

Distributing States' Duties*

Stephanie Collins

This is the final draft of an article whose final and definitive version will be published in
Journal of Political Philosophy

We tend to think states have moral duties: duties to alleviate global warming, protect citizens' moral rights, admit asylum seekers, or wage only just wars.¹ This commonsense view accords with a growing philosophical consensus that states are corporate moral agents, able to bear duties as entities conceptually distinct from—though supervenient upon and constituted by—their members. States have clear membership rules and decision-making procedures that are distinct from the decision-making procedures of members. States are able to act on their decisions, through the actions that their decision-making procedures authorise members to take. States may therefore bear prospective and retrospective responsibility for their decisions and actions. In what follows, I will assume this view is sound.²

Yet problems remain to be solved, if we are to have a full conception of states' duties. One problem arises from the following observation: when states fulfil (many of) their obligations, costs are borne by members—that is, by individuals who are subject to the state's authority. Take the alleviation of global warming. The fulfilment of many obligations states have here—accepting climate refugees, curbing fossil fuel emissions, funding research into

* For helpful feedback, I thank Jonathan Farrell, Avia Pasternak, James Pattison, Liam Shields, two anonymous referees (one of which later identified herself as Anna Stilz), seminar audiences at the London School of Economics and the University of Leeds, and participants at a conference on 'Collectivity' at the University of Bristol. I am grateful to Lea Ypi for her editorial suggestions.

¹ The latter two are also international legal duties (Crawford 2007). I will use 'duty', 'obligation', and 'responsibility' synonymously, to mean 'moral reason with high weight or presumptive power'.

² For such accounts of group agency, see French 1984; List and Pettit 2011; Pettit 2007; Pettit and Schweikard 2006. For arguments that states bear such agency, see Erskine 2001; Goodin 1995, ch. 2; Wendt 2004.

green energy—will lead to costs for members—increased demand on public services, increased fossil fuel purchasing taxes, public funds being diverted to develop green energy, and so on. Some of these costs are direct means to, or components of, the discharge of the state’s duty, while others are incurred as an effect of the state discharging its duty. We can consider both kinds of costs together: for both, the discharge of the state’s duty requires that those subject to its authority lose out. And for both, the state has some control over the distribution of these costs amongst members. The question arises: how should these costs be distributed?³

This article develops a novel answer to this distribution question. To lay foundations, §I outlines a view on which moral agents—individual or collective—acquire duties from numerous and varied ‘sources.’ These sources include causing harm, having the capacity to bring about some good, benefitting from harm, being in an associative relationship, and so on. Different sources have different levels of importance in different contexts, and often several sources combine to generate a duty to perform a particular act.

§II uses this framework to answer the distribution question with regard to small-scale groups, in which implementation is relatively straightforward. It advocates the following answer: the distribution of members’ costs should, insofar as possible, track the source(s) of the group’s duty.⁴ As a first step, costs should be divided amongst duty sources, in proportion to those sources’ importance in generating the state’s duty. As a second step, the costs assigned to each source should be divided amongst members in proportion to their instantiation of that source. This is the ‘source-tracking’ proposal.

³ This question is about distributing amongst members the costs of *group-level* duties. It is thus distinct from more general questions about how to distribute benefits and burdens within society, or what duties members have to one another.

⁴ Other authors have briefly advocated this answer for groups’ wrongdoing-based duties, but none have developed it for states’ duties in general. See, e.g., May (1987, p. 100, 104-6) on criminal punishment of corporations.

§III applies the source-tracking proposal to states. It pays particular attention to the ‘shortfall’ problem, which is rife in long-term, large, and internally diverse groups like states. The problem arises when the group’s duty derives from sources that aren’t instantiated by any member. This is most acute for duties based in historical wrongdoing: perhaps the state did wrong decades ago, but no current member contributed to it, since none were born at the time. Here, an important duty source—perhaps the *only* duty source—doesn’t apply to any member. So costs cannot be distributed in a source-tracking way. The result is a shortfall: the state has a duty, but no individuals should incur costs in its fulfilment. This problem has led recent authors to reject source-tracking distributions for states’ wrongdoing-based duties.⁵

I address the shortfall problem by asking when and why shortfalls matter. I suggest they matter only when sources other than the shortfallen source are also generating moral reasons to for the state to act. Suppose a state bears a wrongdoing-based duty, but no members did wrong. The costs of the duty cannot, under source-tracking, be distributed to members. This shortfall will bother us when and because the state’s members nonetheless benefit from the harm, or have an associative relationship with the victims, or have a strong capacity to contribute to providing victims with a specific relational good. These are sources of duties in themselves, which are in operation alongside wrongdoing. My proposal is that some combination of other duty sources—perhaps benefiting, association, or capacity, amongst others—can be used to ‘top up’ the shortfall, under certain conditions that will be specified. Source-tracking is thus vindicated as a viable method for distributing the costs of states’ duties.

⁵ Stiliz 2011, p. 193-4; Pasternak 2013, p. 365.

I. Sources of Duties

There are several ways of connecting a moral agent, A, to a moral patient, P, such that A is singled out as having a duty to P. Some of these are: causal responsibility, where A contributed to P's situation in a mere causal way; moral responsibility, where A is to blame for P's situation; capacity, where A is especially capable of remedying P's situation;⁶ benefitting, where A benefits from P being in the situation;⁷ and community, where A stands in an associative relationship with P.⁸ We can view each of these connections as a source of moral reason for A to take action with regard to P. Each source arises in many contexts. But only in some contexts will one or more source be strong enough to trigger weighty or presumptive grounds for A to take action—that is, only in some cases will one or more source be significant enough to generate a duty.

Two observations are important for my purposes. First, different sources are weighty in different contexts. To adapt one of David Miller's examples: suppose one of a team of mountaineers becomes injured and must be helped down the mountain. If the injury is serious, perhaps the strongest team member should go with him, in case he needs carrying. Here, capacity is the most important source. If the injury is not serious, perhaps the team member most closely associated with him—his closest friend in the team, say—is the most appropriate duty-bearer. Here, association is more weighty than capacity. If someone caused his injury, then the injury-causer might be the most appropriate person to help, as recompense. And so on.

Second, multiple sources might apply in one context. If one teammate caused the mountaineer's injury, and she's his sister, and she's very strong, then three sources are at play:

⁶ Goodin 1985.

⁷ Butt 2007.

⁸ Scheffler 1997. This list comes from Miller 2007, pp. 100–4. The general idea of multiple connections was outlined in Miller 2001. One might wonder if some, or all, of these connections are reducible to one connection. Even if so, the distinctions are helpful in everyday reasoning and are worth making for that reason.

causation, association, and capacity. All three sources ‘add up’ to determine—or overdetermine—the fact that responsibility belongs with her. The amount of weight of each source depends on, *inter alia*, the extent to which that source is present. For example, if the sister caused the injury, but only in combination with unlikely environmental factors (slippery rocks and worn shoes, say), then causation isn’t so significant. If she’s his sister, but they have been estranged for many years, then association might have less weight. And if she’s the most capable, but not by much, then capacity is less important in the final reckoning. Even if each source is thus insignificant on its own, their combined significance may be large enough to generate a duty for her. It will sometimes—though perhaps only sometimes—be possible to roughly cardinally rank these sources in terms of the weight they lend to the duty.

