

Essentially Aggregative Harm, Restraint, and Collectivization

Political Theory
2024, Vol. 52(1) 34–59
© The Author(s) 2023



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/00905917231185187
journals.sagepub.com/home/ptx



Elizabeth Kahn¹ 

Abstract

Some of the most pressing contemporary social problems result from the amalgamation of a mass of actions that are not intentionally coordinated. Although these *essentially aggregative harms* are foreseeable, it is unclear what moral duties individuals have with regards to them. This paper offers a new analysis of these problems and uses a nonideal contractualist approach to argue in favour of two kinds of duties for individuals. *Collectivization duties* that require individuals to act responsively with a view to ensuring that there are effective governance agents that reliably, fairly, and efficiently prevent these outcomes in the long-term and *duties of restraint* that require individuals to avoid action of a kind that is likely to come together with other actions to cause serious EAH in the immediate future when restraint with regards to actions of this kind could help prevent the outcome from occurring.

Keywords

new harms, responsibility, climate change, structural injustice, coordination duties, collectivization duties, imperceptible difference problem

¹Associate Professor in Political Theory, School of Government and International Affairs, University of Durham, Durham, UK

Corresponding Author:

Dr Elizabeth Kahn, Associate Professor in Political Theory, School of Government and International Affairs, University of Durham, Al Qasimi Building, Elvet Hill Road, Durham, DH1 3TU, UK.

Email: elizabeth.kahn@durham.ac.uk

I. Introduction

Many pressing social problems result from the amalgamation of a mass of actions that are not intentionally coordinated (Lichtenberg 2010; Scheffler 1995).¹ Although these *essentially aggregative harms* (EAHs) are foreseeable and humanly caused, it is unclear what duties individuals have with regards to them.

This essay argues that there are two kinds of ex-ante, pro-tanto duties. When these problems pose a risk of significant harm, individuals primarily have collectivization duties with regards to them. Drawing on Collins' (2013, 2019) concept of collectivization, I propose that individuals have duties to act *responsively* with a view to ensuring that in the long-term there are effective *governing agents* that reliably, fairly, and efficiently prevent these outcomes from emerging. In addition, where significant EAHs are imminent, and collectivization is likely to take too long to address them—and in the meantime, significant damage is likely to be done—individuals have a duty of restraint. This duty of restraint builds on the recommendations of Nefsky (2017) in her analysis of collective impact cases and moral reason. I argue that together these two duties (collectivization and restraint) outline how a conscientious individual should respond to cases of significant EAH in the real world.

My analysis contributes to the existing literature in four ways. Firstly, my definition of essentially aggregative harm brings together a range of cases typically discussed separately and analysed differently in distinct debates in moral philosophy and political theory. I recognise them as sharing common structure and being amenable to a common solution. Examples of EAH include climate change, gentrification, excessive financial risk, microaggressions, transnational labour injustices, global poverty, gender-based issues in medical implant design, battery chicken farming, pollution, engineering faults, and systemic vulnerability to homelessness. These cases have previously been analysed as collective harms (Nefsky 2011, 2015), imperceptible difference cases (Kagan 2011), structural injustices (Young 2011), perfect moral storms (Gardiner 2006), problems of many hands (van de Poel et al. 2012), and moral aggregation problems (Hutchison 2019).

Secondly, the approach I take to assessing candidate duties is innovative and distinct. I adopt an explicitly contractualist framework (inspired by Forst 2012; Scanlon 2000) and use a nonideal approach to assess duties that consider how likely levels of compliance affect what individuals should do.

1. My definition of *essentially aggregative harm* is different to Lichtenberg's (2010) but in keeping with Kahn (2014), as will become clear in section two.

Thirdly, I propose a collectivization solution to the problem that is underexplored in the literature on threshold harms and imperceptible difference cases.² In line with the recommendations of Young with regards to structural injustices, I stress that organising with others to secure political solutions is the best response to these issues. In articulating these duties, I differ from Young's (2011) essentially shared responsibility approach by proposing individual duties to collectivize (Kahn 2018) drawing on Collins' (2013, 2019) conception of collectivization to do so.³

Fourthly, I offer a new argument for preferring a Nefskyan account of duties of restraint over a Parfitian one in cases where EAH is immanent: arguing that a Nefskyan approach should be preferred because it is more likely to protect individuals from EAH given that we can expect significant levels of noncompliance with moral principles that are not imposed by an authority.

Section two and three define the problem of essentially aggregative harm, offer a typology of EAH, and explain the problem EAH presents for moral theory. Section four argues that although governments can offer a more reliable, efficient, and fair solution to EAH than individuals, there is still a need for exploration of what duties individuals have with regards to these problems. Section five outlines the innovative approach used for the evaluation of candidate duties. Section six applies it to potential duties of restraint: arguing that we need a principle that keeps burdens reasonable, distributes them fairly, and can reliably avoid EAH given realistic levels of compliance. Section seven argues *collectivization* offers the best solution to EAH in the long-term. Section eight suggests that with regards to urgent and pressing EAH we must also recognize short-term duties of restraint alongside long-term collectivization duties. The conclusion recaps what has been argued.

2. Problem Analysis

An *aggregative harm* is a morally problematic outcome that results from the amalgamation of a number of distinct acts that are not intentionally

2. This possibility is briefly suggested in Albertzart (2019).

3. This duty draws on the work of Collins (2013). Collins' collectivization duties apply in circumstances where agents can help by collectivizing and there is no existing collective agent able to do so. By contrast, my own account requires contributors to essentially aggregative harm in cases of harm to collectivize in order to prevent the problems they will otherwise continue to help create (Kahn 2018, 2019).

coordinated.⁴ It can be distinguished from harm caused *collectively* by the fact that the contributors lack both the shared decision-making mechanism required to constitute a collective agent and the “we” intention required to be part of informal collective action.⁵ Such harm is *essentially aggregative* if the badness of the morally problematic outcome is more than the sum of the wrongful harm that can be attributed to the individual acts considered one by one (Kahn 2014). In some cases of EAH, the morally problematic outcome is the result of *emergent properties* that only appear when the actions in question *come together*. In other cases, the contributing or constituting acts cause or constitute small amounts of disvalue on their own, but the total disvalue of the outcome is greater than that of the acts considered one by one. In both sorts of cases, it appears that no *agent* or *set of agents* can be identified as outcome-responsible for all the harm done. Hence, these cases present a puzzle for moral theory.⁶ Because EAHs see the occurrence of humanly-caused; foreseeable; avoidable; and morally problematic outcomes, that no agent or set of agents can be held outcome-responsible for, there is a morally problematic backwards-looking responsibility deficit. Worse still, in cases where we have no reason to think the contributory acts are morally impermissible, it is unclear who (if anyone) has an obligation to act differently to prevent these

