

Article

Open Access

Holly Lawford-Smith

What ‘We’?

DOI 10.1515/jso-2015-0008

Abstract: The objective of this paper is to explain why certain authors – both popular and academic – are making a mistake when they attribute obligations to uncoordinated groups of persons, and to argue that it is particularly unhelpful to make this mistake given the prevalence of individuals faced with the difficult question of what morality requires of them in a situation in which there is a good they can bring about together with others, but not alone. I will defend two alternatives to attributing obligations to uncoordinated groups. The first solution has us build better people, who will coordinate their actions willingly and spontaneously when the occasion arises. The second solution has us build better groups, so that when the occasion arises, there is a framework in place for coordinating members into action.

Keywords: Collective obligation; Unstructured groups; Climate change; Bearers of obligations; Poverty; Non-collectives’ obligations; Humanitarian intervention.

1 Climate Change, Poverty, and ‘Our’ Obligations

Bestselling-author Naomi Klein argues that solving the global climate crisis requires global grassroots social activism (Klein 2014). Environmental activist and *Guardian* writer George Monbiot agrees, advocating widespread political mobilisation (Monbiot 2014). Climate ethicists Edward Page and Axel Gosseries argue that the primary obligations to mitigate climate change fall upon the group of those who have benefited most from industrialisation (Gosseries 2004; Page 2012). These are all normative responses to the actual and prospective jointly-caused harms of climate change. They have in common that they attribute obligations to groups of persons, but they are distinct from what are commonly referred

Holly Lawford-Smith, University of Sheffield and Australian National University, Victoria Street, Sheffield, South Yorkshire, UK, e-mail: h.lawford-smith@sheffield.ac.uk

 ©2015, Holly Lawford-Smith, published by De Gruyter.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 3.0 License.

to as ‘collective obligations’ in that the groups to which they are attributed lack features that collectives have.

These kinds of attributions are made in response to other kinds of harms too. In the context of severe poverty, Martha Nussbaum argues that humanity is obliged to provide everyone with certain capabilities (Nussbaum 2007), Bill Wringer argues that the obligations to satisfy positive rights to subsistence fall on the global collective (Wringer 2005, 2014), and Peter Singer and Garrett Cullity both argue that the primary obligation to address poverty bears on the affluent as a group (Cullity 2004; Singer 2009). In the context of humanitarian military intervention, the ‘Responsibility to Protect’ doctrine attributes obligations to intervene upon severe human rights violations to the international community (United Nations 2009; see also discussion in Pattison 2008).

It is tempting, of course, to move from the joint causation of harm to joint responsibility for harm, in the case of climate change, or from joint failures to provide assistance (causation by omission) to joint responsibility for failure to assist, in the case of extreme poverty (and similarly, from responsibility for joint failures to intervene militarily on humanitarian grounds, to joint responsibility for failure to intervene). But we do not make this move lightly even in the case where we are concerned only with the actions of a single individual: usually we want to know not merely whether she caused some harm, but whether she intended to cause it, or reasonably foresaw that she would cause it, or whether her causation of it was negligent or reckless (Miller 2007). In the case of joint causation it is even more complicated. Anthropogenic climate change was caused (and is being caused still) by the actions of millions of people over the many years since industrialisation, global poverty was caused (and is being caused still) by a complicated set of domestic and international actions and failures to act (Cohen 2010). Many – if not most – of these actions would have made no difference to climate change, or poverty, taken in isolation (Sinnott-Armstrong 2005; Lawford-Smith forthcoming(b)). It is hard to see, then, how those actions could count as intentional, reasonably foreseen, negligent, or reckless.

The objective of this paper is to explain why those authors mentioned above, and others besides, are making a mistake when they attribute obligations to the particular groups of persons they do, and argue that it is particularly unhelpful to make this mistake given the prevalence of individuals faced with the genuine question of what morality requires of them in a situation in which there is a good they can bring about together with others, but not alone. Attributing obligations to such groups makes those obligation statements impotent: it requires action from something that cannot act. This is bad for all the familiar reasons that violating ‘ought implies can’ is bad (Ward Smith 1961). But it is also bad because it does not get us any further along in doing what the claims were presumably

meant to do, namely, explicating the normative implications of large-scale harms such as threatened by climate change, or present in global poverty. We should get our story straight, so that we can actually give useful advice that, if acted upon, would see those problems actually being addressed.¹

Only in a very particular subset of cases, such as those to be discussed in Section 2 below, does identifying as a member of a group like 'humanity' or 'the affluent' make it clear what a particular individual needs to do. In many more cases, what any individual needs to do depends on coordination. Addressing either of climate change or poverty will require collective action. This does not mean jumping straight to a collective obligation, for the reasons I will go on to explain. This paper aims to answer the difficult question of what individuals should do when faced with a situation in which collective action is required but no collective capable of action exists. I will argue that there are two better answers (than the positing of obligations to non-collective groups of persons) to what is required in the face of climate change, poverty, and humanitarian intervention: mutual responsiveness, and collective agency. The first solution has us work to increase our responsiveness to one another, so that we might coordinate our actions willingly and spontaneously when the occasion arises. The second solution has us work to form groups of a specific kind, so that when the occasion arises, there is a framework in place for coordinating members into action.

The paper will proceed as follows. In Section 2 I will present two interpretations of claims about groups' obligations, agreeing that one is plausible and setting aside the cases covered by that interpretation to focus in the paper on the other, and in Section 3 I will briefly clarify my terminology in distinguishing groups that are collective agents from non-collective groups. In Section 4 and

¹ A final further reason to discuss the plausibility of such groups' obligations is that they are a solution to an important challenge to consequentialism. The challenge comes from cases in which some harm will come about as a result of many people's actions, where it makes no difference to the outcome whether *a particular* person acts or refrains from acting. So long as the particular person would get something (enjoyment, entertainment) out of performing the action, she has a reason to perform it and no reason not to, but that is true of all such individuals, so all of them act and the harm is brought about. (To illustrate: one person more or less casting a vote in an election makes no difference, but if no one votes there will not be a legitimate government.) One possible solution to this alleged embarrassment for consequentialism is to say there are *group-based reasons* to act or not act, because what the group does makes a difference even if what the individual does makes no difference. Shelly Kagan (2011) sets this solution aside because of the conflict that would arise between the reasons one has as an individual compared with the reasons one has as a member of a group, but he does not comment on what seems to be a bigger problem, namely giving reasons (duties, obligations) to the kinds of groups who seem to lack any ability to act on them (this problem also arises for Sartorio 2004).

Section 5 I will present two apparently plausible reasons for thinking non-collective groups can have obligations, the first to do with reliable production of outcomes, the second to do with responsibility for what one causes. I will give a quick summary of the case against non-collectives' obligations at the end of Section 5, clearing the way for a focus on collective agents, and individuals' duties to create them, in Section 6. There I defend the two solutions mentioned above, namely mutual responsiveness and collective agency. In Section 7 I will conclude.