Relatedly, each source’s weight in a context is affected by the *costs* that discharging the duty would impose. Different sources’ weight will be differently diminished by costs. For example, perhaps more costs can be imposed on agents in the discharge of moral responsibility-based duties than in the discharge of capacity-based duties. If so, the weight that the sister’s capacity lends her duty might be significantly lowered if helping her brother will be costly, whereas the weight of moral responsibility in generating her duty might be only marginally diminished by considerations of cost. In this way, context—which includes the extent to which various sources are instantiated and the different ways in which cost affects each source—determines the sources’ respective weight in the context.

II. Source-tracking Distributions

With this sketch of duty sources in hand, we can turn to examples where groups bear duties and must distribute amongst members the costs of discharging those duties. To settle ideas, this section focuses on a small group: five teenage friends who earn pocket money on weekends, doing odd jobs together for people in the neighbourhood. The group is, on

standard accounts, a group agent: the friends have a shared decision-making procedure, under which group decisions are made via deliberation-based consensus with majoritarian voting if no consensus is reached. This procedure is used to decide who is a member, which services the group advertises, which jobs they accept, who does which parts of which job, and so on. One day, the group is shifting a heavy bookcase. Tim is negligently chatting and not watching his footing. He stumbles in the stairwell and loses his grip. The bookcase slides down the stairs and breaks the foot of a passing child. Plausibly, the group acquires a moral duty to remedy this. The distribution question arises: how should the burdens of this duty be distributed amongst the members?

In the short-term, when the bookcase must be heaved off the child, the answer is easy. A child's foot is at stake, so each member should simply lift as quickly and forcefully as he or she is able: a capacity-relative distribution is best in the immediate aftermath. But what about the longer-term burdens? How should the child's medical bill be paid? Who should bear the brunt of the parents' reprimands? A straightforward answer is that 'the group as such' should bear these burdens. After all, one might think, the group's decision-making procedure makes it an agent, conceptually distinguishable from its members. This might work for reprimands (though even then, some individual will bear the burden of responding appropriately). But it won't work for medical bills. Perhaps the group has accumulated collectively-owned property—mowers, ladders, and so on—which could be sold to pay the bill. However, the group will then have to save for replacement mowers and ladders, or bear the costs of not owning them (such as having to reject certain jobs). Either way, costs will be imposed on members. The question remains of how *these* costs should be distributed.

We might answer that the group may distribute these costs however it likes, as determined by its decision-making procedure. Yet it is counterintuitive that the group has full reign here. Suppose the decision goes to a vote, and the group decides four-to-one that the

poorest member, Gareth, should always pay all medical bills arising from the team's duties. We would assert there's something not right about this; that the group is overlooking important moral reasons; that the group's distribution of burdens should be different. This is consistent with saying Gareth should tolerate that distribution, *if* he voluntarily signed up to the group procedure that produced it. But we should put this voluntarist line of reasoning to one side, since our ultimate concern—states—are not voluntary associations. In order to keep the discussion ultimately applicable to states as non-voluntary groups, we should focus on how distribution *should* go, supposing we were giving advice to a newly-formed group.

Perhaps, then, we should advise a capacity-relative—or maybe an equal—distribution of *both* the long-term and short-term burdens. In support of capacity-relative or equal distributions, we could point out that the acquisition of duties is a risk of acting in the world. We might recommend that members pool this risk, as insurance against individually costly duties. This would result in equal or capacity-relative distribution of the costs of 'claims' made by those in the risk pool. Such risk pooling makes sense when agents cannot control the risks. But the acquisition of many duties—in particular, wrongdoing-based duties—is the result of choice. For these, it looks less sensible to pool risks. Why would diligent Dawn, who is ever-careful about her footing in stairwells, want to pool her risk with Tim? Using a capacity-relative or equal distribution would entail distributing costs to

members who are, by any standard, morally blameless. Indeed they may be the majority. So far as they are concerned their treatment ... may seem as unfair and ethically backwards as the treatment meted out under primitive systems of collective responsibility in which whole tribes or nations are subject to reprisals: after all, in

both set-ups innocent people are called upon to pay the price for the misdeeds and mistakes of their rulers [or fellow members].⁹

In other words, equal or capacity-relative distributions are insufficiently sensitive to individuals' abilities to control their "misdeeds and mistakes" that generate duties. Additionally, to distribute the costs of wrongdoing-based duties in a desert-insensitive way allows individuals to "distance themselves morally and psychologically from the wrong in which they are implicated".¹⁰ This creates a moral hazard: why bother watching your footing in the stairwell, if you know your mates will help pay for any nasty consequences?

One might defend equal or capacity-relative distributions differently, pointing out that the members are "beneficiaries of a common practice in which participants are treated fairly ... and so they should also be prepared to carry their share of the costs..."¹¹ But we are trying to work out *what it is* for members to be treated fairly in the common practice—or at least, in that aspect of the practice that involves distributing the group's duties. The practices of the group are not an all-or-nothing affair, such that members must accept the entire package if that package provides them with net benefit. Instead, we should consider whether the benefits of working together (for example, receiving pocket money) can be pragmatically disentangled from the costs (for example, paying for your friend's negligent stumbling). In the case of the teenagers, these can be separated: if they distribute costs non-equally or not relative to capacity, they will not bring crashing down the entire package of value contained within their shared practice. If the benefits and costs can be disentangled, then there is no

⁹ Crawford and Watkins 2010, p. 290.

¹⁰ Ibid., p. 291.

¹¹ Miller 2004, p. 253. In the case of states, Pasternak (2011, 199-201) similarly argues that "living in a self-ruling political community" provides net benefits such that citizens of democracies should accept equal distribution.

justification for appealing to net benefits in order to block members from asking for the benefits but not the costs.

A third justification for equal or capacity-relative distributions appeals to group solidarity. According to Avia Pasternak, there are “two components of solidary action: first, people who act in solidarity are acting jointly to achieve a worthy social goal [for example, doing jobs for people in the neighbourhood] and, second, they view themselves as equal participants in their joint project.”¹² Pasternak argues that “if ... solidary action is an intrinsically valuable social good that characterizes certain relationships, then in those relationships the equal sharing of the consequential responsibility for the group’s joint actions is in fact an associative obligation of its participants. Equal sharing is an AO [associative obligation] because it is constitutive of the social good of solidary action.”¹³ Pasternak illustrates this with a group pushing a car out of the mud. When they succeed, they each (rightly, she implies) see themselves as deserving an equal amount of the driver’s thanks, because they were acting in solidarity.¹⁴

Can solidarity do this work? I’m not so sure. Suppose the least-helpful member of the car-pushing group tells the owner that the others did more work than her, and therefore deserve more thanks. In other words, suppose she resists an equal distribution. Would we think that she had failed to discharge an associative obligation to share responsibility equally? We would not. We might think that she had no obligation not to take an equal share of praise. But Pasternak’s claim is not that solidary actors are *permitted* to share responsibility equally. It’s that they *ought* to do so. Conversely, suppose the car-pushers shove the car too far, and run over the owner’s cat. If one member was standing in front of the car to check the way was clear, we might think it positively immoral for him to absorb an equal share of the owner’s

¹² Pasternak 2011, p. 200.