-
4. A number of factors always come together to cause a problematic outcome. However, in the law (and in ordinary moral thinking) we typically have no trouble identifying who should be held outcome-responsible in standard cases. For example, imagine a case in which a university class is blown up by a bomb laid under their lecture hall by a terrorist. In this case, there are a number of actions that come together to cause the bad outcome: the teaching administrator assigning the classroom, the lecturer directing the students into the room, the students obeying the lecturer, and the terrorist detonating the bomb. However, we have no trouble in identifying the terrorist as outcome-responsible for the injuries and deaths caused by the bombing. For more on this, see Hart and Honoré (1985) as noted by Lichtenberg (2010).
 5. Where a collective agent (or more loosely coordinated collective) causes harm, it is likely that the collective agent can be held responsible for it. Whether and to what extent responsibility can be distributed to members of the collective is a question that has received much attention. For an in depth discussion of collective responsibility see Isaacs (2011).
 6. Following Miller (2007), I distinguish backward-looking *outcome responsibility* (that attributes outcomes to responsible agents) from forward-looking prospective *remedial responsibility* (which identifies agents with a responsibilities to fulfil going forward) (pp. 82–85).

outcomes from occurring. Such cases present a forward-looking responsibility deficit as well as a backwards looking one.⁷

In cases with both sorts of deficit, our intuitions concerning the importance of protecting significant interests from humanly caused degradation suggests that wrongdoing has taken place and yet our intuitions concerning whether any agent has acted wrongfully suggest no agent has failed to fulfil their duties. There is thus a disconnect between what we believe agents should be protected against (their rights) and what we believe agents can reasonably be expected to do (their duties).

An account of morality that leaves agents vulnerable to predictable, foreseeable, avoidable, and humanly caused harm has a significant deficiency. Hence, it is sensible to consider whether there are any justifiable moral demands that, if complied with, would effectively prevent these foreseeable, morally problematic outcomes from occurring. This paper will seek to identify principles that can fairly govern behaviour in this area so as to reliably avoid these harms (without doing wrong), in order to close the forward-looking responsibility gap.⁸

EAHs are not just interesting moral puzzles. Many pressing contemporary social problems (including anthropogenic climate change and forms of global structural injustice) result from the amalgamation of the actions of masses of people. Deciphering the obligations agents have with respect to EAHs is thus of vital importance (Lichtenberg 2010; Scheffler 1995). A theory based on idealized examples will not, on its own, give us answers with respect to particular cases, as they have complex features that must be taken into account when considering the obligations different individuals have in relation to them. However, dealing with the problem of EAH in the abstract will contribute considerably to these efforts.

7. Along similar lines, Van de Poel et al. (2012) identifies some cases of EAH as “many hands problems” that feature a “gap in the distribution of responsibility.” Meanwhile Collins (2017) diagnoses some cases of EAH as featuring “collective duty gaps”.

8. In doing so we also close the backward-looking responsibility gap if we endorse an account of backward-looking responsibility that attributes responsibility for bad outcomes to those who fail to fulfil moral duties that would prevent the problem if fully complied with. Thanks to Stephanie Collins for pointing this out. If this is the case, then a solution to the problem will mean that these problems are no longer EAHs under my definition, as they do not present a backward-looking responsibility deficit anymore. Thanks to Andreas Bruns for pointing this out.

3. Typology

The category of essentially aggregative harm includes three different sorts of case: threshold, insignificant difference, and compound.

In a *threshold case*, individual actions contribute to a variable that causes or constitutes harm when the level of that variable surpasses a threshold or falls between two thresholds.⁹ Threshold cases can be complex, involving multiple thresholds. There can even be cases in which each contribution makes outcomes worse but there are thresholds that, if crossed, make outcomes significantly worse. There could also be cases where a variable is such that contributions make outcomes better at some levels and worse at others. For example, consider how group size effects the quality of seminar discussion: if the group is too small, the debate suffers, but if the group is too large, it again undermines the quality of debate.¹⁰

Anthropogenic climate change is one high-profile example of a complex threshold case.¹¹ If the release of greenhouse gasses continues, it is likely that global temperatures will continue to increase at a faster rate. If the temperature passes key thresholds, significant changes to our environment will result.

-
9. Nefsky (2017, 2746) uses this terminology in a slightly different way: “for each outcome of the morally significant sort in question, there is some precise *number of acts* of the relevant type needed to bring it about: any less will not be enough to bring it about, and any more will not change things with respect to that outcome.” My definition is more inclusive in that it can include cases where different acts make different levels of contribution and what matters is the quantity of some variable that these acts contribute to the level of, rather than the number of acts performed.
 10. It is interesting to note that in some real-life cases of threshold essentially aggregative harm, like hunting a species to extinction for pleasure, the acts that come together to cause harm are themselves wrongful. These cases are still EAHs because much of the significant harm done only occurs once the wrongful acts come together and it is unclear who is outcome responsible for *all* the harm done (Kahn 2014). However, in these cases we can identify agents as acting wrongfully and we can (and do) expect them to cease their wrongful actions. These cases require separate treatment to that recommended here as there is no reason to favour permissibility over restraint, and thus, they do not present a moral puzzle of the same kind as that which is explored here. My recommendations here do not apply to them. My definition includes these cases whilst Lichtenberg’s does not.
 11. It is my understanding that climate change is a threshold case with multiple unknown tipping points—in keeping with Lenton (2008, 1786–93). It is worth noting that some authors instead see it as what I call an “insignificant difference case” (an imperceptible difference case in the standard terminology). Yet other approaches to climate change characterise it along the lines of what I call a “compound case” (see Gardiner 2006). Thanks to a kind reviewer for pointing this out.

If these *tipping points* are crossed, it is likely that much suffering will result and the continuance of life in some parts of the world will be threatened. However, it is not clear what obligations those who make ongoing contributions to these foreseeable problems have. Nor is it clear who can be held responsible for crossing thresholds ex-post. There is a broad literature on this topic that includes Cripps (2011).

We can also imagine complex cases of EAH where there is no precise threshold that if crossed means that wrongful harm is caused or constituted and yet some levels of a variable are significantly better (from a moral perspective) than others. In these cases (whose existence is controversial [see Kagan 2011]), small differences in the underlying variable make no difference at all to the amount of harm done, but large differences make a morally significant difference. I will call these *insignificant difference cases*¹² because each action makes a very small difference to what occurs, but that difference is too small to alter our moral appraisal of the situation. In these cases, none of the actions taken can be identified as outcome-responsible for the overall harm. Furthermore, none of these actions cause the outcome to be worse than it would otherwise be. Nor is there any fact about the individual contributions that can allow us to hold some contributions, but not others, responsible for the bad outcome.

