might be one that results from the organisation's wrong, without being contemporaneous with it.

The secret, I believe, lies in taking seriously the following idea: when members act within and because of their organisational roles, the organisation acts. Thus, when a current member acts within her role, it's not just that she has some chance of *causing* her organisation to do wrongs, such as creating and sustaining (e.g.) wrongful tenure decisions or racialised inequalities. Such merely causal influence is often scuppered by the multiple realisability, explanatory power, and emergence of the organisation's wrongdoing. Instead, her wrong constitutes a new 'pro tanto' or 'in one respect' organisational wrong. The rest of the chapter develops these ideas and applies them to liberal democratic states.

3.1 Pro Tanto Bads and Overall Bads: The Individual Case

To bring my proposal into view, I will describe an individual case that will provide us with a helpful analogy to the collective case.

At the climax of the novel *The Adventures of Huckleberry Finn*, the main character, Huck, has written a letter to Miss Watson. The letter tells Miss Watson where she can find her runaway slave, Jim. Huck has helped Jim to escape. Huck firmly believes that he did the wrong thing in helping Jim escape, and that he should "see if I couldn't try to quit being the kind of a boy I was, and be better." By sending the letter to Miss Watson, Huck believes, he will "do the right thing and the clean thing." But then Huck reflects upon all the time he's spent with Jim—all the help they have given to one another, all the "talking, and singing, and laughing" and "how good he [Jim] always was"—and Huck feels conflicted about sending the letter. In the end, Huck resolves "'All right, then, I'll go to hell'" and tears up the letter, resolving to "go to work and steal Jim out of slavery again."²⁶

This example has been discussed extensively in the literature on virtue and praiseworthiness at the individual level.²⁷ I'm not concerned with those debates here. Whichever way we assess Huck's character, motivations, and reasoning overall, the following is true: Huck is good in a way and bad in a way. As I will equivalently say: he is 'pro tanto' good and he is 'pro tanto' bad. Huck himself puts this by saying "I warn't square; ... I was playing double. I was letting on to give up sin [by sending the letter], but away inside of me I was holding on to the biggest one of all [that he can't bring himself to send the letter]." (Twain 1884) Huck of course had it the wrong way around. In fact, he is good insofar as he correctly values Jim and their relationship; he is bad insofar as he has incorrect beliefs about what morality requires him to do. But the point both he and I want to press is that he is at once pro tanto good and pro tanto bad.

Bear in mind that Huck's bad beliefs don't *change* as he reflects upon his relationship with Huck: they simply get overridden by Huck's valuation of Jim and their relationship. This valuation is morally good, independently of whether it comes alongside the bad beliefs. Perhaps we want to judge that Huck is good *overall* (or, as I'll say equivalently, 'all things considered'), precisely because the personal feelings eventually override the false moral beliefs. And maybe it's even true that he's better *overall* than a possible Huck who had the good moral beliefs *and* performed the right overall actions (maybe it's better to do the right thing through struggle, rather than easily). Regardless of these overall assessments, the fact remains that Huck's incorrect moral beliefs are, *in themselves*, a bad-making feature of Huck. The point is that Huck (like all of us) has good aspects and bad aspects. He is pro tanto good and he is pro tanto bad.

Now imagine a possible betraying Huck. This Huck feels just as conflicted, in just the same ways, as the original Huck. But this betraying Huck ultimately decides to send the letter and help Jim no more. The betraying Huck, too, is good in a way and bad in a way: he values Jim (and his relationship with Jim) in the way he morally should, and he has false moral beliefs

about whether he should turn Jim in. So the betraying Huck is pro tanto good and pro tanto bad, just like the original Huck. The betraying Huck is, further, bad insofar as the second of these aspects prevails over the first. So he is bad at the overall level, just as the original Huck is good at the overall level. Still, it's true that both Hucks are pro tanto good and pro tanto bad—even though the original Huck is *also* good overall while the betraying Huck is bad overall.

Importantly, the pro tanto goods *and* the pro tanto bads of each of the Hucks remain, even after each has made his respective decision. That is: once the original Huck has resolved to help Jim, it remains true that he is bad insofar as he has incorrect beliefs. And once the betraying Huck has resolved to betray Jim, it remains true that he is good insofar as he has the right valuing attitude towards Jim. These morally significant pro tanto aspects—that are *not* reflected in the all-things-considered actions of the two Hucks—remain morally significant features of each of them. The original Huck is still bad pro tanto; the betraying Huck is still good pro tanto—even once the behaviour of each has been completely settled and is in the past.