¹³ Ibid., p. 197.

¹⁴ Ibid., p. 196-7.

rage, rather than putting himself forward for the most censure. The point is this: an imperative to share responsibility equally does not seem to arise automatically out of solidary action, unless by definitional fiat.

The individualist intuitions that drive these objections suggest the following. *If* the teenagers' duty to remedy the broken foot is purely wrongdoing-based, and *if* Tim is the only one who wrongly stumbled, then Tim should pay the entire bill and absorb all of the parents' reprimands.¹⁵ Suppose Tim stumbled, while Dawn took great care to avoid stumbling. Perhaps Dawn did warm-up stretches, wore oppressively sturdy shoes, craned her neck uncomfortably to watch her footing, and so on. Dawn took on costs specifically to avoid stumbling, and succeeded in doing so. It's not clear how the group might justify a wrongdoing-insensitive distribution to her. If Dawn was also negligent, then she should bear some costs. If Dawn was a model removalist—and if the duty is entirely wrongdoing-based—then she should not pay.

To extend these individualist intuitions to other sources, suppose the group happens to pass a child trapped under a fallen bookcase. The group thereby acquires a capacity-based duty to lift it. Since the group's duty is capacity-based, a capacity-relative distribution seems plausible for both immediate costs and any final costs that might arise. Or imagine the group is moving furniture for Tim and Gareth's elderly parents, whom the other members also know well. The group considers itself to have an association-based duty to do the job for free. Although it is the group's duty, it seems fairest for Tim and Gareth to take on the hardest work. After all, they have the strongest association with the parents.

Of course, not everyone will share these individualist intuitions. But the aim of the rest of this paper is to see how far they can get us when considering states' duties. The

¹⁵ The fact that capacity loomed large in the immediate aftermath might suggest that the group's duty isn't entirely wrongdoing-based. Later in this section, I discuss a multi-source interpretation of this example. §III.C explains what should happen if Tim cannot pay.

general proposal is that the costs of discharging a group agent's duty should be distributed amongst members in proportion to the members' respective instantiations of the source(s) of the group's duty. As per §I's framework, those sources will be different in different cases and often a combination of sources will contribute to the duty's strength. To make the proposal more precise, distribution should proceed in two steps. First, the costs are divided among duty sources, in accordance with those sources' weight. Second, the costs earmarked to each source are divided amongst members in proportion to members' individual instantiations of that source. Cases with one source are easy. If a group's obligation is based purely on causing some harm, then costs should be distributed in proportion with members' causal contributions. Cases with more than one source are trickier: suppose the teenagers' duty is generated by both wrongdoing and capacity, with wrongdoing being twice as weighty as capacity. In this case, first, (roughly) 67% of costs should be earmarked for wrongdoing and (roughly) 33% of costs should be earmarked for capacity. Second, the wrongdoing-earmarked costs should be divided among members in accordance with their varied levels of wrongdoing, and likewise for capacity. This is the 'source-tracking' answer to the distribution question.

III. Application to States

Source-tracking is relatively straightforward for the teenagers. States are more complicated. These are group agents that last for centuries and are usually constituted by millions of people, who vary wildly in their capacities, associations, receipt of benefits, and involvement in group wrongs. Members are often turned into adversaries by the very decision-making procedure that unites them, some members have much more influence over the group's decisions than others, the reasons behind group decisions change with the winds of political favour, and many decision-making methods are casuistic or not explicit. These features make source-tracking difficult.

Yet, just as for the teenagers, I believe source-tracking captures commonsense intuitions about the distribution of states' duties. Suppose the US had an obligation to assist Iraqis whose lives were ruined in the aftermath of the 2003 invasion. Members of the US who protested against the invasion could rightly ask why they should have to contribute to the costs of this assistance. If the state's obligation was grounded wholly in wrongdoing, then the source-tracking proposal would uphold the protestors' complaints and redirect the costs to those members who were wrongly involved in the action. But if the US's obligation was instead grounded in (say) capacity, then the members' question would have a straightforward answer: you should contribute to the assistance because you have the capacity to do so. Either way, the members receive an answer that appeals directly to the reality of their lives as individuals. This is very appealing. Yet even this quick example evokes numerous kinks that must be ironed out, if source-tracking is to apply to states. The rest of the paper examines these.

A. Practicalities

One initial complication concerns the obvious infeasibility of distributing the costs of states' duties in perfect proportion to each member's instantiation of the relevant duty sources.¹⁶ This problem can never fully be addressed—the law is a blunt instrument. But there are feasible ways of approximating this ideal. One is targeted taxes. This arguably already occurs for states' causation-based obligations to alleviate global warming: some states pass this cost onto consumers of petrol, through higher purchasing taxes on that good. The source-tracking view would justify this, stating that individual causers of global warming should bear the costs of the state's fulfilment of the causation-based 'part' of its duty to alleviate global warming. More generally for causation-based duties, if an activity causes the harm, then that

¹⁶ Pasternak 2011, pp. 192–3.

activity can often be taxed. One might think this unpalatably allows past causers to avoid costs, by ceasing to consume the relevant good. But such cessation is itself a cost, justifiably imposed upon past harm causers in virtue of their causation.

Similarly, the capacity-designated percentage of a duty's costs might be paid out of general government funds. The more costs there are that are both earmarked for capacity and paid out of general funds, the more reason there is for that state to have both higher and more progressive taxation on wealth. Such capacity-based duties are sometimes tricky: if the UK has a purely capacity-based duty to provide medical care to climate refugees, should this cost be lumped entirely on doctors? If source-tracking implied so, then it would be unpalatable. Intuitively, taxpayers more generally should bear the costs of paying these doctors well, and so bear costs for the exercise of doctors' capacities. But this just demonstrates the importance of being clear about which capacity generates the duty: it is the UK's capacity to *adequately train and remunerate doctors* that generates its duty, rather than the capacity of particular doctors to provide medical care.

Finally, suppose a state's obligation to reduce carbon emissions is based equally on causation, capacity, and benefitting. We can then imagine a third of the relevant taxes targeting users of petrol, a third of them targeting the wealthy, and a third of them targeting oil or transport companies (or whoever benefits from emissions). Inevitably, such policies will generate some false negatives and false positives. But this need not stop us getting as close to a source-tracking distribution as is feasible.