Parfit's (1986) "harmless torturers" example can be understood as such a case. In Kagan's (2011) version of the example, a thousand people turn up dials that each increase the electrical current running through a victim by a very small degree. The total current running through the victim causes them to experience a high degree of pain; however, no additional dial being turned makes a perceptible difference to the amount of pain they feel (Parfit 1986, 80). In such a case, the morally significant variable (the amount of pain endured by the victim) supervenes on the amount of current running through the wire (the underlying variable). The supervenience is such that small differences in the underlying variable have no effect on the morally significant variable, but larger changes do (Nefsky 2011, 375–79). This case is usually described as one in which there are no precise thresholds that trigger a change in the morally significant variable; instead, big changes in the underlying variable cause changes in the morally significant variable and small changes do not. If small changes in the underlying variable do not make a difference to the morally significant variable and individual acts only make a small difference to the underlying variable, there is no

12. These cases are usually referred to as "imperceptible difference cases." However, what matters in these cases is whether the difference is morally significant rather than whether it is perceptible. The two cases come together where the harm in question is pain experienced by a victim, as is the situation in the famous "harmless torturers" case (Parfit 1986, 80).

chance that any individual act will make a morally significant difference. This makes the claim that any of the actors *wrongfully harm* the victim implausible. It also makes it impossible to isolate any act, or subset of acts, and identify them as *outcome responsible* for the victim's suffering.

Real-life examples that could be characterized as insignificant difference cases of EAH include microaggressions, gentrification, and systemic risk in financial markets. In these cases, no individual contribution makes the problematic outcome significantly worse, and there is no precise threshold or set of thresholds that triggers the existence or intensification of the problematic phenomena. Instead, whether the social problem is caused or instantiated is a matter of extent in a complex way.

With regards to gentrification, there is no affluent individual whose purchasing or renting property in an area leads, on its own, to long-term poorer residents suffering the bads of gentrification. These bads include: no longer being able to afford market rents, being priced out of purchasing housing, being socially excluded from local activities and venues, losing the facilities and businesses they use and instead being offered ones they do not like and cannot afford. Furthermore, it seems *ex-ante* that there is nothing wrongful about any individual choosing to move to an area that is affordable and attractive to them: there is no duty not to move to a cheaper area or an area that one perceives as "gritty" or "up and coming."

Similarly, with regards to microaggressions there is no set number of comments, questions, or jokes that if passed leads to injustice. However, large quantities of these acts undermine the social bases of self-respect for stigmatised groups. Likewise, there is no particular number of actions influenced by unconscious bias that if passed make wages, prices, opportunities and security, significantly worse for members of the minority. Instead, all these small acts come together to maintain a social hierarchy that leads to significant inequalities in opportunity and welfare for members of stigmatised groups and undermines their safety. With regards to the contributing acts to this phenomenon, some are benign or negligibly harmful and not wrongful when considered as isolated incidents: it is when they are repeated time and time again and are pervasive features of social life that they cause a significant moral problem. By contrast, other contributory acts are insulting, unfair, annoying, rude, and disrespectful, even when considered in isolation. With regards to these latter acts, we can easily identify duties to avoid them. With regards to the former explaining why they should be avoided and identifying duties to change is more complex and requires a fuller explanation. The analysis here may not immediately apply to this case because we might have reason to want many of these problematic contributing acts to cease rather than continue, even if they can proceed in small quantities without causing the morally problematic aggregate outcome.

A third type of essentially aggregative harms is a *compound case* where harm results from, or is instantiated by, several different kinds of contributory factors coming together. In a compound case, harm is the result of the amalgamation of *qualitatively distinct* factors, such that the harm only occurs when enough of the required kinds of factor are present. In these cases, we have no reason to identify some factors as natural or normal background and others as outcome-responsible for the problem. This is what distinguishes compound EAHs from regular cases of wrongful harm.¹³

Cases identified as a matter of *structural injustice* are typically of this form. Consider Young's example of how working-class urban single mothers are vulnerable to homelessness in the United States. Young identifies multiple social practices, policies, actions, and norms that together lead to systemic vulnerability. Contributing factors include government rules, the aggregate effect of the decisions of businesses and consumers, zoning board decisions, landlord organization policies, pervasive norms, and common attitudes (Young 2011, 59–62). This makes the problem a complex compound EAH.

In many cases of compound or threshold EAH, the problematic outcome is overdetermined. This means that (as is also the case for all insignificant difference cases) removing any particular action does not change whether the problematic outcome occurs or not. This makes it especially difficult to find any party outcome responsible—as no party's action makes the difference between the bad outcome occurring or not. However, it is important to note that EAHs present a problem for moral analysis whether they are overdetermined or not because even when they are not overdetermined, it remains unclear who should be held responsible for the problematic outcome.

4. Government

This essay considers what (pro-tanto) duties individuals have when they foresee that EAH is likely to emerge from the combined actions of an extremely large number of actors. In focusing on the duties of individuals, one obvious solution has been excluded from consideration: attributing governments responsibility for preventing such harm. Governments could fulfil this

13. There could be complex compound cases in which amalgamations of distinct factors have an effect on a morally significant variable. In these cases, it may be that small changes in the make-up of the compound have no effect on the supervenient property but large changes do. If this is possible, there may be compound cases in which no contribution can alter the morally significant factor, but multiple contributions can create a quasi-compound case in a way analogous to the quasi-threshold cases described in Nefsky (2011).

forward-looking responsibility by assigning duties to actors (citizens, private sector organizations, civil servants, and government departments) and enforcing and/or incentivising compliance with these duties in such a way that reliably, fairly, and efficiently prevents foreseeable EAH.

Functioning governments in well-developed states can avoid many foreseeable EAHs, thus based on the Spiderman “power responsibility” principle, we may think we have grounds to attribute responsibility for avoiding these problems to them.¹⁴ Furthermore, one important rationale for having governments is to prevent these kinds of problems from emerging. As noted by Rawls (1977) and Ronzoni (2009), a vital role of government is to coordinate action to prevent the emergence of background injustice from benign individual transactions. Background injustice is a form of essentially aggregative harm. Governments can also regulate to avoid the development of EAHs that involve a “tragedy of the commons” by enforcing a system of rights or directly intervening through their agents to prevent harm resulting from the aggregation of human actions (Kahn 2014). Furthermore, governments determine the shape of peoples’ legal rights and rule in cases of dispute. We need governments to do this to ensure we can exercise the kind of autonomy required for freedom without wronging others (Reglitz 2019). In doing so, governments impose duties that prevent people from together causing aggregative harm to the vital interests of others.

There are distinct advantages to a government-based solution to EAH over individual duties of restraint. Firstly, a coercive government can solve these problems *more efficiently* and *more fairly* than a moral norm for individuals. This is because they can assign a more complex scheme of roles and responsibilities that more efficiently addresses the problem, and they can implement a more complex scheme for paying for these efforts that more fairly distribute costs.¹⁵

Secondly, having a recognized authority issue laws will result in greater compliance compared to having an informal social norm govern behaviour in this area. This is because the government’s authority allows it to determine a solution that will be complied with by many more individuals than would comply with a pervasive social norm. This is due to a combination of factors that are together sufficient to stabilise a solution. Some people accept the state as a legitimate authority with the ability to create binding moral duties

14. Assuming governments are the kind of agents that can have duties and responsibilities. For a discussion of this see Lawford-Smith (2019).