3.2 Pro Tanto Wrongs and Overall Wrongs: The Organisation Case

I now want to suggest that members of organisations are to their organisations as the different aspects of Huck are to Huck—or at least, they are this way when they are acting within and because of their organisational role. There is one important difference, though: the various beliefs and values of Huck are not actions. So, the aspects of Huck are not *wrongdoings*. I assume that only actions are doings, so only actions are right-doings or wrong-doings. By contrast, organisations are composed of things (members) that can, themselves, act. So, an individual human like Huck can have pro tanto *goods and bads*, while organisations can additionally have pro tanto *rights and wrongs*. Just like the pro tanto goods and bads of a human, the pro tanto rights and wrongs of an organisation can be mutually in conflict.

My suggestion is that—even after the organisation’s decision has been made and enacted—members’ ongoing actions, beliefs, and preferences are aspects of the organisation. This is just as Huck’s conflicted beliefs and values are aspects of himself. So, when an organisation has done wrong, it is still a pro tanto good-making feature of that organisation that it has aspects (members) that oppose that wrong, and it is still a pro tanto bad-making feature of that organisation that it has aspects (members) that fail to disavow that wrong. This is just like how, when the good Huck has ripped up the letter, it is still a pro tanto bad-making feature of Huck that he has wrongful moral beliefs. Yes, this bad aspect has been ‘outvoted’ at the level of overall deliberation. Yes, the all-things-considered deliberation has already been enacted: the letter has been ripped up. But the bad-making feature persists nonetheless. In organisations, these bad-making features can be *actions* of the members, such that it makes sense to speak of the organisation’s *pro tanto wrongdoing*, even if it also has *pro tanto right-doings* and regardless of whether it does *the right thing overall*. In humans like Huck, the pro tanto aspects are not actions, but rather beliefs, desires, values, and so on, so the notion of ‘pro tanto wrongdoing’ does not so readily apply, though the notion of ‘pro tanto bad’ does.

To emphasise, the value of an organisation’s right or wrong pro tanto aspects does not lie in their somehow retaining the potential to reverse or change the organisation’s overall decision, or even to slightly alter the manner of that decision’s implementation. We see this by seeing that the positive value of good (or right) aspects remains, even after the overall bad (or wrong) decision has already been enacted. Again, suppose the betraying Huck has already posted the letter to Miss Watson. Suppose the time for giving tenure to Borderline has already passed. And suppose (as is the case) that New Zealand’s past wrongful treatment of Māori is done. Still, I suggest, it is a pro tanto good-making feature *of Huck* at the later time if he has parts (aspects, components) that hold values opposed to the betrayal. It is a pro tanto right-doing feature *of the committee* if it has parts (aspects, components) that oppose the rejection.

And it is a pro tanto right-doing feature of *New Zealand* if it has parts (aspects, components) that oppose the wrongful treatment.

Now, one might think there is a disanalogy between Huck and organisations: in Huck's case, everything that's done (believed, preferred) by some part of him is also done (believed, preferred) by him. In contrast, in the organisation case, a member might do (believe, prefer) something without the organisation itself doing (believing, preferring) that thing. So it's false that it's an action—even a 'pro tanto action'—of *the organisation* that members oppose an organisational wrongdoing after the fact.

I am not sure whether everything that's believed or valued by aspects of Huck are believed or valued by Huck simpliciter. But let's suppose that all actions and attitudes really do 'transfer up' from the aspects to the person, in the case of an individual. Allowing this, we face the problem: certainly not all actions 'transfer up' from the sub-organisational to the organisation: a member of the tenure committee might love her daughter, without it being true that the tenure committee loves her daughter, even pro tanto. So why think that the actions of members are pro tanto actions of organisations?

In answer: the only member actions that 'transfer up' to the organisation are the actions (and beliefs, values, etc) that are performed (held) by the member *within and because of her role in the organisation*. If—but only if—actions (and attitudes) are held in this way do they 'transfer up' from the member to the organisation itself, and thereby become the organisation's pro tanto actions (and attitudes).