B. Individual-level Instantiation

A different issue is source-tracking's presumption that a group's duties will usually be instantiated by individual members. In large-scale and complex groups like states, individual-level instantiation might seem like the exception, not the rule. If that's correct, then source-

tracking in states looks hopeless: most sources of states' duty sources just don't apply to individuals. A couple of examples motivate this thought.¹⁷

Wrongdoing-based duties are a first example. Consider the 1979 Air New Zealand crash on Mount Erebus. There, 257 people died as a result of inadequate company organisation. This led to an incorrect flight path being taken, through no fault of any individual.¹⁸ Or consider 1987 the sinking of the *Herald of Free Enterprise* off the coast of Belgium, where 193 people died as a result of a "body corporate ... infected with the disease of sloppiness."¹⁹ This sloppiness ultimately led to a crewmember being asleep in his cabin when he should have been closing the bow-door. In such cases, the group's duty seems to be sourced in wrongdoing. But all individuals did what they should have done according to their prescribed roles in the organisation, and they couldn't have known that they should have done differently. They did no wrong. So, it seems, the costs cannot be source-tracked. While these examples concern companies, it is easy to see how this problem extends to states.

Notice, though, that for groups to cause anything, there must be individual-level goings-on. These individual-level goings-on may not be 'causes' proper, particularly if the group-level phenomena (the airline's disorganisation, ferry company's sloppiness) are realisable in multiple different ways at the individual level.²⁰ But the individual-level occurrences should not be ignored. They do help constitute the group-level phenomena in its actual realisation. So we should ask: what individual-level occurrences constituted the airline's development of inadequate organisation, or the ferry company's slide into sloppiness? These things did not happen *ex nihilo*. Rather, individuals would have voted for bad reasons, turned blind eyes, heeded bad advice, believed implausible predictions and excuses, failed to hold each other to account, failed to challenge norms, failed to pick up slacks, failed to install

¹⁷ I thank Anna Stilz and an anonymous reviewer for pressing such examples.

¹⁸ Stilz 2011, p. 193.

¹⁹ Pettit 2007, p. 171.

²⁰ List and Menzies 2009.

sufficient checks and balances, and so on.²¹ When states do wrong, many ordinary members are guilty of just these kinds of failings. These individual ‘enactors’ should bear costs, on the source-tracking proposal’s application to wrongdoing-based duties.²² And it does not matter that states’ wrongdoings are often overdetermined by such individual-level enactors: causes that overdetermine outcomes are plausibly still linked to the outcomes, in the sense of ‘linked’ relevant to wrongdoing-based duties. For example, an overdetermining cause is a ‘NESS’—“a necessary element of a set of antecedent actual conditions that was sufficient for the occurrence of the consequence.”²³ Of course, it might be difficult to identify these individual enactors, especially if their connections to the wrong did not violate their prescribed role in the group—but that is an issue of practicalities, not an issue of non-existent individual-level instantiation.

Association is a second duty source where individual-level instantiation may seem doubtful in states. Perhaps Germany can bear associative relations—and so associative duties—to France, without any non-governing Germans having associative relations with—or even knowing—any non-governing French. This inter-state associative relation is constituted by explicit agreements, intertwined histories, geographical proximity, similar values and cultures, and so on. These don’t seem to have individual-level correlates. Intuitively, though, the costs of Germany’s associative obligations to France can legitimately be passed on to Germany’s non-governing members. Source-tracking, it seems, cannot give us this intuitive result. However, several of features of states’ associative relations do have analogues at the individual level. These can occur even without the individuals knowing one another. French

²¹ Hindriks 2009; Mäkelä 2007.

²² The language of ‘enactors’ is used by Pettit (2007, p. 189ff.) and List and Pettit (2011, pp. 161-3), when they endorse the co-existence of individual-level and group-level contributors to group wrongs.

²³ Wright 1985, p. 1802. Wright (1985) and Stapleton (2008) each argue that the law should use the NESS account of causation, especially in tort cases. For different accounts of what is wrong with being an overdeterminer of harm, see Kagan (2011) and Barry and Øverland (forthcoming).

and Germans have intertwined genealogies, live relatively nearby to each other, and have similar values and cultures. They deeply affect one another's lives, even if neither party knows it. These facts can be used to justify and implement a source-tracking distribution of their states' mutual associative duties.

So we should not be so quick to doubt that ordinary members will instantiate the sources of states' duties. We can instead approach each case with the assumption that individual-level instantiation is the rule, not the exception. To take the exceptions as a general reason to avoid source-tracking, as and when it is possible, seems unnecessarily cautious. That said, source-tracking should say something about the exceptions. To do this, we must turn to the third and biggest complication for source-tracking in states: shortfalls.

C. Shortfalls

i. Topping Up Shortfalls: Two Conditions

When examining shortfalls, duties sourced in historical wrongdoing serve as a stark and compelling case study. My discussion will therefore focus on these, though any resolution of them must generalise to other sources. Suppose, then, that a state committed a grave wrong 200 years ago—slavery or colonialism, for example. Assume we can assess the counterfactuals and establish that present-day individuals face incredible hardship as a result of the wrongdoing. Suppose the state thereby incurs a wrongdoing-based duty to compensate these present-day people, although no present-day individuals contributed to the wrong. The source-tracking model is silent on how the state should pass the cost of its duty onto its members. It seems the state bears a duty, but no individual should bear costs through its fulfilment. The duty is thus 'shortfallen'.

Shortfalls have led recent authors to reject source-tracking for wrongdoing-based duties. As Anna Stilz argues, “[s]queamishness about implicating individuals ... seems

rarefied when we turn our attention to the uncompensated wrongs perpetrated by states in their citizens' names."²⁴ Likewise, Pasternak entertains the idea of source-tracking the costs of states' obligations to repair harms. But, she says, this has "one obvious drawback: [... it] is likely to create 'liability shortfalls', whenever it would be hard to identify those individuals within the state who are more responsible for its misconduct, or to extract the necessary resources from them."²⁵ This imperative to repair historical wrongdoing makes wrongdoing-based duties the litmus test for any response to shortfalls.

I suggest we start addressing the shortfall problem by asking why shortfalls bother us. For example, why do we care that, on the source-tracking model, present-day white Australians are not liable to costs as a result of the wrongdoing-based duties the Australian state has to Aborigines? The past wrongdoing of Australia obviously has something to do with it. But perhaps it is not the end of the story. Consider Leif Wenar's examples:

We do not hear calls for reparations from the Germans for the Allied bombings of Dresden, or from the United States for the Japanese surprise attack on Pearl Harbor, or for that matter from the United States for the British arson of the White House in 1814. These are events which many at least believe were historical injustices, but we do not hear demands for reparations from these believers because there is little sense that those who could potentially claim reparations under backward-looking principles suffer from an unjust distribution of rights or resources *at present*.²⁶

²⁴ Stilz 2011, p. 196-7.

²⁵ Pasternak 2013, p. 365.

²⁶ Wenar 2006, p. 402, emphasis in original.