15. For a discussion of how acting together can lower the costs of action and increase motivation, see Lichtenberg (2014).

and thus comply out of a sense of obligation. Other people fear the government's coercion and thus are likely to obey its dictates to avoid sanction. With a government-enforced system, people generally can have more confidence that others will comply and thus that their efforts will not be in vain, and this too will increase compliance. All of which means a governance solution lowers the risk of insufficient uptake to prevent EAH.

Thirdly, in many cases there will be several distinct norms that could be practiced to avoid EAH. This means that in the absence of an authority with the ability to coordinate action, it is likely that different individuals will comply with different norms. Worse still, many are likely to choose to comply with no norm because they doubt others will do so and thus suspect their own efforts will be pointless. An important advantage of government action is that it can create a decisive answer to who must do what, which allows an individual to collaborate with others reliably, expecting them to, in turn, play their part. This means that the state has a better chance of succeeding in avoiding EAH compared to individuals independently complying with what they take to be their duties unilaterally or through informal organisation. For these reasons, we should conclude that a governance-based solution is the most attractive way to tackle serious EAH.

Despite the salience and attractiveness of governance solutions, the role of this paper is to identify the duties of *individuals* rather than governments with regards to EAH. There are two rationales for this focus. Firstly, in the cases that motivate this inquiry, government responses have not been forthcoming: the problems of anthropogenic climate change, gentrification, and overfishing have emerged despite the existence of functioning governments. In some cases, this is because the problems are transnational in scope and thus there is no existing government with the right jurisdiction to tackle these problems. In other cases, problems with political systems, vested interests, power, commerce, corruption, war, and competition could explain governments' failure to act. In cases of government failure, for any of these reasons, it is worth considering what duties individuals and other moral agents have with regards to foreseeable EAH.

A second, reason to investigate the duties of individuals with regards to EAHs is that governments are ultimately human creations that can be brought about or disbanded by people. Governments rely on people for their establishment and continued existence. Given this fact, it is worth considering what moral duties individual people have with regards to EAHs. This should include examining whether individuals have a duty to help to create, maintain, or reform these organisations to address EAH. In doing so, we will gain a theory of individuals' duties with regards to EAH that can apply regardless

of the particular facts concerning governance in their jurisdiction. Such a theory can apply across political contexts in space and time and therefore is more fundamental.¹⁶

5. Nonideal Contractualist Approach

This essay assumes that moral duties are multiple and pro-tanto (Ross 1930). Neither a commitment to nonaggregation (typical of some contractualist approaches) nor a commitment to maximising the good (favoured by rule utilitarian approaches) is assumed. To assess candidate duties, I will examine each one as part of a scheme of pro-tanto universal duties that includes most of the duties recognized in “common sense morality”: duties of nonharm, promise keeping, fairness, reciprocity, charity, promoting justice and the common good, and solidarity. I will adopt a nonideal approach to assessing rival schemes of duties: what will be judged and compared is the likely instantiation of a set of norms rather than the norms considered in the abstract. Thus, general compliance and awareness of the norms alongside partial non-compliance will be assumed when comparing different regimes of duties. This is done so that the analysis can make recommendations based on a realistic view of a reasonable society—one in which most, but not all, people generally comply with moral norms in so far as it is not prohibitively burdensome to do so. Thus, this will be a work of nonideal moral theory in which full compliance is not assumed. As a result, the moral duties recommended will be better suited to real-world scenarios than principles that emerge from more idealised theorising that assumes full compliance.

In each case, what will be examined is whether the proposed scheme of duties can be “reasonably rejected” (Forst 2012; Scanlon 2000). It will be assumed that a scheme of duties can *only* be reasonably rejected if there is an alternative scheme of duties that we have reason to prefer all things considered. I presume that grounds for preferring one scheme over another can be based on concerns of fairness and goodness.¹⁷ Balancing these demands is a

-
16. I think that all collective moral duties need to be derived from individual duties (although they need not be reducible to them). I see collectives (including governments) as a means through which individuals can act in complex ways and achieve complex goals (including the discharge of moral duties) that they could not otherwise fulfil. I realise this is controversial, and thus my analysis here does not rest on readers sharing this conviction.
 17. I am not assuming that these are the only grounds on which a scheme of duties can be assessed; however, they are the concerns that will be appealed to in my analysis here.

perpetual problem for moral and political philosophy. In this essay, I will not propose a particular rubric for determining our duties in all circumstances. Doing so is not necessary to solve the problem at hand and I am sceptical with regards to whether there could be an infallible rubric for determining justifiable norms.

6. Duties of Restraint

Restrictive Approach

One way for a set of moral principles to reliably avoid EAH is to prohibit all actions of types that are known to sometimes come together to cause EAH. This approach involves working out which actions come together with others to cause EAH and requiring agents to refrain from taking these actions at all times. Such an approach is highly restrictive, and thus I will call it “the restrictive approach.”

This principle, if generally complied with, will lead to significant EAH being avoided whenever it is predictable. However, this strategy will place significant restrictions on what agents can do. The costs of refraining from all actions that are known to *sometimes* come together with others to cause or constitute EAH would be very great indeed. This principle would, for example, avoid the bads of gentrification by banning any affluent person from purchasing or renting a home in a working-class neighbourhood because such actions can lead to gentrification if too many people take them.

There are likely to be more nuanced principles that can avoid banning so much action whilst still securely avoiding foreseeable and significant EAH. Other things being equal, these principles should be preferred to the restrictive principle because they allow people more liberty in determining their own lives whilst still effectively avoiding EAH. Those who are asked to give up activities that are valuable to them under a ban on actions that sometimes come together with others to cause EAH could reasonably reject the restrictive principle in favour of a less demanding option that can offer reasonable assurances that significant EAH is avoided (all other things being equal).

Nefskyan Approach

A less restrictive approach would ban actions that are known to sometimes come together to cause or constitute EAH *only* in cases where we have reason to think there is a good chance that an EAH of this kind will in fact occur, and its occurrence will in part result from actions of the kind in question. This approach bans actions of a particular kind when we have good reason to

believe actions of this kind are indeed likely to come together with other factors to cause or constitute EAH in current circumstances.

Nefsky has proposed that in cases where a good or bad outcome results from the aggregation of a mass of actions, there is reason to restrain from “x-ing” in cases where your act of x-ing could be part of what brings about bad consequence Y, and it is possible that Y will not be brought about in part because of a lack of x-ing.

When applied to the question of duties of restraint with regards to cases of essentially aggregative harm, Nefsky’s approach recommends we only have reason to restrict action in cases where the EAH outcome is “up in the air”—where we do not know if it will or won’t occur already. Furthermore, we only have reason to restrict those actions that help to promote the bad outcome (Nefsky 2017).