2 Interpreting Claims about Groups' Obligations

Neither Monbiot nor Klein sound like they are saying something particularly implausible when they speak or write about political mobilisation or grassroots social activism as the solutions to anthropogenic climate change or extreme poverty. Monbiot finished his 2014 SPERI Lecture by saying:

So you say to me, "Well what do we do instead? You produce these arguments against trying to save nature by pricing it, by financialisation, by monetisation. What do you do instead?" Well, ladies and gentlemen, it is no mystery. It is the same answer that it has always been. The same answer that it will always be. The one thing we just cannot be bothered to get our bottoms to do, which is the only thing that works. Mobilisation. It is the only thing that has worked, the only thing that can work. Everything else is a fudge and a substitute and an excuse for not doing that thing that works. And that applies to attempts to monetise and financialise nature as much as it does to all the other issues we are failing to tackle.

The question is *what* 'we'? 'We' cannot be bothered; 'we' need to stop making excuses; 'we' must mobilise. There are two ways to interpret this claim. The first is as ascribing an obligation (or set of obligations) to *every* person in the relevant set. The relevant set here might be individuals passionate about environmental activism, it might be all individuals remotely capable of activism whether they are passionate about the environment or not. Let us assume it is the latter. More important than how to demarcate the relevant set is whether obligations are being attributed to all members or only some, and to all unconditionally or only conditionally. If they are being attributed to all unconditionally, then what Monbiot is really saying is that *every* person remotely capable of environmental activism must perform actions that could plausibly be described as 'mobilising'. The problem is that mobilising, like collectivising, is not something any person can do alone (I will talk about this more later in the paper). Furthermore, even if it was (let us say we counted a one-person protest, a one-person blockade, a one-person defacing of a billboard, etc., as an act of mobilising) it is far from clear that it is *true* that every person ought to do those things. It might be enough that a

sufficient number of people in the relevant set perform the actions, which would create a collective action problem given the potential desires of some (or all) to be the ones exempted from the subset; it might be positively harmful for more people than necessary to perform the actions, say because they will get in each other's way and hinder the attempt; and it might be futile for one or a mere few to act without wider support.²

To give a toy case to make the point, suppose you are with two friends, and the three of you come across a broken-down car in need of a push-start. It would take all three of you pushing to move the car, which is large, and stopped on a slight incline. But unfortunately you know the driver, who your two friends dislike immensely. While you chat with the driver about what happened, your two friends make it absolutely clear that they do not intend to help, and walk a little further down the road where they sit belligerently, shooting malicious glances back at the driver from time to time. It cannot be the case in such a situation that you should go alone to push ineffectually at the back of the car. It is true that it would be good if the three of you gave the car a push-start, and it is true that were the other two to push at the back of the car you should push too. But from neither of those does it follow that you have an unconditional obligation to push at the back of the car. What you ought to do – whether you ought to do anything at all – depends on what your two friends do. Such an obligation is at best conditional: if they do – or if they signal willingness to do – then you ought to do. Your obligation to act depends on the satisfaction of the conditional's antecedent.

There are only two kinds of cases for which it follows straightforwardly that all persons in the relevant set must perform actions. The first is the kind of case in which it would take a contribution from every person to secure the good. For example, imagine that a local supermarket will prohibit the use of plastic bags in order to reduce the number of bags going to landfill only if none of its customers resist the change: any single complaint would result in the prohibition not being made. In that case when we say 'the supermarket's customers should not complain', we really mean 'for every single one of the supermarket's customers, it is the case that she should not complain'.

² For discussion of cases structured in each of these ways, see my (Lawford-Smith 2012). Note that in that paper I am concerned with what is required of individuals who are members of genuine collective agents, asking how collectives' obligations *distribute* to their members. I am concerned in this paper with groups that do not count as collectives, and therefore do not bear collective obligations in the first place. So the question here is not one about distribution, but those same structures (joint-necessity, threshold good, threshold-with-harm, incremental good) occur in thinking about the kinds of contributions that individuals would have to make together in order to secure some particular outcome.

The second is the kind of case in which every person's actions make a difference to the securing of the relevant good. INGOs have made huge efforts in recent years to communicate to individuals the good that their charitable donations will secure. Suppose that the Against Malaria Foundation were to make a claim of a similar flavour to Monbiot's or Klein's, for example they named a figure and asked that the global community donate it in order to provide enough Long-Lasting Insecticidal Nets (LLINs) for everyone affected. This does not have the same structure as the plastic bags case, because there are a great number of different ways in which different constellations of individuals could make donations together amounting to the figure named by the Foundation, and on which at least some people would not have to donate anything at all. (Securing the end is one thing, securing it according to the optimally fair distribution is another). But it has a different structure that similarly permits the interpretation that all individuals should donate, namely that the world is a bit better for every individual who does so. The more funding that is provided – short of the figure named by the Foundation – the more LLINs can be distributed to those affected, and in turn the more malaria is prevented. So when the Against Malaria Foundation says 'the global community should provide N billion dollars' they are also saying 'every person should make a donation – the more the better'.

If a case has either of the two structures just mentioned,³ then it is reasonable to interpret the claim about an obligation of an uncoordinated group as being equivalent to a claim about the obligations of each member of that uncoordinated group.⁴ Whether poverty has this structure depends in large part on whether we think of it as a holistic problem ('global poverty') or a cluster of local problems. On neither model would a solution require *every* person in the world to act. On the latter model it will *sometimes* be true that the more people who act, the better; and it will sometimes be false. In the wake of some natural disasters,

³ I take it that these will be the *only* cases where it is reasonable to interpret the claim such, for the consequentialist reason that it being better if more people rather than less perform the action is the *explanation* of why all ought to perform it (unless the alternative explanation holds, namely that the good outcome is tied to *all* performing it). But a non-consequentialist may resist this conclusion on the grounds that there are some cases where people have duties to act even if their action would not make a morally relevant difference to the outcome. For example, a non-consequentialist may wish to interpret the claim "left-leaning Scots should vote Green!" as being equivalent to the claim that any left-leaning Scot should vote Green (regardless of what the others are doing, regardless, even, of whether the election is already won, or already lost). For such people, this will provide a further kind of case in which non-collective obligation claims are equivalent to claims about all members of the non-collective.