On this basis, Wenar surmises that “valid claims for reparation rest at the deepest level ... on reasons we have to improve our current relations so that we can get along better in the future.”²⁷

My suggestion, in the spirit of Wenar’s, is that states’ obligations are often not entirely what they seem. This is not as strong as Wenar’s suggestion that non-wrongdoing sources are the *only* sources of states’ reparative duties. I accept that wrongdoing might be a—or the only—source of some duties of states. In cases without shortfalls, this poses no problem. But in cases where duties are shortfallen, any legitimate demand that members bear costs to make up the shortfall must derive from duty sources that are not shortfallen. In other words, in order to justify imposing shortfallen costs on members, we must assert that non-shortfallen duty sources are adding some weight to the state’s duty. We can then justifiably require members to make up the shortfall, on the basis of their individual instantiations of these non-shortfallen sources.

There are two main conditions here. First, the non-shortfallen sources must have genuine weight in generating the state’s duty. Such sources often do this when and because there is deep damage to the current (or projected future) relationship between the state and the would-be reparation receivers. An apt description of such ‘damaged relationships’ is given by Margaret Urban Walker. She describes damaged relationships as those in which participants lack *confidence* that they share the same standards for reciprocal treatment, or lack *trust* in another to abide by those standards, or lack reasons to be *hopeful* that inadequate

²⁷ Ibid., p. 396. Sources other than wrongdoing are similarly used to ground reparations for historical injustice by: those in Kumar and Tan (eds) 2006; Miller 2007; Thompson 2002; Vernon 2003. The justifications generally do not appeal to responsibility shortfalls (although Kukathas (2006, pp. 340–1) alludes to these), but rather problems with assessing counterfactuals, with identifying victims’ descendants, or with retributive justice generally. Of course, many authors working on historical injustice do not side-line wrongdoing (e.g., Sher 2005; Simmons 1995).

relations will be remedied if they come about.²⁸ Of course, in some contexts the fact that a relationship is damaged doesn't generate any duties at all. Damaged relationships are most troubling when parties have frequent or avoidable interaction, or where confidence, trust, and hope between them are particularly low. In such cases, duty sources such as association, capacity, and benefit will often—perhaps in combination with one another—generate duties for the parties to repair the relationship, independently of wrongdoing.

Like Wenar, I believe this holds in cases of states' historical wrongdoing where reparations are intuitively appropriate. In such cases, the state has a damaged relationship with the intuitive reparation recipients. The relationship of the US to black Americans, of Australia to Australian Aborigines, of Germany to allied states post-WWI, of Iraq to Kuwaitis after the First Gulf War, or even of the US to Iraqis now, are all damaged relationships. These are also instances in which sources other than wrongdoing are instantiated by the state: they are cases where there is a *close association* between the state and its partner in the damaged relationship, or where the state *benefits* from being in the damaged relationship (say, because it has power over the other party), or simply where the state has the *capacity* to repair the damaged relationship. A state's duty to repair its damaged relationship can come from a range of, or combination of, such sources—even if we ignore the duty-generating power of wrongdoing. Agreeing with this is not hard: simply imagine a relation between a state and a population, in which general confidence, trust, and hope are very low, due to exogenous factors. Would we say there are reparative (literally, repairing) obligations? I think would. If so, then something other than wrongdoing must be generating them.

So much for the first condition on topping up shortfallen wrongdoing-based duties: the state itself must have a duty grounded in non-wrongdoing sources. I've suggested this

²⁸ 2006a, p. 384; for in-depth discussion, 2006b.

will typically be a duty to repair the damaged relationship between the state and the would-be reparation-receivers. The duty to repair this damaged relationship might be grounded in association, benefit, capacity, or some combination of these. The second condition has two parts. First, (i) enough of the state's members have a damaged relationship with the would-be reparation-receivers, where the individual-level damaged relationship is causally or constitutively connected to the state-level damaged relationship. Second, (ii) enough of these members instantiate the sources that ground the state's reparative duty, where the individuals' instantiation of these sources is causally or constitutively connected to the state's instantiation of these sources. The causal or constitutive connection between individual-level and state-level facts mirrors the teenagers' individual properties, which caused or constituted their group's instantiation of duty sources.

Again, in cases where reparations are intuitively appropriate, this second condition seems to hold. Take the US state's evident obligation to repair its damaged relationship with black Americans. As Walker explains, relationships between black Americans and white Americans are damaged in ways that are similar to the relationship between black Americans and the US state.²⁹ Consider insidious practices such as racial profiling. This is practiced by state officials (prominently police) and by ordinary citizens on the street. One might object that many white Americans know very few—perhaps zero—black Americans. How can their relationship be damaged, if there is no relationship to speak of? In reply: a relationship is damaged if it lacks confidence, trust, and hope. This doesn't require interactions. It can be assessed counterfactually, by asking whether a lack of confidence, trust, and hope *would* be exhibited *if* the parties were to interact. Research on individuals' propensity for racial profiling, stereotyping, and implicit bias suggests that many members of the US do have

²⁹ 2006a, pp. 387–92.

racially-aligned damaged relationships in this sense, even if they are strangers.³⁰ Moreover, there is plausibly a causal or constitutive link between the damaged relationship at the individual level and at the state level: individuals' implicit racism both reinforces and is reinforced by that of state officials. This is not least because state officials just are individuals, who inevitably influence and are influenced by the racist implicit attitudes of those around them.³¹ Condition (i) seems to hold.

Additionally, enough members of the US state instantiate the same duty *sources* as the US state, with regard to their duty to repair their damaged relationship. White Americans are *associated* with black Americans through geographical and social proximity, white Americans *benefit* from the damaged relationship (through 'white privilege'), and white Americans generally have a greater *capacity* (money; political and social power) to repair the relationship. And these individual-level sources plausibly perpetuate, and/or are perpetuated by, the state's instantiation of the same sources. Indeed, it's not clear what it would mean for a state to be geographically proximate, or receive a benefit, or have some capacity, without some members also being geographically proximate, or receiving redounded benefits, or having the capacity to do their part in the exercise of the state's capacity. The individuals' duty sources partly constitute the state's duty sources, even if though latter are not entirely reducible to the former. So condition (ii) also holds. It is therefore legitimate for the state to impose costs on white Americans in the discharge of its duty to repair its relationship with black Americans. This can be done without reference to the shortfallen source, wrongdoing.

This works for other shortfallen sources. Consider §III.B's example: Germany's associative obligation to help defend France from invasion. Imagine this duty arose purely

³⁰ For an overview, see Brownstein 2015.