According to this principle when agents know that EAH is unlikely to result either from their action in combination with those of others, or from actions of the same kind as theirs alongside other actions, they are free to continue to take that action. Thus, they can freely take action that if combined with other actions would cause harm, if they know those other actions are unlikely to take place. Thus, I can explore a wilderness area on foot if I reliably know others will not be doing so, even if the wilderness is of significant value and would be destroyed if explored on foot by thousands of individuals.

However, according to this principle, if it is clear that there is a significant risk that EAH will occur (but it is not certain it will occur) and, if it does occur, its cause will include actions of the kind I am taking, then I have reason to stop taking actions of this kind (Nefsky 2017, 2753).¹⁸ This principle means that in cases where the bad outcome is inevitable and thus cannot be stopped by restraint with regards to the kind of action an agent is contemplating, they are free to act. This solution is more attractive than the restrictive principle as it allows agents to act when they know that their restraint cannot help to prevent essentially aggregative harm. It is significant that, according to this principle, individuals have reason to practice restraint even if the impending EAH is likely to be overdetermined such that their restraint alone

18. Of course, we need details regarding what is meant by “actions of a particular kind.” The idea is to capture acts that could be part of the cause of an EAH (that is likely to occur) but might not end up being the actions that can be identified as causally responsible. This is typically the case when many actions are taken and only some of them are needed for the outcome to occur and only some of them will end up being the cause of the outcome if it does occur.

will make no difference to the amount of suffering that is likely to occur. Nefsky's account does more than recommend restraint where an individual can prevent a bad outcome through their restraint alone.

One notable problem with an approach that forbids all actions that Nefsky notes we have reason to avoid is that it still asks individuals to sacrifice a lot.¹⁹ It calls on all agents to avoid actions that may help to produce EAH whenever such harm is risked, and it could be prevented by sufficient restraint of the kind contemplated alongside other factors. A principle that decreased the amount of restraint required from people whilst still effectively preventing foreseeable EAH would be even better (given that it was not objectionable for some other reason). Such a principle would make some people better off by increasing their freedom and lessening the burdens of fulfilling their obligations without increasing harm to others or the risk of harm to others.

Parfitian Approach

Parfit, in "Five Mistakes in Moral Mathematics," identifies an obligation that could be better than the Nefskyan principle because it requires less restraint. Parfit (1984) proposes that there is an obligation not to take action that is part of a set of acts that *together* harm people (pp. 70, 86). He then notes that this creates a problem regarding which acts should be included in the set identified as harming. He seeks to avoid including irrelevant individuals in the set of actors who can be identified as together harming—by suggesting that it is only *the smallest group* that together makes outcomes worse than they would otherwise be—that should be identified as required to practice restraint.²⁰

What happens when we apply this principle to cases of EAH? General compliance with this principle is a way of avoiding EAH without requiring as much restraint as demanded by the Nefskyan approach. Thus, this principle appears to offer a pareto improvement over Nefskyan approach (improving the liberty of some without increasing the danger to others),

19. Nefsky herself recognises an imperfect duty of restraint (Nefsky, 2021). Whether such a principle can deliver us from EAH depends on the extent to which it is taken up.

20. Parfit proposes his principles as an extension of the act consequentialist account. He aims to supplement act consequentialism to deal with cases in which it gives unintuitive recommendations. In this essay, I take the liberty of treating Parfit's account as a set of principles that should be adopted for the regulation of people (Parfit 1984, 71).

and thus it appears *prima facie* that it should be preferred over Nefsky's principle, all else being equal.

To understand why, consider a mass compound case of EAH in which various moral agents release one of three individually benign chemicals into a lake, and the subsequent reaction between those chemicals creates a gas that kills six people. The case is such that each chemical necessary for the reaction is released by more than one agent but, for the problematic effect to result, we only need some nonzero amount of all three chemicals to be released into the lake. In this case, twenty agents release chemical A, thirty release chemical B, and forty release chemical C. Parfit's approach suggests that the smallest group of agents who could prevent the problem by altering their action should have restrained themselves. Thus, the twenty agents that release chemical A (and only these agents) should be regarded as violating the moral demand not to together impermissibly harm. This is because they are the smallest group of whom it is true that they could have avoided the problem by acting differently. By contrast, the Nefskyan approach would require all agents releasing any of the three chemicals to restrain themselves. This is because each of these agents takes action of a kind that could be combined with other actions in a pattern of restraint that could prevent an outcome from occurring, and at the time of their release of the chemical whether or not the poison would be produced was still "up in the air," it could have been avoided.

Parfit's method minimises the number of people who must restrain themselves. In doing so, it allows for EAH to be avoided whilst inconveniencing the smallest number of people. This is a promising approach to solving the problem of duties to avoid EAH. However, we may still have reason to prefer the Nefskyan principle, all things considered. Following text explains three reasons why this is the case.

Firstly, consider reliability. The Parfitian approach only avoids EAH if there is *full compliance*. In the chemical drop case, if any of the agents who plan to release chemical A violate the duty identified by Parfit's principle, the EAH will still take place. In threshold cases of EAH, the same problem emerges: when our principle only holds the smallest group liable to practice restraint, one single act of noncompliance is enough to pass the threshold and trigger the harm. This issue is particularly worrisome given the role Parfit suggests this principle should play: reforming common-sense morality to avoid the kinds of harm that can emerge from the amalgamation of our actions in mass societies (Parfit 1984, 86).

Even in a fairly moral and upstanding society, we can expect a degree of noncompliance with moral principles due to unwillingness, negligence,

misunderstanding, or disagreement on what is required. This means that a principle that requires 100% compliance to be successful is unlikely to reliably prevent foreseeable EAH. In contrast, principles that call for greater restraint, like the Nefskyan principle, are more likely to reliably avoid EAH given partial noncompliance. Given this fact, those at risk from EAH could reasonably reject Parfit's solution in favour of a set of principles that is more costly but has a greater chance of avoiding EAH in realistic conditions of partial noncompliance.

A second concern with Parfit's principle is that it pays no attention to the way the costs of avoiding EAH are *distributed*. Parfit's principle, as it stands, could be incredibly burdensome to some individuals whilst leaving others relatively well off in a way that is unfair. In such a case, those most heavily burdened could reject this approach in favour of one that more fairly distributes costs if such a principle is available.