We then face a second problem. There might seem to be a contradiction: the organisation can both do X (because a member does X within and because of her role) and do not-X (because a different member does not-X within and because of her role). But there is no contradiction. These are *pro tanto* actions, not overall actions. There is no contradiction in an individual having conflicting desires and beliefs: consider a vegan who desires to eat bacon,

while desiring not to harm animals. This analogous to Huck having (contradictory) beliefs and valuings, both of which are *pro tanto* aspects of him. The two contradictory aspects of Huck can (logically speaking) co-exist at the *pro tanto* level. We'd only get a logically troublesome contradiction if the contradiction persisted at the overall level. In neither the vegan case, nor Huck case, nor the organisational case, does the contradiction persist at the overall level. So there is no contradiction in the members' contradictory actions all 'transferring up' to become *pro tanto* actions of the organisations.

What does it mean for an action to be 'within and because of' a member's role? When a member performs an action (or has an attitude) *within* that member's role, this action (attitude) is one that the organisation's structure and decision-making procedure sanctions someone with one's role to have, where the organisation sanctions it because the action (attitude) mirrors, promotes, fulfils, or respects (or so on) the organisation's own actions (attitudes). When an action or attitude is held or performed *because* of one's role, it is held or performed precisely because (one believes) the organisation sanctions it as part of one's role.

These 'within' and 'because' conditions can come apart: I can perform an action that's *within* my role, without performing it *because* of my role. I might meet with students outside of my demarcated office hours, where this is sanctioned by the university as part of its goal of getting high student satisfaction scores. But I needn't do this *because* the university sanctioned it: we can imagine that I would make myself widely available regardless of the university's goals. Here, I act within my role without acting because of my role. Likewise, I can act because of my role without acting within my role. This is what I do if I strike over increased workload: my action is *caused by* what the university sanctions as part of my role, without my actions being *what* the university sanctions as part of my role. I will assume that, when I act *both* within *and* because of my role, my action is an action of my organisation. Likewise for attitudes: when

I believe or prefer *both within and because* of my role, my attitude is an attitude of my organisation.²⁸

Crucially, there are some attitudes and actions that members have or perform within and because of their roles, where the organisation does not dictate the precise contours of the attitude or action. For example, when I write papers within and because of my university role, my university (in a sense) writes papers. But the university doesn't say that I must write a paper on collective wrongdoing rather than anything else. Instead, it is up to me exactly which papers I write; exactly how I act within and because of my role. This leaves space for me to do wrong 'in my university's name'—for example, if I were to write an inaccurate paper on the unblemished good of New Zealand's treatment of Māori. It is still true that the organisation's action of 'doing research' is multiply realisable, explanatory of my action, and emergent. But I have scope to directly make the organisation do a pro tanto wrong—even if this does not translate into an overall wrong, say because the university apologises for my paper.

Bringing all this together, I can sometimes make it the case that my organisation does wrong. I do not (often) do this by *causing* overall and all-things-considered organisational wrong. Instead, I do this because my actions, performed within and because of my role in the group, are wrong—even if my group did right overall. These actions constitute pro tanto actions of the organisation, such that the organisation has pro tanto done wrong. This occurs regardless of whether these wrongful actions are reflected in the organisation's overall decisions. This is just as Huck is pro tanto bad (in virtue of his bad beliefs), regardless of whether his bad beliefs are reflected in his overall decision (that is, regardless of whether he ultimately decides to betray Jim).

In this way, I can make my organisation do wrong simply by acting—within and because of my role—to endorse an organisational wrong that is already 'done and dusted.' Call this 'done and dusted' wrongdoing, 'W₁.' This might be something like refusing a good tenure

candidate, or perpetrating racialised injustice. W_1 generates obligations for members to disown W_1 after the fact. It generates these obligations because, when a member acts within and because of her organisational role to *fail* to disavow W_1 , she makes it the case that the organisation itself *pro tanto* fails to disavow its own wrongdoing. If a member allows this failure to happen, then this is a new wrong, W_2 , which is done by her. W_2 in turn constitutes a new wrongdoing, W_3 , of the organisation itself. W_3 might not be done by the organisation all-things-considered. But a wrongful aspect (value, belief, action) of an entity is wrongful, even if that wrongful aspect is overruled or ignored at the level of the entity's all-things-considered decision about what to do. The member's obligation is to prevent this new organisational failure, and she does wrong if she does not discharge that obligation. The point of disavowing is not to resist, prevent, counteract, or compensate for W_1 .²⁹ That would require action at the all-things-considered group level, which members (often) cannot touch. But organisations' *pro tanto* beliefs and valuing can be valuable in themselves.