³¹ The causal or constitutive relationship is difficult to prove. But sociological research has found, for example, that an increase in media reporting on racial profiling led to a decrease in searches of black Americans by highway patrols. The increase in media reporting was itself a result of increased awareness of implicit racism amongst individuals. (Warren and Tomaskovic-Devey 2009)

from an associative relation that lacks individual-level correlates: an agreement or promise between the states, made before any present-day Germans were born. To justify imposing the costs of this duty upon Germans, I would suggest appealing to duty sources other than agreement or promising. For example, plausibly, any state's duty to honour age-old treaties is partly sourced in the benefits for all of *pacta sunt servanda*. Insofar as members have moral reason to produce these benefits when they can, and insofar as their ability to produce these benefits partly constitutes Germany's ability to do so, the 'part' of Germany's duty that is sourced in this capacity can be distributed in a source-tracking way.³² Alternatively or additionally, Germany's capacity to avoid damaging future France-Germany relations might play duty-generating roles in this case, with similar individual-level correlates.

ii. Topping Up Shortfalls: The Proposal in Detail

These two conditions lay the ground for my more detailed proposal. Ideally, we should implement the source-tracking proposal, outlined in §II. However, if some source(s) of a state's duty are instantiated by so few members that we cannot impose upon them the entirety of the costs earmarked for the shortfallen source, then we should impose what we can upon these few members on the basis of that shortfallen source. We should then consider whether other, non-shortfallen sources are weighty enough to generate a duty to make up the shortfall. If they are, then we should normalise the weightings of these sources—that is, make them sum to 100%, retaining their relative weightings—and distribute the rest of costs according to these new normalised weightings.

As a simplified example, suppose Australia's duty to pay \$1billion in reparations to Aborigines is based 50% in wrongdoing, 30% in capacity to repair relations, and 20% in benefitting from harm. Suppose (falsely) that absolutely no member instantiates wrongdoing.

³² I thank James Pattison for suggesting this.

We then ignore wrongdoing when distributing costs, and normalise the remaining sources: the duty is now 60% capacity-based and 40% benefit-based. We then implement §II's source-tracking proposal for this new ratio of sources, *if* association and benefit have enough aggregate absolute weight to generate a duty to pay the \$1billion independently of wrongdoing.

This 'if' clause prevents over-inflating the newly normalised duty sources, and will be met only when the original duty is overdetermined. We should normalise and re-earmark the full costs *only if* the non-shortfallen duty sources are strong enough, in aggregate but independently of the shortfallen source, to generate a duty to pay the shortfallen cost. And even once normalised, the total cost that may be earmarked to any one source is subject to a ceiling, which is determined by the absolute weight of that source. The 'cost ceiling' of any given source might not be hit if there are other sources that both generate the state's duty and can be used to distribute costs to members. But once some sources are ignored (because they are shortfallen), then the costs that are earmarked to the remaining sources will come closer to—though never exceed—their respective ceilings. The point of this is that normalising does not enable us to attribute more liability to members than individual-level considerations would require them to bear.

However, this means that any duty generated by the non-shortfallen sources might demand less than would be demanded if those duty sources were working in tandem with the shortfallen source. Perhaps capacity and benefit together can demand only \$0.5billion in reparations, rather than the \$1billion they would demand alongside wrongdoing. If so—and if there really are *no* individuals who instantiate wrongdoing—then only this less expensive duty should be discharged. Crucially, though, if we balk as such a result in a particular case, then we should consider whether capacity, or benefit, or some other source, is in fact doing

more work than we initially thought, and can indeed justify the \$1billion imposition. The role of the ‘if’ clause is to force us to consider this.

A slightly different treatment of this example is required if there are some living members who do instantiate wrongdoing, but it would be objectionably burdensome to lump these few members with 50% of the total costs. Here, we should impose whatever cost is appropriate on these remaining wrongdoers, in accordance with their respective wrongdoing. The source-tracking proposal is neutral about what costs are appropriate in such cases: perhaps, as at law, wrongdoers should sometimes be severally liable *in full* for repairing the wrongs they commit with others.³³ If so, the few remaining wrongdoers might be liable for the full 50%. But suppose these wrongdoers are such that we should give them together only 10% of costs. Then, this 10% should be imposed on them. The remaining \$0.9billion should then be re-normalised for capacity and benefitting, and so earmarked 67% to capacity and 33% to benefitting. If there is a strong victim-focused imperative to see Aborigines receive the \$1billion, then it will be plausible to let these sources carry this weight.

There are two main reasons for responding to shortfalls in this way. One is that it allows the non-shortfallen duty sources to exercise the strength they would anyway have exercised in the absence of a shortfall. It is thus properly responsive to those sources. The “if” clause ensures this is true. I can see no reason to deny these sources the duty-generating strength they have independently of the shortfallen source, in a situation where that strength allows us to make up shortfalls that are, *ex hypothesi*, morally troubling. A second reason is that it is justifiable to the individual cost-bearers. It appeals to their attributes as individuals, and to duty sources that apply to them as individuals, without needing to rely on moralised conceptions of membership, or concoct or presume any kind of solidarity between them and their state or fellow members. In this way, we can respond to victims’ claims—the claims

³³ On the law here, see Cane 2002, p. 178.

that ignite our concerns about shortfalls—while also being able to justify the imposition of costs on individual members of the duty-bearing state, by evoking sources they actually instantiate.

Pragmatically, of course, in some contexts only a non-cardinal ordering of sources' importance may be possible. And, more generally, it will rarely be plausible to produce precise numbers, by asserting that association is, say, twice as important as benefiting. These facts make source-tracking difficult to implement in a highly specific way. Ultimately, contextual debates will have to be had between those with differing views about sources' weightings. Compromise between these differing views will inevitably have a role in determining final distributions. In some cases, though—such as that of black Americans or Australian Aborigines—the same cost-bearers and rights-bearers will be implicated, regardless of which precise combination of sources we use to top up the shortfall.

iii. Two Objections

One objection to this 'topping up' method is that the duty the state ends up discharging is not its original duty. There remains a shortfall in some sense, since the group's wrongdoing-based duty never gets discharged *qua wrongdoing-based duty*—and likewise for other shortfallen duties. The state discharges a different duty: a capacity-based duty, or a benefit-based duty, or some combination of these—but it never discharges the duty that derived from the particular shortfallen source. This is true. But in cases where shortfalls bother us (and not all will), the crucial concern is that victims should receive recompense to the full extent feasible. This is what motivates the shortfall problem, as seen in the quote above from Stilz. And this concern is resolved, regardless of what duty source we use to resolve it.

One might respond that, for victims of historical injustice, “it will not do simply to lump them in with the poor more generally, for their claim is of a very specific

nature.”³⁴ Unfortunately, though, if that ‘specific’ (i.e., source-indexed) claim cannot be resolved without imposing unfair costs on others, then that specific claim will have to go unresolved. Crucially, though, there is still a specific claim that is acknowledged. This is a claim based on a very particular capacity, or particular benefits, or particular association, or so on. To explain why various sources—like capacity, association, or benefit—are duty generators for particular collectives and individuals particular contexts, we will have to make reference to historical, relational, and material details. Thus the new claim will be highly specific. This prevents victims from being “lumped in with the poor more generally”—even if the newly-acknowledged specific claim is different from the specific claim of having been wronged.