Consider again the case of the chemical drop, in this example the smallest group who could avoid the EAH by restraining themselves consists of the twenty agents who release chemical A. Parfit's account suggests that these individuals are obliged to avoid releasing the chemical. If they fail to do so, they are each guilty of wrongdoing because they together harm others. In contrast, according to the Parfitian principle, those who release the other chemicals do nothing wrong, are not duty bound to restrain themselves, and should not be morally criticized for their choices. This seems *prima-facie* unfair. Those who release chemicals A, B, and C stand in a very similar relation to the harm caused. Yet Parfit's theory suggests only a subgroup of them acts impermissibly. It seems to be very much a matter of luck whether an agent is called on to practice restraint in this example. There does not seem to be anything this set has done to warrant their having the burden of preventing the harm rather than some other subgroup. It could be argued that it is *unfair* to ask some contributors to avoid contributing, and not others, when there is no feature of their action that makes them uniquely liable to absorb a disproportionate share of the burdens of avoiding the problem. Those so burdened could object to the fact that they are singled out for additional burdens. Thus, the principle could be criticized for the way it distributes the burdens of avoiding EAH. This is not to say that it would be impossible to justify this distribution of costs and the principle that determines it, all things considered. However, it does suggest more work needs to be done to establish the preferability of this means of preventing EAH over alternatives.

There is one additional issue that must be addressed by Parfit's account: it is unclear how his principle can work in insignificant difference cases. This is a problem that he noted awareness of in his response to critics (Parfit 1986

referenced in Nefsky 2019). Nefsky argues for her account partly because of this issue with Parfit's. In insignificant difference cases, there is no precise number of agents who can together prevent the harm from occurring because there is no precise number of actions that change the outcome. This means that Parfit's principle is indeterminate in these cases. It is unclear how many people must restrain themselves to ensure the outcome does not occur. There is no smallest group that can avoid the outcome through restraint.

Revert to the Nefskyan Approach

For insignificant difference cases, we could revert to the Nefskyan approach: recognising pro-tanto duties of restraint in cases where EAH is risked (but is not inevitable) that fall on all agents who will otherwise take action that will help to cause EAH. Such an approach will lead to significant amounts of restraint but will reliably prevent these insignificant difference cases from resulting in harm. The demandingness of this duty could be mitigated by allowing a few exceptions to these duties to those with strong reason to take action of these kinds (for example due to need). This works as long as these exemptions are not so widespread as to result in a significant risk of EAH in the long-term. An example of such a policy and exemption is requiring people not to drive petrol or diesel cars in the city centre to avoid harm to those with asthma but having an exemption for those with physical disabilities (Abel et al. 2021, 135). Alternatively, we could instead have a more rough and ready principle calling on agents to avoid these acts most of the time. Such a principle allows people to exempt themselves in cases where they have strong reason to act. Along these lines, Nefsky (2021) suggests an imperfect duty calling on people to restrain themselves most of the time but allowing for some contributions.

Reverting to the Nefskyan approach addresses the reliability and fairness issues identified previously with Parfit's approach. The Nefskyan approach treats all those contemplating action that will help to promote EAH equally. Furthermore, by pro-tanto obligating everyone who promotes the outcome, it increases the likelihood of success in conditions of partial compliance. This is particularly true in cases where what we need is a large enough number of agents to practice restraint, and it doesn't matter so much who practices restraint.

Given the vital importance of avoiding EAH, the priority is to ensure reliability. This requires our principle to stand up to partial noncompliance. It also needs for it not to impose the sort of burdens on individuals that undermine compliance levels. Furthermore, who is obliged and what they must do needs to be easy to identify for the duty to be realisable. We have reason to

favour a Nefskyan principle over a Parfitian solution, all things considered. Although less efficient and calling for more restraint, The Nefskyan principle is more reliable in cases of partial noncompliance and fairer in the distribution of burdens. Most importantly, it is easier to identify who has duties as we only need to identify restraint as being of a kind that could prevent the outcome (in combination with other acts of restraint).

However, I do not believe that the Nefskyan principle identifies *all* of the duties we have with regards to EAH. When it comes to the sorts of mass cases that Scheffler, Lichtenberg, and myself are chiefly concerned with, it is my contention that there is another important duty that I have not yet discussed that has been neglected in the literature on collective harm and imperceptible difference. In the next section, I explore the case in favour of recognizing collectivization duties in addition to the duty of restraint identified by the Nefskyan approach.

7. Collectivization

In the face of the threat of EAH, individuals can play their part in solving the problem by complying with the Nefskyan duty: practicing restraint by avoiding actions that promote foreseeable future EAH, most of the time. However, this means taking on significant sacrifices without assurance that this strategy will succeed in avoiding EAH. In such a scenario, individuals will understandably lack confidence that enough others will do likewise to prevent EAH from occurring.

An attractive alternative to such unilateral restraint is for individuals to *work together* to establish and maintain a governance-based solution that fairly and efficiently prevents foreseeable EAH. This means acting in ways that promote a governance solution and acting responsively to those who do likewise. If enough individuals take such steps, a governance solution will be established and maintained. If we think this is what agents should do with regards to EAH, we need to recognize that individuals have *coordination duties* with regards to foreseeable and avoidable EAH. These duties require them to take responsive steps to establish and maintain a fair and reliable collective solution to prevent foreseeable and avoidable EAH.²¹

A coordination duty is not a duty to take an action or set of actions that can be fully specified in advance. Rather, it is a duty to act responsively, adjusting

21. Collins' (2019) later account renames what she previously called 'collectivization duties' 'coordination duties'.

what one does to relevant features of the situation in question and the responses of others to it. Collins develops the concept to explain the duties individuals have in cases where there is a morally pressing task that cannot be fulfilled by any existing agent individually (including collective agents of formal and informal kinds). My account of when and why individuals have coordination duties uses Collins' model to outline the sort of action required by the duty. Coordination duties share with forward-looking responsibilities, like those posited in (Young 2011) with regards to structural injustice, the fact that their focus is on the fulfilment of a particular end or a particular task. However, instead of charging any agent (or aggregate of individuals) with responsibility for ensuring the task is done or the end is achieved, coordination duties require the agent to take responsive steps *towards* achieving the end or discharging the task with others: responding to the steps others take with the aim of making it the case that they together fulfil the task or achieve the end. Such an account has the advantage, over a shared responsibility approach (Young 2011), of avoiding the worry that it will lead to judging an individual based on whether they succeed in producing an outcome that they need others to cooperate with them to achieve.²²

In most cases, by establishing and maintaining a governance-solution, individuals can reliably prevent EAH with less costs than unilateral restraint or spontaneous collaboration between willing agents. In these cases, coordination duties require that individuals ensure there is a collective agent that prevents foreseeable EAH. This makes the coordination duties they have *collectivization duties* (because they require agents to use an existing collective or to produce a new collective and then use it to address the harm).²³

22. For a discussion of this concern regarding Young's account of responsibilities with regards to structural injustice, see Nusbaum's forward to Young (2011) and Kahn (2018).

23. Collins' (2019) work distinguishes two kinds of coordination duties: those that require collective agents (collectivization duties) and those that require informal coordination (responsiveness duties). In cases where there is an obvious, salient, and fair solution and sufficient willingness to prevent the problem, an informal process of responsive action, norm formation, and maintenance can prevent EAH. In these cases, a formal governance agency and the use of coercive power will be unnecessary and thus may be undesirable given the coercion, and bureaucracy they involve. In these cases agents have responsiveness duties (Collins 2019). However, in most cases the best solution to EAH will involve governments preventing it from occurring. This is due to the advantages of a governance-based solution over social-norm-based solutions outlined in section 4.