⁴ Peter French endorses this interpretation of claims about the obligations of uncoordinated groups in (French 1975) and (French 1979, p. 210).

for example, the challenge to relief efforts is not raising enough money via donations, but rather organisations getting through to provide the necessary support (facing difficult logistical challenges, bypassing corrupt governments, etc.). In such cases it will not be true that the more people who donate, the better. Climate change is not such that it requires every person in the world to act, and *neither* is it necessarily such that the more people who act the better (see discussion in Lawford-Smith forthcoming(b)). At least when it comes to taking actions to reduce personal greenhouse gas (GHG) emissions, or undertaking political actions in order to change our government's environmental policy, we come face to face with the futility problem. What we need to do is act together, not for any one person to act regardless of what everyone else does.

Thus in the rest of the paper I will set aside all the cases in which it is plausible to read an attribution of an obligation to a non-collective group as shorthand for attributing obligations to every member of the group, and I will focus on the cases in which it would be futile for members of the non-collective group to act alone, or in which the actions of a subset of the members would be sufficient to securing the good, or in which the actions of more than a subset of the members would create an obstacle to securing the good. In cases such as these, we are left with a problem in making sense of the attribution of an obligation to that group. I will argue that the attribution itself is a mistake.

3 What Kinds of Groups Can Bear Obligations?

Many people accept that there are some groups which can be the bearers of obligations. The first question to ask is what separates a group which can be the bearer of obligations from a group which can not be. The next is which camp groups like those mentioned earlier – 'humanity', 'the global collective', 'the affluent', 'the global community' – fall into. The usual way to answer the first question is to defend some set of conditions for collective agency, such that when a group meets those conditions it can be the bearer of obligations, and when it does not it can not. Because 'ought implies can', groups have to be able to fulfill the alleged obligations in order for those obligations to actually bear on them. The conditions for collective agency are supposed to capture what it would take for a group to be able to fulfill obligations. There are many proposals (e.g. Held 1970, p. 479; Preda 2012), out of which I find the most plausible – although needing some important modifications – to be those in (Pettit and Schweikard 2006). They say: "A group of individuals will constitute an agent, plausibly, if it meets conditions like the following. First, the members act jointly

to set up certain common goals and to set up a procedure for identifying further goals on later occasions. Second, the members act jointly to set up a body of judgments for rationally guiding action in support of those goals, and a procedure for rationally developing those judgments further as occasion demands. And third, they act jointly to identify those who shall act on any occasion in pursuit of the goals, whether they be the group as a whole, the members of the group individually, certain designated members, or certain agents that the group hires” (Pettit and Schweikard 2006, p. 33).

This set of conditions is overly strong, in that it requires individuals to ‘act jointly’ at each stage, which in turn requires each individual to have the intention that the group produce a particular end, the intention to do their bit in the production of that end, the belief that each of the others intends to do her bit in the production of that end, and the intention of doing their bit because of that belief (Pettit and Schweikard 2006, p. 23). But surely one individual can act as, for, or on behalf of the group, so that the conditions for collective agency can be fulfilled without requiring joint action at each stage.

With the required modifications, designed to weaken the conditions in the direction just mentioned, the conditions (which I take to be necessary) will be:

1. An individual (alone or together with others) establishes: a set of non-conflicting goals that are intended to be common, and a procedure for assigning roles whose fulfillment will enable the pursuit of those goals.
2. At least two individuals (usually including the individual(s) in 1.), intend the goals, the procedure, and the outputs of the procedure to bear upon their own decision-making goals and procedures.
3. At least one individual from 2. or 1. assigns roles – to the individuals in 2. and in accordance with the procedure in 1. – whose fulfilment will enable the pursuit of the goals in 1.

Forming such a group is not particularly arduous; the most important thing is the intention of the individuals in 2. to take the outputs of the group’s decision-making procedure into account in their own decision-making. This does not mean the group’s decisions have to have trumping weight. 2. might be met even if the individual considers fulfilling his role as a member of the group but ultimately decides to pursue a project more important to him personally.⁵

⁵ Note that both the original conditions and this modified version require a group decision-making procedure. See (Collins 2013, p. 244) for another alternative formulation of the original Pettit and Schweikard conditions. For the view that groups other than those with a shared decision-making procedure can have obligations, see (Feinberg 1968; Held 1970); for the view that individuals should act *as if* they were in such a group see (Woodard 2007, Ch. 4 and 5).

It is clear that the kinds of groups mentioned earlier (again: 'humanity', 'the global collective', 'the affluent', 'the global community') will not meet the modified conditions proposed above (nor do they meet the original Pettit and Schweikard conditions), and therefore do not count as collective agents.⁶ If it were clear-cut that only collective agents could bear obligations then this would be the end of the discussion. But I want to take up the challenge in Wringe (2014, p. 174–177), to not simply take it as self-evident that only agents can be the bearers of obligations. Several authors have denied that there must be collective agency before there is collective (or at least 'joint') obligation. Explicit defenders of this view include Bill Wringe (2010, 2014), Felix Pinkert (2014, forthcoming), Anne Schwenkenbecher (2013, 2014), and David Killoren and Bekka Williams (Killoren and Williams 2013). I take the most interesting and difficult challenge to come from the cases discussed by Virginia Held (1970) and Stephanie Collins (2013), which I discuss in the next section under the label of 'mutual responsiveness'.

In what follows, I will discuss what I take to be the strongest argument in favour of non-collectives' obligations. I will argue that ultimately, the most plausible cases can be explained in terms of individuals' obligations to act responsively, and that we should dismiss talk of non-collectives' obligations in favour of (a) individuals' obligations, to be responsive and/or to form collectives capable of acting in pursuit of the relevant goods, and (b) collectives' obligations (once formed) to pursue those goods.

4 Spontaneous Coordination and Mutual Responsiveness

In the next two sections I want to talk about a couple of sequences of reasoning that I think partially explain the thought that non-collectives can bear obligations. In this section I will focus on a sequence that goes as follows. Sometimes outcomes are brought about by the cumulative actions of individuals, and on some theories of action the bringing about of that outcome by the non-collective will be classified as a non-collective's action.⁷ Actual implies possible *a fortiori*, so if 'ought implies can' were to be read as requiring only unrestricted possibility,

⁶ From here I will use 'collectives' instead of 'collective agents', and 'non-collectives' for groups not meeting these conditions.

⁷ See the more detailed discussion of action in Section 5.

then these actions of non-collectives would be candidates for obligatoriness, namely, to be such that they ought not to be performed, or such that if they are performed the non-collective is responsible after the fact for their performance. Now we need only to generate the in-principle obligation, for example, ‘university students ought to protest against rising fees’. We check the outcome of university students in fact protesting against rising fees against ‘ought implies can’ (on the weak reading) and notice that it is not knocked out. There we have a non-collective’s obligation.