A second objection asks: what if there is a shortfall and the ultimately-evoked duty sources are instantiated by numerous states and their members? This is problematic in those cases where it seems clear that one state alone bears the duty in the first instance, with others having duties only if the first state reneges. To illustrate, suppose Myanmar brutally oppresses its minority Rohingya population. Generations pass and the individual oppressors die. Yet the relationship between Myanmar and the Rohingya is still badly damaged. My approach to shortfalls suggests that the present-day state has a duty to repair the damaged relationship, and (at least some of) its present-day members have duties to bear costs. The state-level and individual-level duties are based on capacity, association, or benefit—not wrongdoing. This is fine as far as it goes. Yet meanwhile, suppose Australia (and Australians) have the greatest capacity to help the Rohingya, Bangladesh (and Bangladeshis) are most strongly associated with the Rohingya, and China (and Chinese) have benefited the most from the initial injustice.³⁵ Under source-tracking, it seems we cannot assign the duty to

³⁴ Kukathas 2006, p. 332.

³⁵ I thank an anonymous reviewer for this example, which I have adapted slightly.

Myanmar (and Myanmars) alone. Rather, all these states (and their members) have duties. Intuitively, this is the wrong result.

My response comes in two steps. First, recall that multiple duty sources *add up* to determine the overall strength of an agent's moral reasons to perform a particular action. It might be true, of each duty source, that some other state (and its members) instantiates that source more strongly than Myanmar (and Myanmars) do. But the fact that Myanmar instantiates all three sources is likely to make it the group agent with the strongest *overall* moral reason to repair their relationship with the Rohingya. If Myanmar's members also instantiate all three sources, then the costs can likely be distributed in a source-tracking way. This will place responsibility primarily at Myanmar's door.

This reply will not always work. Suppose Australia is by far the most capable of helping the Rohingya. We can then imagine that Australia has a stronger duty, sourced in capacity alone, than the Myanmar's duty, which is sourced in capacity, association, and benefitting combined. It then seems Australia's duty is stronger than Myanmar's: a bad result. This leads to the second step in the reply, which emphasises the intimate connection between being *a participant in a damaged relationship* and *having a capacity to repair that relationship*. 'Capacity to repair a damaged relationship' does not mean 'gross domestic product' or 'general international political influence.' In politics, as in personal life, outsiders' capacities for relationship repair tend to be conditional on the willingness of relationship participants to exercise their capacities for relationship repair. This makes it very odd to assert that Australia's capacity *to repair the Myanmar-Rohingya relationship* could possibly be stronger than Myanmar's capacity to do so. This is consistent with Australia having the strongest capacity to *more generally assist the Rohingya*. If Australia has that capacity, then Australia has the strongest duty to generally assist the Rohingya, while Myanmar has the

strongest duty to repair the Myanmar-Rohingya relationship. And that looks like the right result.

These objections highlight two details about my proposal. First, we must always be clear about the content of the duty—about which outcome or action the duty ranges over—since this will affect which duty sources apply, to which agents, with what strength. For example, is the action ‘help the Rohingya’, or ‘repair the Myanmar-Rohingya relationship’? If the former, Australia’s capacity looms large; if the second, Myanmar’s capacity, association, and benefits gain more traction as duty sources. Each agent has a duty over their respective outcome. Second, we must be aware that contextual details can furnish particular agents with highly specific capacities, associations, and benefits, which can act as sources of weighty duties to exercise those specific capacities, act in accordance with those specific associations, or disgorge those specific benefits. Thus while Bangladesh might be closely associated with the Rohingya, and might have strong duties *to aid the Rohingya* on that basis, the Bangladesh-Rohingya association doesn’t oblige Bangladesh to *fix the Myanmar-Rohingya relationship*. The Myanmar-Rohingya association, however, does oblige Myanmar to aim at that particular outcome. Likewise, while China might have benefited the most from the injustice, and so have duties to disgorge or compensate for those particular benefits, Myanmar will have received different benefits, that give them a stronger duty to compensate for those particular benefits.

D. Reductionism?

This response to shortfalls evokes the fourth and final aspect of source-tracking’s application to states. It concerns reductionism. ‘Reductionism’ is a slippery term, which I use to refer to the view that “[e]very particular object [e.g., state’s duty] in our social ontology is identical to (‘the same as’) some object [e.g., individual duty or set of individual duties] in the individual-

level ontology. (There is at most a difference in description.)”³⁶ Reductionism is controversial—particularly when the relevant social object is a state’s duty.³⁷ If source-tracking entailed reductionism about states’ duties, then all non-reductionists would have reason to reject source-tracking. And source-tracking might seem to entail reductionism. After all, once we’ve divvied up the costs amongst members, it seems we can ignore the state’s duty altogether.

Source-tracking does not entail reductionism. Source-tracking answers the question: how should the costs of groups’ duties be distributed amongst members, *when such costs must be distributed?* This question allows that the italicised clause can be false. For example, a state’s duty might require that various government departments be re-structured. We can imagine this not burdening any member: the members managing the restructure are fairly remunerated and everyone is happy with their new position. Such a state-level duty implies no costs for members. So that duty does not fall under this paper’s guiding question. This does not mean the duty doesn’t exist. And it doesn’t imply anything about whether the state’s restructuring duty is identical to a set of individual-level duties to enact the re-structure.³⁸ This point establishes that source-tracking is, at least, neutral on reductionism.

Two other points demonstrate that, as I have formulated it, source-tracking positively endorses non-reductionism. First, the state’s duty has explanatory power. It explains why there are costs that need distributing in the first place: they are the costs necessary for discharging the state’s duty. The state’s duty also explains why we must ask after specific duty sources of individuals: we should ask after those sources that generate the state’s duty. And the state’s duty explains why we should *not* ask after other specific duty sources of individuals: we should ignore those individual-level duty sources that the state does not

³⁶ List and Spiekermann 2013, p. 633. This is one of four kinds of reductionism they outline, but is the most pertinent here.

³⁷ See references in fn. 2.

³⁸ I examine this question in Collins unpublished.

instantiate. The latter sources may generate individual-level duties, but the discharge of those individual-level duties would not contribute to the discharge of any duty *of the state*. The duty of the state is central to the source-tracking story.

A second non-reductionist feature of source-tracking is that it endorses the ‘multiple realisability’ of groups’ duties. Multiple realisability is incompatible with reductionism.³⁹ Take the UK’s duty to reduce carbon emissions. This duty persists despite members being born and dying, and changing jobs and social roles. Under source-tracking, such changes mean that changes should occur in the UK’s distribution of costs. This means the state’s duty entails different distributions—is realised in different ways at the individual level—at different times. Yet each distribution is a realisation of one overarching duty, which belongs to the state rather than its members. Several facts unite these different distributions into one overarching state-level duty: the distributions are enforced by the state; the distributions include only those who are under the state’s authority; and the distributions allow the state to ‘tick off’ a particular duty on its moral ‘to-do list’. These facts pick out the different cross-temporal distributions as each being a distribution of—a realisation of—one and the same duty of the state. The state’s duty is thus included in the source-tracking ontology, at least as I have formulated it. A more individualist version may be possible. But the non-reductionist version maintains the intuitive connections between (and, therefore, existence of) the duties of our states and our duties as individuals.