There are, of course, some caveats. The individual's obligation not to commit W_2 is not absolute and overriding. If her role is strictly defined, then the organisation's edicts fully determine which actions and attitudes of hers will count as 'within and because of her role.' There might then be no possible action that is *both* within-and-because-of-her-role *and* a disavowal of her organisation's wrongdoing. In such a case, she does not acquire the disowning obligation—and so, she does not do a wrong that results from W_1 (unless, of course, she has an active role in that wrongdoing, or gets 'on the hook' in one of the ways outlined in Section 2, or voluntarily joined (or remains in) the organisation knowing that it would (likely) do wrong). This is the right result: members who have absolutely no scope for resistance within their organisational role—and are not liable in one of the ways just bracketed—should not be on the hook when their organisations do wrong. But insofar as an organisation leaves a member discretion on how, exactly, she acts *qua* organisation constituent—and insofar as there is some

disowning action within that discretion—then the obligation to disown arises, as does the wrongdoing if she does not disown.

What about temporal distance? Is this an excuse? No. Assuming I have discretion within my role to disown W_1 , my failure constitutes a new organisation-level pro tanto wrongdoing, regardless of whether I perform this action before, during, or after W_1 . This is just as Huck's wrongful belief exists—and is bad—before, during, and after his decision to betray Huck. Such a belief would be bad even if Huck acquired it only once he was an old man.

4. Application to Liberal Democratic States

Let's return to New Zealand's wrongful treatment of Māori. There are two separate wrongs that play the W_1 role. First, there is the W_1 wrong of disenfranchising Māori until 1867. Second, there is the W_1 wrong of failing to adequately address ongoing inequalities between Māori and Pākehā. Both are multiply realisable; both explain the actions of members; and both are emergent (insofar as their controlled and predictable nature relied or relies upon the organisational structure). One is historically distant, while the other is contemporary. What does my framework say about which contemporary citizens have done wrong in each of these cases?

Following the solutions considered in the Mr Borderline case, some members will have done wrong in straightforward ways. First, some will have reasoned badly when deciding what their inputs into New Zealand's decisions would be. Second, some members are steerers. Third, some members control for the fact that many New Zealanders (most particularly, Māori) are not adequately represented in decisions about how to deal with racialised inequalities. With these solutions on the table, few current individuals are on the hook.

So who else is on the hook? My proposal is that all those members who do not act within and because of their role to disavow a state-committed wrongdoing thereby create a new

pro tanto state-level wrongdoing. Now, “acting within and because of one’s role to disavow a state’s wrongdoing” itself requires that disavowing a wrongdoing is something that a member *can* do within and because of their role. This is why I’m dealing with liberal democratic states only. I understand these to be states in which citizens have the right to vote, free speech, freedom of association, freedom of thought and conscience, and so on (except, often, children and criminals—but no one is suggesting that *these* individuals do wrong when the state does wrong).³⁰

In liberal democratic states, citizens are given discretion about how, exactly, to perform their role—where that role is “voicing their views—whatever those are—on how the state should be run.” There are numerous ways to voice one’s views, insofar as there are numerous views one can hold. So members have discretion over exactly how to fulfil their role (in particular, discretion over whether to use their role to disavow wrongdoings of their state), and where they do wrong if they do not use that discretion to disavow their state’s wrongdoings. If they do not use their role to do this, then they are rendering it the case that their state (pro tanto) does not disavow its own wrongdoing.

How much cost must a citizen bear in her disavowal? Is her private condemnation enough, or must the disavowal be public? To the second question, the answer is that the disavowal must be public—or else, she is not disavowing *within* and because of her role as citizen. One’s role as citizen does not extend to one’s private thoughts (at least in a truly liberal state), so using these thoughts to disavow will not be enough to make it the case that *the state* disavows.

On the question of how much cost the citizen must bear in her disavowal, I can provide no definite answer here—except to say that this is not a unique problem for the framework I propose. It is always vague exactly how much cost one is required to bear in the fulfilment of any duty (for example, must I sacrifice my finger if this is necessary to rescue a drowning

child? My whole hand?). In general, though, we can say that the worse the state's initial overall wrongdoing (W_1) is, the more cost a citizen must bear in her disavowal. This means that, for any state-level wrongdoing that is not egregious (genocide or ethnic cleansing, for example), the cost the citizen can be morally required to bear in her disavowal will not be extremely high. For egregious wrongs, the potential cost might be high. But even for such egregious wrongs, the question kicks in of how demanding morality is in general. There are general constraints on morality's demands, and these constraints will operate here.