There are two versions of this step from actual to possible, and possible (plus desirable) to obligatory. One is implausible in a way that we can probably all agree on, and one is much more challenging. The implausible version goes as follows. A local cattle farm will go out of business if it receives orders from less than 60% of its regular customers this month (and this will be a good thing, in terms of methane emissions). Quite by chance, without coordination between its regular customers (who do not know each other), and as a low-probability confluence of different factors, those who would normally have ordered from the farm are each prevented from doing so: some fall sick, some have to travel, some are gifted produce from alternative sources, etc. (see also examples in Jackson 1987, p. 93). By fluke, ‘withhold custom sufficient to drive the cattle farm out of business’ became an action of the group’s. Because it was actual, we can see that it was possible; because it was possible, we can see that it would not have been ruled out as impossible if put forward as the content of an obligation of the group’s.⁸

A response to this reasoning can be given by way of a strengthening of ‘ought implies can’. Theorists interested in a possibility constraint for obligation do not usually – ever? – have in mind mere logical, metaphysical, or nomological possibility. A simple demonstration: one of the university students’ flatmates locks her in the bathroom on the day of the big fee protest, and her friend calls to say “where are you? You ought to be here protesting with us!” The student responds “I can’t! Tamati locked me in the bathroom”. It would be counterproductive for the friend to say “well you can, it’s perfectly consistent with the laws of nature that you be here rather than there right now”. The ‘ought implies can’ constraint is meant to knock out obligations in cases where the in-principle obliged individual lacks the ability, or the opportunity, to act (for a fairly recent survey, see Vranas 2007). This more circumscribed account of ‘can’ that features in ‘ought implies can’ is specifically about agency. Properties such as ‘ability’ and ‘opportunity’ do not apply to non-collectives, because the counterfactuals usually used to assess

⁸ A version of this reasoning appears in Gunnar Björnsson’s ‘Coinciding Clean Hands’ case (Björnsson 2014, p. 110), although he later explicitly rejects ‘mere chance’ as being sufficient to justify an obligation (Björnsson 2014, p. 114).

whether they apply come out as nonsensical. Can the non-collective group 'university students' protest? Let us see. Would it tend to succeed in protesting with a probability greater than 0.5 (or higher than a contextually-specified threshold) if it tried? Hard to say, for it is hard to say what its trying might involve. This invokes mental properties familiar from the case of individual agency, of which there are functional equivalents in the case of collective agency, but which are conspicuously absent in the case of non-collective groups.

But what if we rejected this model of understanding groups' abilities in a way analogous to our understanding of individuals' abilities, and instead looked for something like reliable production? An action meets the 'can' constraint if, in a high proportion of nearby worlds, it would be performed. Returning to the example of the cattle farm, most of the counterfactuals that both hold the relevant features of the regular customers' situation fixed and include the successful driving of the cattle farm out of business will come out as false, because a small amount of variation to the conditions will be enough to disturb the low-probability confluence of factors necessary to spontaneously producing that outcome. For example, bad weather might have seen some of the travel plans disrupted; sickness might have missed one or more customers or been less severe; etc. That is why the non-collective's action just described is not particularly difficult to dismiss as a candidate for being obligatory. The action 'can' eventuate, but only flukily.

There is a much more challenging case, involving mutually responsive individuals, in which it looks like non-collective groups' actions can arise reliably.

Consider the following, from Virginia Held (1970):

Assume that there are seven apparently normal persons in a subway car: none is acquainted with any other; none are sitting together. The second smallest person of the seven rises, pushes the smallest to the floor, and, in full view and hearing of the remaining five, proceeds to beat and strangle his victim. [...]he remaining five persons do nothing for, say, five minutes, at the end of which time interval the smallest person is dead. [...]No one of the five, acting alone, could have subdued him; it is extremely probable that action by two or more of the group to subdue him would have succeeded, with no serious injury to themselves; the group was not so numerous that, if each member had acted, confusion would have resulted (Held 1970, p. 476–477).

Held thinks that the persons in the subway car can be assumed to share the end of preventing the attack. Any reasonable person would want the attack to not be happening in the first place, anyone being attacked would want the bystanders around them to help them. She also thinks there is no particular mystery about what needs to be done. The assailant needs to be restrained; it does not matter very much how exactly that is achieved (whether the largest person pins his arm back and the second-largest sits on his legs, or whether it is the second-largest

who takes the arms, for example). Held has also made sure to eliminate any of the usual preventers of action, such as undue risk of being hurt, or hindering the rescue. Thus, the five persons in the subway car not involved in the attack have the obligation together to restrain the assailant and protect the victim, and if they fail, as they in fact do in the case as Held imagines it, they are culpable together for that failure. Her case generalises: accidents and attacks are often the kind of thing that can be assisted or prevented by persons acting together; it is usually reasonable to assume that persons will want assistance to be provided, and prevention to be achieved; and it is usually more or less clear what needs to be done. In such cases, clusters of individuals can act together without being coordinated by one of their number (or an external agent) (see also the beach rescue case in (Collins 2013)). Felix Pinkert and Anne Schwenkenbecher have both argued with reference to Held's case that groups of agents that are not themselves agents in their own right can yet be the bearers of obligations (Schwenkenbecher 2013, 2014; Pinkert 2014).

Extrapolating from this, we can see that in certain kinds of (idealised) societies, non-collectives' actions in pursuit of the good will be commonplace. All persons in such a society have strong commitments to particular values as a matter of character, integrity, or personal virtue. All persons in such a society are well-versed in what is required in situations where those values are compromised. Every now and then, a situation would arise in which individuals could advance the good only by acting together, and this situation would act as a kind of trigger: the persons in that society would simply step into the familiar roles and the situation would be resolved. To illustrate, imagine that members of a certain cultural group had internalised the value of being prepared to die to defend their own and others' freedoms. Then imagine that this group is the subject of a civil war, and taken hostage in forced-labour camps by members of the dominant political group. Even without the capacity to communicate and so coordinate an uprising, we might imagine that all it would take for there to be an uprising is any single prisoner disobeying a rule or order: a single act of resistance could trigger similar acts of resistance in all other prisoners, because all other prisoners are prepared to die to protect their own and others' freedoms, and the triggering action is sufficient to 'coordinate' others around that particular moment and location for the uprising. This shows that the adoption of different moral norms by groups of persons – in this case unconditional/categorical rather than conditional/consequentialist – will affect the extent to which we judge that a non-collective can reliably perform an action. So non-collectives can act reliably. If that is all it takes to meet an alternative version of 'ought implies can' not modelled on human mental states, then non-collectives' obligations are back in the picture.