What this precisely requires individuals to do will depend on the nature and reliability of existing governing institutions and political agents. In some cases efforts should focus on working with others to reform existing governance agencies, whereas in other cases a better approach is to establish new governing agents. With regards to how to bring this about, there will be circumstances in which agents can support or join an existing political party and help it to win power. In other cases, it is a better strategy to pressure the existing government through petitions and protests. In yet other cases, individuals would do best to form new organisations, and then pressure the government or seek power themselves. Which strategy should be taken will depend on concrete details of the scenario faced. Hence, their first step to discharging their collectivization duties will be to investigate the situation and assess possible options for preventing future EAH.

For this collectivization solution to EAH to work, there must be sufficient agents willing to act responsively to both bring about and then maintain governance solutions. This requires some good will and willingness to take on burdens. However, it is my contention that this is not a utopian suggestion that cannot be fulfilled in the long-term.

The argument in favour of collectivization over unilateral restraint relies on the idea that it is realistic to believe that individuals can establish a shared institution and compliance mechanism that can solve EAH without wronging anyone and without causing even worse outcomes. To assert this is to have faith in the potential of governance agencies to do good all things considered (*contra* Hayek 2014) and to believe that utilizing coercion to organize social life can be morally permissible (*contra* Nozick 1974).

There are important arguments suggesting democratic systems of governance incentivize abuse and offer little reason for nonabusers to engage (Brennan 2016). Furthermore, coercive authorities have committed great crimes against humanity and violated rights in the past. These concerns give us reason to be wary of government action to prevent EAH. However, the size and significance of contemporary EAH gives us reason to establish governance agencies to address them. In establishing such solutions, we must be attentive to issues of accountability, balance of powers and thoughtful institutional design in order to limit the threat posed by governance institutions.

The collectivization approach offers a promising solution to EAH. However, questions remain regarding how demanding these duties can be as well as on what basis and to whom they can be assigned. It could be that it is those who will otherwise contribute to threatened EAHs who have the primary coordination duties (Kahn 2013, 2014, 2018, 2019). Alternatively, duties

could fall on all those able to co-ordinate.²⁴ Other options for assigning duties include assigning them to those benefiting from these harms or the actions that contribute to them. Alternatively, a general duty to collectivize could be preferred to minimize burdens for each individual, fairly distribute burdens, and maximize the chance of having sufficient compliance to succeed.²⁵ How we should decide between these options is a complex question that I will not address here.

In most cases of significant EAH, coordination should be preferred to a Nefskyan restraint-based approach in the longer term because it is likely to minimize the burdens of reliably avoiding EAH in the long term and do a better job of fairly allocating burdens.²⁶ This is because a collective solution, in many cases, will be able to find a more efficient way to avoid EAH than straightforward restraint (as discussed in section four). Thus, in many cases of EAH, the Nefskyan principle can be reasonably rejected in favour of collectivization duties.²⁷

8. Imminent Threats

A significant worry with regards to the collectivization solution to EAH is that in cases of impending EAH there is a significant risk that coordination efforts will not proceed quickly enough to avoid significant harm being done. In cases where EAH is imminent and collective solutions do not offer a

-
24. Collins' (2019) own solution works along these lines: basing duties on ability. She identifies agents as having a duty when they are part of a set of individuals that could act together to solve a morally pressing problem and there is no existing agent or other putative group better placed to do so.
 25. Jubb (2012) recommends an equal distribution of duties to all. He argues in favour of this, over assigning duties to contributors, in light of the lack of control contributors have over their contribution.
 26. However, in the short-term it could be incredibly costly as the system resists reform due to vested interests in the status quo, as well as the difficulties in setting up a bureaucracy to ensure compliance or administer services.
 27. Nefsky's (2017) approach to collectively caused outcomes (p. 2753) suggests agents have moral reason to take some action *x* when *x*-ing could be part of a collective benefit, whether the benefit occurs or not is still up in the air, and it is possible that the outcome will fail to come about in part due to a lack of *x*-ing. The concept of collectivization goes beyond this by requiring that agents take a series of responsive actions. This enables complex collective action-based solutions to be recommended rather than discrete actions that could aggregate to produce a collective benefit. This enables the explicit suggestion that what is needed is the establishment and maintenance of a governance solution.

reasonable prospect of quick success, some individuals need to be recognized as having an obligation to practice restraint, until a collective action solution has been established. For this reason, a moral theory that only recognized collectivization duties, with regards to EAH, could be reasonably rejected because it leaves some vulnerable to suffering EAH. Thus, an additional duty must be recognized that requires restraint with regards to actions that are likely to help to cause or constitute EAH in the short term. We should recognize a duty along the lines of the revised Nefskyan principle suggested in section 6 that applies where serious EAH is immanent. My proposal is that this duty operates alongside collectivization duties to produce fair solutions that can prevent EAH long-term. Once fair, just and effective governance solutions adequately protecting against EAH are in place, the moral duty of restraint will not be operative. Instead, individuals will have duties to monitor governance agencies and ensure they do their job, whilst also complying with the duties they are assigned by those agencies. Both these demands follow from their collectivization duties they have with regards to EAH: they describe how individuals can coordinate to establish governing institutions that prevent EAH and, once those institutions are in place, ensure they continue to fulfil this role.²⁸

9. Conclusions

In a globalized world of mass societies, avoiding humanly caused suffering requires a better understanding of the obligations individuals have in relation to essentially aggregative harm. The analysis here argues that they have two types of duties. On the one hand collectivization duties that require individuals to act responsively to establish and maintain collective agencies that fairly, efficiently, and reliably prevent EAH. It has been argued that this solution is best in the long-term because it gives better security that EAH will be avoided whilst also enabling the efficient and fair distribution of the burdens of avoiding it. However, it was noted that in circumstances where serious EAH is immanent and establishing a governance solution will take time, there are also individual duties to refrain from actions that help to promote EAH. These pro-tanto duties (drawing on Nefsky's analysis of collective harm cases) recommend restraint when EAH is threatened, one's restraint is part of a set of omissions that could prevent the EAH when combined with other actions, and whether the EAH will occur is "up in the air."

28. In maintaining that collectivization duties continue even once collective agents that govern are in place, I differ from the approach of Collins and Lawford-Smith (2016).

Acknowledgments

I am grateful to the readers and audiences who offered helpful feedback on the ideas that became this paper at presentations in Bristol, Frankfurt, Glasgow, Graz, Leeds, Leiden, Manchester, Melbourne, Montreal, and York. I want to particularly thank Stephanie Collins, Christopher Finlay, Matthew Rendall, Julia Nefsky, Andreas Bruns, the Economic Ethics Network, the Durham Centre for Political Thought, and the editors and anonymous reviewers at *Political Theory* for their assistance in developing the paper.