IV. Conclusion

When a state has a duty that requires the distribution of costs amongst members, how should the costs be distributed? I have answered: in proportion to the members’ respective

³⁹ On multiple realisability as a reason for non-reductionism about groups’ duties, see List and Pettit 2011, pp. 161-3. On multiple realisability generally, see List and Menzies 2009.

instantiations of the sources of the state's duty. Sometimes it is not possible to give the costs of a state's duty a distribution that tracks all sources of the state's duty—as in cases of historical wrongdoing. Here, we should acknowledge that such shortfalls matter in cases that cry out for someone to act, even if no individual instantiates the shortfallen source. This suggests that something other than the shortfallen source is doing work; that some other duty source is generating the imperative to make up the shortfall. The proposal is that we should understand states' seemingly shortfallen duties as grounded in duty sources that are instantiated by a sufficient number of members, and so do not result in shortfalls. If no such sources are weighty enough to generate a duty in a given case, then should we consider cost-sharing among many states. However, my hunch—not fully defended here—is that sources other than wrongdoing can carry much more weight than we often think, both in generating states' duties and in enabling source-tracking distributions of those duties' costs. If so, then where ever there are victims, there will almost always be states and individuals who bear duties.

I lack space to fully discuss how this might be implemented in international law. The general suggestion is this: the less correlation there is between the sources of a state's duty and the sources instantiated by members, the more hesitant international legal institutions should be to allow the duty's costs to be distributed to members. Such institutions might instead assign the duty to the state on the basis of another duty source, which does apply to members. This would require an expansion of international law's duty bases. A compatible option would be for international legal institutions to dictate or advise how the costs of discharging a state's duty should be distributed. Targeted taxes are the obvious suggestions. Such tactics allow us to do justice to both parties: those to whom the duty is owed, and those on whom the costs ultimately fall.

References

- Barry, Christian and Gerhard Øverland. Forthcoming. "Individual Responsibility for Carbon Emissions: Is There Anything Wrong With Overdetermining Harm?" In Jeremy Moss (ed.), *Climate Change and Justice*. Cambridge University Press.
- Brownstein, Michael. 2015. "Implicit Bias." In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <<http://plato.stanford.edu/archives/spr2015/entries/implicit-bias/>>.
- Butt, Daniel. 2007. "On Benefiting from Injustice." *Canadian Journal of Philosophy* 37(1), 129–152.
- Cane, Peter. 2002. *Responsibility in Law and Morality*. Oxford: Hart Publishing.
- Collins, Stephanie. Unpublished. "Organisations' Duties and Members' Duties."
- Crawford, James and Jeremy Watkins. 2010. "International Responsibility." Pp. 283–98 in S. Besson and J. Tasioulas (eds), *The Philosophy of International Law*. Oxford: Oxford University Press.
- Crawford, James. 2007. *The International Law Commission's Articles on State Responsibility: Introduction, Text and Commentaries*. Cambridge: Cambridge University Press.
- Erskine, Toni. 2001. "Assigning Responsibilities to Institutional Moral Agents: the Case of States and Quasi-states." *Ethics & International Affairs* 15(2), 67–89.
- French, Peter. 1984. *Collective and Corporate Responsibility*. New York: Columbia University Press.
- Goodin, Robert E. 1985. *Protecting the Vulnerable*. Chicago: University of Chicago Press.
- Goodin, Robert E. 1995. *Utilitarianism as a Public Philosophy*. Cambridge: Cambridge University Press.
- Hindriks, Frank. 2009. "Corporate Responsibility and Judgment Aggregation." *Economics and Philosophy* 25, 161–177.

- Kagan, Shelley. 2011. "Do I Make a Difference?" *Philosophy and Public Affairs* 39(2), 105-141.
- Kukathas, Chandran. 2006. "Who? Whom? Reparations and the Problem of Agency." Pp. 330–341 in Kumar and Tan (eds).
- Kumar, Rahul and Kok-Chor Tan, eds. 2006. *Reparations*. Special Issue of *Journal of Social Philosophy* 37(3), pp. iv–v, 323–482.
- List, Christian and Kai Spiekermann. 2013. "Methodological Individualism and Holism in Political Science: A Reconciliation." *American Political Science Review* 107(4), 629–643.
- List, Christian and Peter Menzies. 2009. "Non-Reductive Physicalism and the Limits of the Exclusion Principle." *The Journal of Philosophy* CVI(9), 475-502.
- List, Christian and Philip Pettit. 2011. *Group Agency*. Oxford: Oxford University Press.
- Mäkelä, Pekka. 2007. "Collective Agents and Moral Responsibility." *Journal of Social Philosophy* 38(3), 456–468.
- May, Larry. 1987. *The Morality of Groups*. Notre Dame: University of Notre Dame Press.
- Miller, David. 2001. "Distributing Responsibilities." *Journal of Political Philosophy* 9(4), 453–471.
- Miller, David. 2004. "Holding Nations Responsible." *Ethics* 114, 240-268.
- Miller, David. 2007. *National Responsibility and Global Justice*. Oxford: Oxford University Press.
- Pasternak, Avia. 2011. "Sharing the Costs of Political Injustices." *Politics, Philosophy & Economics* 10(2), 188–210.
- Pasternak, Avia. 2013. "Limiting States' Corporate Liability." *Journal of Political Philosophy* 21(4), 361–81.

- Pettit, Philip and David Schweikard. 2006. "Joint Actions and Group Agents." *Philosophy of the Social Sciences* 36, 18–39.
- Pettit, Philip. 2007. "Responsibility Incorporated." *Ethics* 117, 171-201.
- Scheffler, Samuel. 1997. "Relationships and Responsibilities." *Philosophy & Public Affairs* 26, 189–209.
- Sher, George. 2005. "Transgenerational Compensation." *Philosophy & Public Affairs* 33, 181-201.
- Simmons, A. John. 1995. "Historical Rights and Fair Shares', *Law and Philosophy* 14, 149–84
- Stapleton, Jane. 2008. "Choosing What We Mean by 'Causation' in Law." *Missouri Law Review* 73, 433–480.
- Stilz, Anna. 2011. "Collective Responsibility and the State." *Journal of Political Philosophy* 19(2), 190–208.
- Thompson, Janna. 2002. *Taking Responsibility for the Past: Reparations and Historical Injustice*. Cambridge: Polity Press.
- Vernon, Richard. 2003. "Against Restitution." *Political Studies* 51, 542–57.
- Walker, Margaret. 2006a. "Restorative Justice and Reparations." Pp. 377–95 in Kumar and Tan (eds).
- Walker, Margaret. 2006b. *Moral Repair: Reconstructing Moral Relations after Wrongdoing*. Cambridge: Cambridge University Press.
- Warren, Patricia Y. and Donald Tomaskovic-Devey. 2009. "Racial Profiling and Searches: Did the Politics of Racial Profiling Change Police Behaviour?" *Criminology and Public Policy* 8(2), 343-369.
- Wenar, Leif. 2006. "Reparations for the Future." Pp. 396–405 in Kumar and Tan (eds).

Wendt, Alexander. 2004. "The State as Person in International Theory." *Review of International Studies* 30(2), 289–316.

Wright, Richard W. 1985. "Causation in Tort Law." *California Law Review* 73, 1788–1813.