But lest I have now made the case for non-collectives' obligations look too promising, let me now turn to the problems. The first is that mutual responsiveness only works in a subset of the cases where we might want non-collectives' obligations. It works (and sometimes works very well) in cases where everyone can be presumed to share an end, and where what each person needs to do is roughly the same (or where the differentiated roles are a matter of common awareness). But arguably the most pressing contemporary case, namely widespread action against climate change, is not like that at all. Granted, there are things that individuals can do independently and unilaterally. But if we were to figure out what ideally ought to be done by 'humanity' (all people everywhere) and then distribute the roles necessary to see that achieved, the roles would surely have nothing like the content individuals come up with when acting independently and unilaterally (*contra* Jamieson 2007). Still, a subset is something – showing that mutual responsiveness does not lead to reliable fulfillment of ostensible non-collective obligations in all cases is not a major blow when the battle is being fought over whether there are any non-collective obligations at all. One option, obviously, would be to concede those cases, and modify the strong claim 'non-collectives cannot bear obligations' to a weaker claim, 'there are very few cases in which non-collectives can be the bearers of obligations'. I still prefer the stronger claim, as I hope the discussion makes clear.

Two responses are available to the challenge posed by this subset of cases. The first is to say that the conditions that make mutual responsiveness possible simply *suffice for* collective agency. The group of individuals is in fact coordinated, just in the rather unusual way structured by social norms and common knowledge rather than by leaders or decision-making procedures. Referring back to the conditions outlined in Section 2, social norms and the broader social context would do the work of assigning 'roles... whose fulfillment will enable the pursuit of the goals'; the facts about important values being compromised would do the work of 'goals... intended to be common', and the individuals' dispositions toward being good people would take care of 'intend[ing]... the goals, the procedure, and the outputs of the procedure to bear upon their own decision-making goals and procedures'. This would be to concede that there are the grounds for obligation in these cases of what looked like non-collectives' actions (actions brought about by mutually responsive members), but to show that the grounds are there precisely because there actually is the kind of agency or agency-replacement necessary for a group to be the bearer of obligations. In such cases, the groups count as collectives, rather than non-collectives.

The second, which I am more sure about, is that we can tell the full story about what ought to be the case without recourse to attributing obligations to non-collectives. Mutual responsiveness increases the chances of successful

avoidance of harm; it makes the world better. But responsiveness is a property of an individual: she has it when she notices that a desirable end could be produced only by several people acting together, and she coordinates her own actions around what she sees others doing. A responsive student notices that other students are gathering and painting signs for the protest, she joins them and takes up the task of attaching the finished signs to long pieces of wood in order to make them into placards. Nobody needed to direct her to do so, she simply looked at what was being done and what needed to be done and made a judgement about where her efforts would be most useful. When several responsive individuals are clustered together and a triggering situation arises, mutual responsiveness will likely result in successful group action. If we want to see more successful group action on climate change, or poverty, or humanitarian intervention, we either need to work towards being more mutually responsive (and making others more mutually responsive), or we need to work towards there being more coordinated groups, capable of action in their own right.⁹ The latter strikes me as slightly more reliable, for the reasons already given that sometimes the roles are not clear and coordination is needed. But exactly what we do not need is to simply claim that non-collectives have obligations. This will be unhelpful when this does not tell any member of such a group what to do, and it will be unnecessary when the evaluative picture together with the dispositions of the relevant persons already make it likely that action will be forthcoming (see also discussion in Collins 2013).

5 Action, Responsibility, and Normative Symmetry

The second sequence of reasoning that seems to vindicate non-collectives' obligations goes as follows. One of the things that creates forward-looking responsibility, i.e. the responsibility to discharge particular duties or obligations, is backward-looking responsibility, i.e. causation of or causal contribution to the states of affairs the duties or obligations are aimed at remedying.¹⁰ Individuals

⁹ Here the 'we' is those who care about securing the morally relevant ends, e.g. mitigating and adapting to climate change, alleviating poverty, intervening militarily on humanitarian grounds when doing so is necessary. I have not said that there are duties to work towards mutual responsiveness, or there being more coordinated groups, for their own sake. I have said only that if there are duties relevant to the end of mitigating and adapting to climate change, they are not duties of non-collective groups, but rather duties of individuals, to be more responsive, or to take steps toward collectivising (see further discussion in Section 6).

¹⁰ Forward-looking responsibility may be created by other things too; here I need only that causation of harm is one of those things. See discussion in Miller (2007).

can have obligations (I will switch to using 'responsibility' for the backward-looking and 'obligation' for the forward-looking in order to avoid confusion between the two) as a result of what they have caused, *even when* they lacked the relevant kind of control over what they caused. So too for non-collective groups. Some outcomes can only be produced by the cumulative actions of individuals, and when those outcomes are produced, the non-collectives who produced them can have obligations as a result of what they have caused. But responsibility (retrospective) and obligation in virtue of causation (prospective) are symmetrical: one cannot have an obligation as a result of performing an action that one had no obligation not to perform. If the non-collective is responsible, then it had an obligation. It is responsible, therefore it had an obligation. University students failed to protest and the fees went up; they are responsible (as a non-collective group) for the fees going up; therefore they must have had an obligation to protest against the fees going up. Again, we seem to have indicated a non-collective's obligation.

The first thing we need, to get this idea off the ground, is a permissive account of action, such that the cumulative actions of individuals can count together as a single action. Frank Jackson has argued that a number of actions performed by distinct individuals can in some cases be considered together to be an action (Jackson 1987, p. 93), so for him there is no problem in describing groups of individuals' actions themselves as actions. Sara Chant's discussion of unintentional collective action supports non-collectives' actions, with several persuasive examples (Chant 2007, p. 249–250). Peter Dalton (1993) argues for 'extended action', actions made up of sub-actions performed over an extended temporal or geographical period. Philip Pettit and David Schweikard propose conditions for 'joint action' (2006, p. 23–24). On their analysis, joint actions are actions performed by groups whose members have certain beliefs and intentions. So long as a non-collective group is composed of members who each satisfy the relevant belief and intention conditions, there will be joint actions attributable to the non-collective group itself, where there is not necessarily also collective agency (2006, p. 23, 33, and 35).¹¹

The second thing we need is an answer to the question of whether responsibility (backward-looking) is premised on mere causation, or causation plus some further thing (such as control). Considering only the case of individual human agents for a moment, this question has been divisive in discussions of moral luck. Are individuals responsible for everything they cause, even if they were asleep, temporarily insane, blinded by a jealous rage, drugged, stressed, framed by

¹¹ Note that one consequence of being permissive in this way about the application of the term 'action' will be that action and control come apart in a way they might not have on a restricted linguistic usage, which may have serious implications for responsibility. See also (Held 1970, p. 476).

others, overtired, misled, drunk, non-culpably ignorant...? The question has led Michael Zimmerman, for example, to propose a distinction between the 'scope' and the 'degree' of responsibility for what persons bring about, where two persons behaving identically may end up with responsibility of the same degree (indexed to control) but differing scope (indexed to outcomes) (Zimmerman 2002).