The general constraints on morality's demandingness operate not just on the costliness of *specific* disavowing actions, but also on the costs an individual must bear across her whole life, to disavow the multitude of wrongdoings done by all her organisations. These considerations of demandingness limit the size and the number of W_2 wrongdoings (and, therefore, W_3 wrongdoings). If it would be overly costly for an individual to learn about W_1 , or if she is blamelessly a member of many wrongdoing organisations so that she cannot possibly disavow all the wrongdoings, then this will cancel some of her obligations to disavow. She will then not have done wrong for *every* wrong of *all* of her organisations. But the precise scale of these constraints is not something I can hope to measure here, since they are determined by morality's general demandingness and the relative importance of W_2 and W_1 as compared with other obligations of the individual. Certainly, if W_1 is easily discoverable, and disowning it is cheap, then an individual does wrong if she does not disown W_1 .

These considerations of demandingness raise another question: do citizen-victims of state injustice have a duty to disavow? Do Māori have a duty to disavow New Zealand's racialised injustice? In their capacity as state members, yesn. This accords with recent (though very different) arguments that victims have duties to resist injustice.³¹ It is likely, though, that considerations of demandingness—and of the fair distribution of the costs of organisations'

duties amongst their members—will mean that victims’ duties are far less onerous than those of non-victim members.

In sum, this account vindicates the following common-sense intuition about citizenship in liberal democratic states: when a liberal democratic state does wrong, wrong is done by all those citizens who endorse, condone, or acquiesce in the wrongful policy—but wrong is not done by those citizens who condemn the wrongful policy (assuming these condemners are not liable in one of the more straightforward ways). In cases of historical wrongdoing—and in cases where the state’s contemporary wrong is highly removed from ordinary citizens, such that they cannot change its overall ongoing decisions—this common-sense intuition is difficult to endorse. By seeing the actions of citizens—including endorsing, condoning, acquiescing, and disavowing—as constituting actions of the state, we can explain how those who do not disavow the wrongdoing do a distinctive kind of wrong. It is not the wrong of having contributed to the state’s past or contemporary injustice. It is instead the distinctive wrong of failing to make it the case that the state (*pro tanto*) disowns that past or contemporary injustice. Of course, common-sense does not recognise the idea of a *pro tanto* wrongdoing—and especially of an organisation doing contradictory things *pro tanto*. But these ideas can justify, and thus vindicate, the common-sense intuition spelled out above.

Despite this vindication of common-sense, one might wonder whether the disavowing citizen can really get herself off the hook so easily. Aren’t we *all* implicated in our states’ acts, whether we like it or not? Can it be so easy as attending a rally or signing a petition? After all, perhaps the state’s wrongdoing was committed as a reasonable means to, or constituent of, a permissible state action. The citizen might endorse that permissible action. If so, then perhaps she is ‘on the hook’ for all wrongful actions the state takes towards that permissible action, even if she acts to disavow those wrongful actions.

To fill in the blank on this ‘permissible action’ we might think, with Anna Stilz, that each citizen of a democratic legal state has committed to the end of formulating “just and nondominating laws that take everyone’s interests into account and to which each citizen has the chance to contribute.”³² Stilz argues that:

When a democratic legal state abuses its power ... even dissenters are liable for repairing harms their state inflicts. That is because even these dissenters require the framework of law that the state provides, if their rights are to be publicly defined and enforced. So long as state power is being exercised in a way that takes account of their fundamental constitutional interests, these dissenters, along with other citizens, authorize the institution.³³

Thus, “membership alone will suffice to implicate the citizen’s will.”³⁴ Or we might think, with Avia Pasternak, that citizens “share the common goal of living in a self-ruling political community.”³⁵ This amounts to intentional participation in the state’s overall activities, and “when citizens intentionally participate in their state, they gain authorship of their state’s activities”.³⁶ To be clear, these authors might agree that citizens sometimes have obligations to disavow their states’ unjust acts. But—because of citizens’ relation to their state as a whole and its overall political project—citizens cannot avoid nonetheless having authorship over the state’s unjust acts.

Now, having *authorship* over a state’s unjust acts is not the same as being *wrongful* or *blameworthy* in relation to those acts. Nonetheless, someone might build on (and depart from) Stilz’ and Pasternak’s views, in the following way: If we endorse the broad goal of, say, ‘governing our society democratically,’ then no amount of petitioning can justify us in claiming to have *fully* disavowed the components or implications of the state’s pursuit of that broad goal. The broad goal comes first, in the sense that our commitment to it trumps the normative force

of any disavowal of the means that the state takes to it. So, *all* members are on the hook for W_1 , even those who disavow it.