Declaration of Conflicting Interests

The author declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: The author received financial support from the *Justitia Amplificata* Centre for Advanced Studies and the University of Durham when developing the ideas for this paper.

ORCID iD

Elizabeth Kahn  <https://orcid.org/0000-0001-9723-1224>

References

- Abel, Will, Kahn, Elizabeth, Parr, Tom, & Walton, Andrew. 2021. *Introducing Political Philosophy: A Policy-Driven Approach*. Oxford: Oxford University Press.
- Albertzart, Maike. 2019. "A Kantian Solution to the Problem of Imperceptible Differences." *European Journal of Philosophy* 27: 837–51.
- Brennan, Jason. 2016. *Against Democracy*. Princeton, NJ: Princeton University Press.
- Collins, Stephanie. 2013. "Collectives' Duties and Collectivization Duties." *Australasian Journal of Philosophy* 91 (2): 231–48.
- . 2017. "Filling Collective Responsibility Gaps." *The Journal of Philosophy* 114(11): 573–91.
- . 2019. *Group Duties: Their Existence and Their Implications for Individuals*. Oxford: Oxford University Press.
- Collins, Stephanie, & Lawford-Smith, Holly. 2016. "The Transfer of Duties: From Individuals to States and Back Again." In *The Epistemic Life of Groups*, edited by Fricker Miranda & Brady Michael, 150–72. Oxford: Oxford University Press.
- Cripps, Elizabeth. 2011. "Climate Change, Collective Harm and Legitimate Coercion." *Critical Review of International Social and Political Philosophy* 14 (2): 171–93.

- Forst, Rainer. 2012. *The Right to Justification*. New York, NY: Columbia University Press.
- Gardiner, Stephen. 2006. "A Perfect Moral Storm: Climate Change, Intergenerational Ethics and the Problem of Moral Corruption." *Environmental Values* 15 (3): 397–413.
- Hartt, Herber Lionel Adolphus, & Honoré, Tony. 1985. *Causation in the Law*. 2nd ed. Oxford: Clarendon.
- Hayek, Friedrich. 2014. *The Road to Serfdom: Text and Documents: The Definitive Edition*. Abingdon: Routledge.
- Hutchison, Katrina. 2019. "Gender Bias in Medical Implant Design and Use: A Type of Moral Aggregation Problem?" *Hypatia* 34 (3): 570–91.
- Isaacs, Tracy. 2011. *Moral Responsibility in Collective Contexts*. Oxford: Oxford University Press.
- Jubb, Rob. 2012. "Contribution to Collective Harms and Responsibility." *Ethical Perspectives* 19 (4): 733–64.
- Kagan, Shelly. 2011. "Do I Make a Difference?" *Philosophy & Public Affairs* 39 (2): 105–41.
- Kahn, Elizabeth. 2013. "Global Poverty, Structural Injustice and Obligations to Take Political Action." PhD dissertation, University of York, York.
- . 2014. "The Tragedy of the Commons as an Essentially Aggregative Harm." *Journal of Applied Philosophy* 31: 223–36.
- . 2018. "Injustice and Collectivization in World Politics." *Global Justice: Theory, Practice, Rhetoric* 11 (2): 29–50.
- . 2019. "A Structural Approach to the Human Right to Just and Favourable Working Conditions." *Critical Review of International Social and Political Philosophy* 22 (7): 863–83.
- Lawford-Smith, Holly. 2019. *Not in Their Name: Are Citizens Culpable for Their States' Actions?* Oxford: Oxford University Press.
- Lenton, Timothy M., Held, Hermann, Kriegler, Elmar, Hall Jim, W., Lucht, Wolfgang, Rahmstorf, Stefan, & Schellnhuber Hans, Joachim. 2008. "Tipping Elements in the Earth's Climate System." *Proceedings of the National Academy of Sciences of the United States of America* 105 (6): 1786–93.
- Lichtenberg, Judith. 2010. "Negative Duties, Positive Duties, and the "New Harms"." *Ethics* 120 (3): 557–78.
- . 2014. *Distant Strangers: Ethics, Psychology and Global Poverty*. New York, NY: Cambridge University Press.
- Miller, David. 2007. *National Responsibility and Global Justice*. Oxford: Oxford University Press.
- Nefsky, Julia. 2011. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan." *Philosophy & Public Affairs* 39 (4): 375–79.
- . 2015. "Fairness, Participation, and the Real Problem of Collective Harm." *Oxford Studies in Normative Ethics* 5: 245–271.

- . 2017. “How You Can Help, without Making a Difference.” *Philosophical Studies* 174 (11): 2743–67.
- . 2019. “Collective Harm and the Inefficacy Problem.” *Philosophy Compass* 14(4): e12587.
- . 2021. “Climate Change and Individual Obligations: A Dilemma for the Expected Utility Approach, and the Need for an Imperfect View.” In *Philosophy and Climate Change*, edited by Budolfson, Mark, McPherson, Tristram & Plunkett, David, 201–21. Oxford: Oxford University Press.
- Nozick, Robert. 1974. *Anarchy, State, and Utopia*. New York, NY: Basic Books.
- Parfit, Derek. 1984. “Five Mistakes in Moral Mathematics.” *Reasons and Persons* 1: 55–83.
- . 1986. “Comments.” *Ethics* 96: 832–72.
- Rawls, John. 1977. “The Basic Structure as Subject.” *American Philosophical Quarterly* 14 (2): 159–65.
- Reglitz, Merten. 2019. “A Kantian Argument Against World Poverty.” *European Journal of Political Theory* 18 (4): 489–507.
- Ronzoni, Miriam. 2009. “The Global Order: A Case of Background Injustice? A Practice-Dependent Account.” *Philosophy & Public Affairs* 37 (3): 229–56.
- Ross, William David. 1930. *The Right and the Good*. Oxford: Oxford University Press.
- Scanlon, Thomas M. 2000. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Scheffler, Samuel. 1995. “Individual Responsibility in a Global Age.” *Social Philosophy and Policy* 12 (1): 219–36.
- Van de Poel, Ibo, Fahlquist, Jessica Nihlén, Doorn, Neelke, Zwart, Sjoerd, & Royakkers, Lamber. 2012. “The Problem of Many Hands: Climate Change as an Example.” *Science and Engineering Ethics* 18: 49–67.
- Young, Iris Marion. 2011. *Responsibility for Justice*. New York, NY: Oxford University Press.

Author Biography

Dr Elizabeth Kahn is associate professor in political theory and director of the Centre for Political Thought at the University of Durham (UK). Her research centres on questions of individual duties with regards to systemic injustice. She has previously published on human rights, global injustice, extreme poverty, and duties in a “tragedy of the commons” scenario. In 2021, she coauthored a textbook (with Abel, Parr and Walton) introducing political philosophy by examining live public policy dilemmas. Her next project will examine questions of legitimacy for activists acting outside of official channels for political participation.