If we think control matters, we will think two individuals with responsibility of very different scope (let us say both fired their guns, but only one of their bullets hit and killed the intended victim, so one is responsible in scope for a death, the other only for an attempt) nonetheless have responsibility of the same degree, namely for that part of the world they exercised control over: the firing of the gun at their victims with the intention of killing them (Lewis 1989). If we think causation is all that matters, we will think the two individuals have responsibility of very different degree: one is responsible for a death, one for an attempt. This would presumably legitimise differential blame and punishment (as in fact we actually see in practice in most modern legal systems).

So far, we have several accounts of action permissive enough to count clusters of actions that bring about particular outcomes (usually those that involve harm, or fail to avert harm) as non-collectives' actions, and we have an account of responsibility that tracks causation alone – say, the causation involved in a non-collective's action bringing about a harmful outcome. Does this get us all the way to non-collectives' obligations? Before I answer that (but to kill the suspense, the answer is 'no!') let us put the issue in terms of a case, to bring out the important elements of the question more clearly.

Imagine that one hundred different individuals on one hundred consecutive days drop litter into a very small lake, and by the hundredth day the fish population is finally killed off by the pollution. Making use of Jackson's or Dalton's accounts of actions (Pettit and Schweikard's account is much stronger and will exclude this case), we can say the non-collective group of those hundred individuals 'polluted the lake', and thereby brought about the outcome of having 'killed off the fish population'. Making use of the 'causation alone' account of responsibility, we can say that the non-collective group is responsible for what it did. But this raises two interesting questions, one of which is problematic enough to be a reason to reject the claim that there is responsibility in such a case.

The first question is what the non-collective group's responsibility means for those who compose the non-collective. Surely, given that none of the hundred individuals were acting in full knowledge of each other's actions or with the intention to kill off the fish population (let us say they had good reason to think they were the only one littering in that way), they cannot be held equally responsible (i.e. dividing the overall harm of the loss of the fish population between the hundred of them in blaming, or punishing, or seeking reparation). It seems that

the only permissible distribution of responsibility back to the non-collective's constituents is proportional to each constituent's contribution. But in that case we did not need the story about what the non-collective did, and what the non-collective is responsible for. Each person is responsible for what she did, but that does not add up to the sum total of what the non-collective caused. (In other words, each of the hundred is responsible for littering, but the blame, punishment, and reparation appropriate to littering will very likely not add up to be proportional, in total, to the killing off of the fish population, which is a permanent future deprivation for those who would have made use of the lake). Non-collectives' obligations do not add anything worth having.

The second question is whether the symmetry between responsibility for actions and obligations not to perform those actions does not rather imply that we made a mistake in attributing responsibility. If there is symmetry, then whenever there is responsibility for an action after the fact, there will also have been an obligation to not have performed that action before the fact. And in reverse, wherever there is an obligation not to perform an action, there is responsibility for having performed it if that obligation is violated. The case for symmetry strikes me as strong: it is unfair to hold a person accountable (to blame her, to punish her, to require reparation from her) for an action that she had no obligation not to perform. Thus when we suspect responsibility, we should assume also obligation, and then have a look to see whether that would have been plausible. The non-collective presented above polluted the lake and killed off the fish population. It seems it is responsible for that. But that would mean it had an obligation not to pollute the lake and thereby kill off the fish stock. But what 'it' is that? There was just an uncoordinated group of individuals, acting independently and unilaterally, who just happened to engage in similar enough behaviour at a similar enough time and place, that their actions taken together caused a harmful outcome – the pollution of the lake and collapse of the fish population. The only things in the vicinity capable of bearing obligations are the individual constituents of the non-collective. Thus by *modus tollens*, there is no responsibility in this case. Precisely: if there is responsibility, then there is obligation. There is no obligation (by ought implies can, no bearer of the obligation so no capability of fulfilling it) therefore there is no responsibility.¹²

¹² An alternative to denying that there is an obligation, compatible with the symmetry argument, is to say that there is an obligation in virtue of other features. All that is precluded by the symmetry argument is that there is an obligation in virtue of there being responsibility (causation). That leaves open that the non-collective has the obligation because it is 'best-placed'. Of course, I would deny that it can have the obligation in virtue of these other features, because the fact that it lacks control would preclude those equally. But the *possibility* is there.

Of course, the way to get around this is to reject symmetry (see discussion in Miller 2007). Then we get an obligation not to perform an action without being committed to there being responsibility for having performed that action; and *vice versa*, if we want it, responsibility for having performed an action without being committed to there having been an obligation to not have performed it. Without a commitment to symmetry, the critic of non-collectives' obligations would not be forced to either deny there is responsibility in order to deny there was an obligation (as I am), or explain the obligation in virtue of some other features (although I would deny non-collectives' obligations *tout court* for all the reasons explained so far). Those engaged in the debate over moral luck (e.g. Zimmerman 2002) believe there is responsibility even when there is no control. But they do not say much about whether there were obligations not to perform the actions for which there is then responsibility. Whether or not they explicitly reject symmetry, they appear to have failed to notice that there is a strong case for it.

Those who think that our moral obligations are subjective, tied to our epistemic states (or our 'reasonable' epistemic states), will deny that there are any obligations under many of the conditions I mentioned above, such as temporary insanity, non-culpable ignorance, or having been drugged. Those who think our moral obligations are objective, and hold regardless of what we know or believe, may allow that there are obligations under these kinds of conditions (on this distinction see discussion in Smith 2010a,b). They can excuse the violation of objective obligations in some cases, and they can explain responsibility as present in those cases where an excuse is not available. But it seems that in application to the non-collective presented above, there would be a pretty good case for excusing, namely something approximating non-culpable ignorance seems to be present in the fact that the non-collective's constituents were not organised, acted independently and unilaterally, did not expect that others would be acting in the same way, did not assign a high credence to the proposition that their own littering would be an actual cause of the eventual pollution of the lake and collapse of the fish population, and so on. So even accepting asymmetry (although I think we should not) does not put non-collectives' obligations on solid ground. On a subjective reading there would be no obligations because there is a functional equivalent of non-culpable epistemic failure (the group is not an epistemic agent so this is hardly surprising); on an objective reading even a *pro tanto* obligation would likely be excused, for the same reason.

In summary, although we might want to attribute actions to non-collectives under certain conditions, we are not entitled to move from actions to obligations. Some actions eventuate without being produced in the way we usually think necessary for obligation, e.g. control, or reliable production. In such cases, we might agree of the action that the non-collective produced it, without agreeing that the

non-collective ought not to have produced it. Each individual constituent of a non-collective can be obliged not to perform certain actions, both unconditionally, and conditional upon her beliefs about what the others will do, depending on the case, but the individuals together (the non-collective) cannot be obliged to not perform certain actions, or not bring about certain outcomes, because they (it) cannot; there is no infrastructure for coordination, or organisation, or even the mere dissemination of information.