In response, I would make two points. First, this objection risks mis-describing the actions of states that most citizens support. That is, it is an open empirical question whether most citizens of a given liberal democratic state have a commitment to anything like the broad democratic goal. It is also an open question whether that commitment should be seen as foreclosing, or undermining, their disavowal of the various wrongful ways their state might pursue that goal. Whether the broad and permissible goal trumps (disavowals of) the specific and wrongful one is a question about their priority in citizens' minds, not to be settled by conceptual fiat. Second, even if citizens are committed primarily to a goal like 'governing ourselves democratically,' they can (and should) still disavow the actions of their states that are *not* reasonable means to, or components of, this goal. It is perfectly coherent for a citizen to say: "yes, let's govern ourselves democratically—but not like this! I disavow this!"

5. Conclusion

Who does wrong when an organisation does wrong? I have focused on members, particularly members with no influence over the organisation's overall wrong. Historical injustices are prime examples. But the fact that organisational wrongdoings are multiply realisable, explanatorily powerful, and emergent should lead us to think that these cases arise more often than we might think. What's more, these distinctness-supporting facts should lead us to conclude that it is not a simple conceptual truth that current members will have done wrong when an organisation does wrong.

On my proposal, it's not true that all current members do wrongs that follow from their organisation's wrongs. Instead, those current members do wrong who do not act within and because of their role to disavow the organisation's wrongdoing. They acquire an obligation to disavow because, if they do not, this will amount to a new organisation-level wrongdoing,

where the organisation (in one aspect, i.e., pro tanto) itself fails to disavow its own prior wrongdoing. If correct, this requires us all to be vigilant and vocal about our organisations' past, present, and potential future wrongdoings.

Acknowledgments

For feedback on drafts, I thank Jonathan Farrell, Luara Ferracioli, Anne Schwenkenbecher, and the Political Theory Seminar at University College London. A presentation of Section 4 received helpful comments from an audience at Princeton University.

References

- Arpaly, Nomy. *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford: Oxford University Press, 2002.
- Bennett, Jonathan. "The Conscience of Huckleberry Finn." *Philosophy* 49(188), 1974: 123–34.
- Björnsson, Gunnar and Kendy Hess. "Corporate Crocodile Tears? On the Reactive Attitudes of Corporate Agents." *Philosophy and Phenomenological Research* 94(2), 2017: 273–98.
- Braham, Matthew and Martin van Hees. "Responsibility Voids," *The Philosophical Quarterly*, 61(242), 2011: 6–15.
- Bratman, Michael. *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press, 2014.
- Chant, Sara. "The Special Composition Question in Action." *Pacific Philosophical Quarterly* 87(4), 2006: 422–41.
- Collins, Stephanie. "Distributing States' Duties." *Journal of Political Philosophy* 24(3), 2016: 344-366.

- Collins, Stephanie. "Duties of Group Agents and Group Members." *Journal of Social Philosophy*, 48(1), 2017: 38–57.
- Collins, Stephanie. "Filling Collective Duty Gaps." *The Journal of Philosophy*, 114(11), 2017: 573-591.
- Collins, Stephanie. Forthcoming. "The Government Should Be Ashamed: On the Possibility of Organisations' Emotional Duties." *Political Studies*.
- Collins, Stephanie and Holly Lawford-Smith. "We the People." Manuscript.
- Copp, David. "The Collective Moral Autonomy Thesis," *Journal of Social Philosophy*, 38(3), 207: 369–88.
- Cullity, Garrett. *The Moral Demands of Affluence*. Oxford: Oxford University Press, 2004.
- French, Peter. *Collective and Corporate Responsibility*. New York: Columbia University Press, 1984.
- Gilbert, Margaret. *On Social Facts*. London: Routledge, 1989.
- Hess, Kendy. "The Free Will of Corporations." *Philosophical Studies* 168(1), 2014: 241-260.
- Hess, Kendy. <this volume>
- Hindriks, Frank. "Corporate Responsibility and Judgment Aggregation." *Economics and Philosophy* 25(2), 2009:161–77.
- Hodgson, Geoffrey M. "Institutions and Individuals: Interaction and Evolution." *Organization Studies* 28, 2007: 97–116.
- Jackson, Frank and Philip Pettit. "Structural Explanation in Social Theory." Pp. 97-132 in D. Charles & K. Lennon (eds.), *Reduction, Explanation, and Realism*. Oxford: Clarendon Press, 1992.
- List, Christian and Kai Spiekermann. 'Methodological Individualism and Holism in Political Science: A Reconciliation.' *American Political Science Review* 107(4), 2013: 629–43.