For all the reasons just given, talk of non-collectives' obligations is at best superfluous – in that it does not add anything more than talking about the obligations of each individual would – and at worst misleading, in that it perpetuates the idea that non-collectives can be the bearers of obligations, which in turn encourages impotent solutions to major collective action problems like climate change, and cooperation problems like widespread poverty, in which non-collectives like 'the global populace' or 'humanity' are ascribed an obligation to prevent global climate change, and to end global poverty. Such non-collectives are not able to act (although they are made out of things that are able); such non-collectives lack the control over their actions necessary for obligations (even though sometimes they do produce actions); and there is nothing that they ought to do, because there is nothing that they could do (even though in societies of mutually responsive persons, there is something each could do to reliably produce the relevant outcome).

6 Duties to Collectivise

What is left if we dismiss the idea of non-collectives' obligations? Where there are not collective agents, we should be clear about the obligations that exist – for individuals – because of the goods that could be pursued if there were. I have already mentioned that one solution to failures of joint action in cases where joint action could secure morally important goods is to work to increase our responsiveness to one another, to work toward making sure that social values (and thus what counts as a violation of those values) and the roles necessary to resolving a violation are well-publicised. When individuals are willing to act together and they know what they each need to do in order to secure a joint outcome, all we need is a triggering event to get joint action off the ground. This solution will be good enough some of the time, but as I said earlier, it does not work in cases where the roles are diverse and complicated, where coordination or organisation are required, and where the values are not settled or violations are non-obvious. For that reason, it is worth thinking about an alternative solution. I will develop

the suggestion in Collins (2013) that individuals have duties to take steps to collectivise, in the first instance. First we must take steps to bring into existence the kinds of groups that can reliably produce desirable outcomes; then those groups will come to bear obligations to act in pursuit of morally important goods (such as the mitigation of climate change, and the alleviation of widespread poverty) – and these obligations will in turn have implications for the groups’ members. The kinds of groups that can reliably produce desirable outcomes are collective agents. Duties to take steps to collectivise are duties to work towards the creation of collective agents.¹³

These are obligations that fall upon individuals, to take steps towards collectivising, e.g. by signaling conditional willingness to cooperate (Lawford-Smith forthcoming(a)), acting as a trigger or responding to another’s triggering when others have signaled conditional willingness (this step takes us from willingness to action), persuading, remonstrating, and motivating others to take individual steps towards collectivising. Once a collective agent is brought into existence, it can use formal and informal mechanisms to distribute roles to members jointly sufficient for seeing its obligations fulfilled. In terms of outcome, there is no disagreement with Wringer (2014),¹⁴ who is one of the most vigorous defenders of non-collectives’ bearing of obligations. But the way of getting there is very different. Wringer thinks the non-collective group that is the global community bears an obligation to secure the subsistence rights of the poor, and that this obligation distributes to the members of that group – everyone, everywhere – in one of two ways, one of which is to coordinate in the way required to allow the fulfilling of others’ subsistence rights (Wringer 2014). These are essentially individuals’ obligations to collectivise, but on Wringer’s story such obligations are grounded in the obligations of the wider group, whereas on my story (and also that of Collins 2013) such obligations are grounded in the good that could be secured *were there* groups capable of acting in particular ways.

A defender of non-collectives’ obligations might be less than enthusiastic about replacing non-collectives’ obligations with individuals’ collectivisation duties, for at least three reasons. First of all, there is something deeply morally problematic about our response to challenges like climate change, or in the way we fail to respond to urgent situations when we find ourselves among other capable persons who are not responding. If all the projected harms of climate change come to pass, or if a person is beaten to death in front of twenty witnesses

¹³ Although they can also be duties to join, reform, and check existing collective agents or their goals and procedures. See discussion in Collins and Lawford-Smith (forthcoming).

¹⁴ Although there is a serious disagreement with Wringer (2010), which argues that those who bear obligations are not necessarily those responsible for fulfilling them.

any combination of which could have helped, we somehow want to be able to say that many people failed together, that there is something they could have done as a group that would have avoided those negative outcomes. If all we can say is that many individuals failed in their individual duties to collectivise, it seems like we miss our target. What has gone wrong in the world is that there has been avoidable suffering, not that persons merely failed to, e.g. signal to one another in the ways that they should have.

The critic of non-collectives' obligations – me! – has to bite this bullet. We are still able to make evaluative claims, e.g. to lament the fact that there has been suffering. But we are not permitted the prefix 'avoidable', because we do not think that there was any agent (or non-agent for that matter) under whose control the outcome fell. That means the only normative claims we can make are attributions to individuals of collectivisation duties (or attributions of responsibility in light of their failure to act upon those duties). I suspect the fact that this is so dissatisfying is exactly what has led theorists to propose non-collectives' obligations. But dissatisfaction does not justify saying something false.

Second of all, the motivation for believing collective obligations to be disjunctive once distributed to members (in the way defended at greater length in Lawford-Smith 2012) might be thought to apply equally to collectivisation duties. The critic of non-collectives' obligations must resist ascribing unconditional duties to individuals when there is some good that could be pursued were all such individuals to have (and fulfill) categorical duties. That is because in many of these kinds of cases, what any one individual ought to do depends on what the others do. But so too for collectivisation duties. Collins (2013) acknowledges that an individual cannot 'collectivise' alone (that is why she talks in terms of the duty to take steps towards collectivising, and why I have been careful to unpack collectivisation duties in more detail here as duties to signal, to trigger, to persuade). But there is more to this objection than that observation. The idea is that just as duties to perform a part of a collective action must depend on whether others will do their respective parts, so too must collectivisation duties. Just as there is no point forming the first link in a human chain to rescue a drowning person in the knowledge that others will not come and complete the chain, there is no point signaling conditional willingness to cooperate in the knowledge that others will not signal their own.

There are two responses to this worry. The first is that there is a conceptual difference between the two. Individuals cannot produce particular collective outcomes alone, while they can fulfill the duties related to collectivising alone. There may well be situations in which an individual should not fulfill those duties, for example when the costs of doing so would far outweigh the benefits, or when the case has been rigged so that the collective action will surely fail if she performs the

actions (imagine that if she signals conditional willingness to cooperate she will be shot, a situation which is actually close to reality in e.g. forced labour camps or under political regimes with dense networks of spies). But she can perform her collectivisation duties in a way that she often cannot perform the action-disjunct of her distributed collective obligations.