- List, Christian and Peter Menzies. "Non-Reductive Physicalism and the Limits of the Exclusion Principle." *The Journal of Philosophy* 151(9), 2009: 475–502.
- List, Christian and Philip Pettit. *Group Agency*. Oxford: OUP, 2011.
- Mäkelä, Pekka. "Collective Agents and Moral Responsibility," *Journal of Social Philosophy* 38(3) 2007, 456–68.
- Pasternak, Avia. "Sharing the Costs of Political Injustices." *Politics, Philosophy & Economics* 10(2), 2011: 188–210.
- Pasternak, Avia. "Limiting States' Corporate Liability." *Journal of Political Philosophy* 21(4), 2013: 361–81.
- Pettit, Philip. "Responsibility Incorporated." *Ethics* 117, 2007: 171–201.
- Richie, Katherine. "What Are Groups?" *Philosophical Studies* 166(2), 2013: 157–72.
- Schwenkenbecher, Anne. <this volume>
- Stilz, Anna. *Liberal Loyalty: Freedom, Obligation, and the State*. Princeton University Press, 2009.
- Stilz, Anna. "Collective Responsibility and the State." *Journal of Political Philosophy* 19(2), 2011: 190–208.
- Strand, Anders. "Group Agency, Responsibility, and Control." *Philosophy of the Social Sciences* 43(2), 2012: 201-224.
- Vasanthakumar, Ashwini. "Epistemic Privilege and Victims' Duties to Resist their Oppression." *Journal of Applied Philosophy*, forthcoming.
- Wringe, Bill. <this volume>
- Zoller, David J. "Distributing Collective Moral Responsibility to Group Members." *Journal of Social Philosophy* 44(4), 2014: 478–97.

¹ Elsewhere, I argue that group agents' wrongdoings require that some individual did wrong, at some past time (Stephanie Collins, "Duties of Group Agents and Group Members," *Journal of Social Philosophy* 48(1) (2017): 38–57). In other places, I argue that we can justify imposing costs on current members regardless of their wrongdoing. (Stephanie Collins, "Distributing States' Duties," *Journal of Political Philosophy* 24(3) (2016): 344–66; Stephanie Collins. "Filling Collective Duty Gaps." *The Journal of Philosophy*, 114(11) (2017): 573–91) By contrast, my current aim is to demonstrate how *current* members *do wrong*.

² To be clear: while W_1 can occur without any present-day member doing wrong, W_3 relies on W_2 .

³ Geoffrey M. Hodgson, "Institutions and Individuals: Interaction and Evolution," *Organization Studies* 28 (2007): 97-116.

⁴ Kendy Hess. <This volume>

⁵ Katherine Richie, "What Are Groups?" *Philosophical Studies* 166(2) (2013):157-72.

⁶ Christian List and Philip Pettit, *Group Agency* (Oxford: Oxford University Press, 2011), 21.

⁷ Kendy Hess, "The Free Will of Corporations," *Philosophical Studies* 168(1) (2014):241-260.

⁸ The different decision rules organisations can use are outlined in List and Pettit, *Group Agency*, ch. 3.

⁹ Carol Rovane, *Bounds of Agency: An Essay in Revisionary Metaphysics* (Princeton: Princeton University Press, 1998).

¹⁰ The wrongdoings of these collectives are the concern of, for example, Garrett Cullity, *The Moral Demands of Affluence* (Oxford: Oxford University Press, 2004); Bill Wringe, <this volume>.

¹¹ Anne Schwenkenbecher, <this volume>.

¹² Michael Bratman. *Shared Agency: A Planning Theory of Acting Together* (Oxford: Oxford University Press, 2014).

¹³ Margaret Gilbert, *On Social Facts*. (London: Routledge, 1989).

¹⁴ Philip Pettit, “Responsibility Incorporated,” *Ethics* 117 (2007): 171-201; Gunnar Björnsson and Kendy Hess, “Corporate Crocodile Tears? On the Reactive Attitudes of Corporate Agents,” *Philosophy and Phenomenological Research* 94(2) (2017): 273–98.

¹⁵ The three reasons that follow adapt the discussion in Stephanie Collins, “‘The Government Should Be Ashamed: On the Possibility of Organisations’ Emotional Duties,” *Political Studies*, forthcoming.

¹⁶ Christian List and Kai Spiekermann, ‘Methodological Individualism and Holism in Political Science: A Reconciliation.’ *American Political Science Review* 107(4) (2013), 629–43. Christian List and Peter Menzies, “Non-Reductive Physicalism and the Limits of the Exclusion Principle,” *The Journal of Philosophy* 151(9) (2009):475-502. Frank Jackson and Philip Pettit. “Structural Explanation in Social Theory.” Pp. 97-132 in D. Charles & K. Lennon (eds.), *Reduction, Explanation, and Realism*. (Oxford: Clarendon Press, 1992).

¹⁷ For related causal-explanatory points, see List and Spiekermann, ‘Methodological Individualism,’ 637-40.

¹⁸ Anders Strand, “Group Agency, Responsibility, and Control,” *Philosophy of the Social Sciences* 43(2) (2012): 201-224.

¹⁹ Sara Chant, “The Special Composition Question in Action.” *Pacific Philosophical Quarterly* 87(4) (2006):422-441

²⁰ This example comes from David Copp. “The Collective Moral Autonomy Thesis,” *Journal of Social Philosophy*, 38(3) (2007):369–88, pp. 379-80. It is an instance of the ‘discursive dilemma’, discussed in List and Pettit, *Group Agency*, ch. 2.

²¹ Frank Hindriks. “Corporate Responsibility and Judgment Aggregation.” *Economics and Philosophy*. 25(2) (2009):161–77

²² Matthew Braham and Martin van Hees. “Responsibility Voids,” *The Philosophical Quarterly*, 61(242) (2011):6–15

²³ David Copp, “The Collective Moral Autonomy Thesis,” *Journal of Social Philosophy*, 38(3) (2007):369–88, p. 380.

²⁴ David J. Zoller, “Distributing Collective Moral Responsibility to Group Members.” *Journal of Social Philosophy* 44(4) (2014): 478–97, p. 490.

²⁵ Mäkelä Pekka, Collective Agents and Moral Responsibility,” *Journal of Social Philosophy* 38(3) (2007):456–68, esp. p. 463.

²⁶ Mark Twain, *The Adventures of Huckleberry Finn* (1884, available at <
<https://americanliterature.com/author/mark-twain/book/the-adventures-of-huckleberry-finn/chapter-31>>), chapter 31.

²⁷ See, for example, Jonathan Bennett, “The Conscience of Huckleberry Finn,” *Philosophy* 49(188) (1974):123-134; Nomy Arpaly, *Unprincipled Virtue: An Inquiry into Moral Agency* (Oxford: Oxford University Press, 2002)

²⁸ Admittedly, this is a large assumption. Peter French makes a similar suggestion for when individuals’ actions are redescribable as corporations’ actions (Peter French, *Collective and Corporate Responsibility* (New York: Columbia University Press, 1984) pp. 44, 55).

Elsewhere, I argue that all and only my actions that are within and because of my role can fulfil my organisation’s duties. (Collins, “Duties of Group Agents.”)

²⁹ I am therefore bracketing duties to resist injustice—understood as duties to try to *prevent* injustice.

³⁰ Are citizens of liberal democratic states members of states in the same way that committee members are members of a committee? I think the answer is ‘yes’, though I can’t argue for

that here. See Stephanie Collins and Holly Lawford-Smith, “We the People,” unpublished manuscript.

³¹ E.g., Ashwini Vasanthakumar, “Epistemic Privilege and Victims’ Duties to Resist their Oppression,” *Journal of Applied Philosophy* forthcoming.

³² Anna Stilz, *Liberal Loyalty: Freedom, Obligation, and the State* (Princeton University Press, 2009), p. 195.

³³ Anna Stilz, “Collective Responsibility and the State.” *Journal of Political Philosophy* 19(2) (2011): 190–208, p. 206.

³⁴ Stilz, “Collective Responsibility,” p. 208.

³⁵ Avia Pasternak, “Sharing the Costs of Political Injustices.” *Politics, Philosophy & Economics* 10(2) (2011):188–210, p. 199.

³⁶ Avia Pasternak, “Limiting States’ Corporate Liability,” *Journal of Political Philosophy* 21(4) (2013): 361–81, p. 374.