The second is that discharging individual duties to collectivise is significantly cheaper than discharging distributed collective obligations (again focusing on the action-disjunct). Signaling conditional willingness might be as cheap and easy as nodding your head in response to another's raised eyebrow, or adding your name to an online coordination forum. The fact that they are so cheap and easy to discharge means that there will be fewer occasions for failure to discharge them (fewer instances in which as *pro tanto* 'oughts' they are weighed up against others and yet do not emerge as the all-things-considered ought; less instances in which as all-things-considered oughts their fulfillment would come at disproportionate cost, entailing that failure to fulfill them would be excused).

Finally, why think that the fulfillment of collectivisation duties is likely to lead to the existence of a collective, and subsequently to collective action? One need only be moderately sceptical about the prospects of interpersonal cooperation to think that all relevant individuals could discharge their duties related to collectivisation and nonetheless there could fail to be a collective or collective action. Well-intentioned and willing-in-principle individuals might still fail to agree upon the collective's aims, or on a way of distributing roles to members, or about which members should take a leadership role. We might all be willing in principle to intervene upon an assault to save a victim's life, but none of us is willing to be the person to get closest to the assailant, in case we get seriously physically hurt. In some situations there are natural leaders – individuals with more charisma, or authority, or forcefulness of character. But in other situations there are not, and in those it seems that even if all relevant individuals signaled their willingness to cooperate, and one such individual triggered that willingness into action, there might still be anarchy with respect to what the group should actually do, and how it is to be done.

In response, I think a proponent of collectivisation duties must concede that they are not a surefire way to secure collective action. The more structure there is already in place, the more likely it is that the fulfillment of collectivisation duties will lead to collective action, e.g. a constitution or manifesto, established social norms that dictate which roles must be filled for the performance of a given task, processes for the election of leaders or executive subgroups, clear and specific ends. It also seems that the more flexible potential members are, the better (the failure imagined here is largely a result of strong personalities failing to agree on how collective ends should be secured). The less structure there is already

in place, and the less flexible potential members are, the less likely it is that the discharging of collectivisation duties will lead to collective action. There might be ways to strengthen the account of collectivisation duties so that their fulfillment more reliably produces collective action, but it is hard to see exactly how these would go (for example, a duty to 'be more flexible' is unappealing, first of all because it is about what persons should be like rather than what they should do, and second of all because if all persons in a group were maximally flexible there would arise just the same problem as if all persons were maximally inflexible, namely nothing would get done). It would be useful to think more about the mutual responsiveness mentioned in the last section to this end, because that only requires an initial actor, around whom others can coordinate their actions.

7 Conclusion

There are better and worse motivations for attributing obligations to non-collectives. The worse motivations simply fail to see an alternative way of accounting for the obvious moral importance of persons acting together against certain harms than to attribute obligations to act to the non-collective itself. The better motivations simply assume that because (on some theories) there is non-collective action, there is also non-collective responsibility and – by symmetry – non-collective obligation; or because there is action, and in some cases even reliable action, there is candidacy for obligation, and in combination with desirability, actual obligation. But the better motivations make important mistakes, and the worse motivations are unnecessary because there is an alternative way to theorise about the obligations that exist in light of goods that could be produced were there collective agents, or were there mutual responsiveness of the right kind within a population. That is to think about individuals' duties to collectivise.

Acknowledgments: This work is supported by the European Commission and the Australian Research Council. I am grateful to audiences at the Cooperation & Equality Research Group at the University of Manchester, the Centre for Ethics & Metaethics and Centre for Aesthetics Joint Seminar at the University of Leeds, and the Philosophy Seminar at the University of Sheffield; and to Stephanie Collins, Wolfgang Schwarz, Geoffrey Brennan, David Estlund, Jimmy Lenman, John Maier, Richard Healey and Filippo Ferrari for helpful discussion; to David Estlund, Anca Gheus, Ryan Cox, David Birks, Stephanie Collins, Mark Budolfson, and Frank Hindriks for written comments on the (many!) prior drafts of this paper; and to two anonymous reviewers for the *Journal of Social Ontology* for helpful comments and suggestions.

Bibliography

- Björnsson, Gunnar (2014): “Essentially Shared Obligations”. In: *Midwest Studies in Philosophy* XXXVIII. p. 103–120.
- Chant, Sarah (2007): “Unintentional Collective Action”. In: *Philosophical Explorations* 10. No. 3, p. 245–256.
- Cohen, Joshua (2010): *Philosophy, Social Science, Global Poverty, in Thomas Pogge and his Critics*. Cambridge: Polity.
- Collins, Stephanie (2013): “Collectives’ Duties and Collectivization Duties”. In: *Australasian Journal of Philosophy* 91. No. 2, p. 231–248.
- Collins, Stephanie and Holly Lawford-Smith (forthcoming): “The Transfer of Duties: From Individuals to States and Back Again”. In: Michael Brady and Miranda Fricker (Eds.): *The Epistemic Life of Groups*. Oxford: Oxford University Press.
- Cullity, Garrett (2004): *The Moral Demands of Affluence*. Oxford: Clarendon Press.
- Dalton, Peter (1993): “Extended Action”. In: *Philosophia* 24. p. 253–270.
- Feinberg, Joel (1968): “Collective Responsibility”. In: *Journal of Philosophy* 65. No. 21, p. 674–688.
- French, Peter (1975): “Types of Collectivities and Blame”. In: *The Personalist* 56. p. 160–169.
- French, Peter (1979): “The Corporation as a Moral Person”. In: *American Philosophical Quarterly* 16. No. 3, p. 207–215.
- Gosseries, Axel (2004): “Historical Emissions and Free-Riding”. In: *Ethical Perspectives* 11. p. 36–60.
- Held, Virginia (1970): “Can a Random Collective of Individuals be Morally Responsible?”. In: *Journal of Philosophy* 68. No. 14, p. 471–481.
- Jackson, Frank (1987): “Group Morality”. In: Philip Pettit, Richard Sylvan and Jean Norman (Eds.): *Metaphysics and Morality: Essays in Honour of J.J.C Smart*. Oxford: Blackwell.
- Jamieson, Dale (2007): “When Utilitarians Should Be Virtue Theorists”. In: *Utilitas* 19. No. 2, p. 160–183.
- Kagan, Shelly (2011): “Do I Make a Difference?”. In: *Philosophy and Public Affairs* 39. No. 2, p. 105–141.
- Killoren, David and Bekka Williams (2013): “Group Agency and Overdetermination”. In: *Ethical Theory and Moral Practice* 16. No. 2, p. 295–307.
- Klein, Naomi (2014): *This Changes Everything*. New York: Simon & Schuster.
- Lawford-Smith, Holly (2012): “The Feasibility of Collectives’ Actions”. In: *Australasian Journal of Philosophy* 90. No. 3, p. 453–467.
- Lawford-Smith, Holly (forthcoming(a)): “Unethical Consumption & Obligations to Signal”. In: *Ethics & International Affairs*.
- Lawford-Smith, Holly (forthcoming(b)): “Difference-Making and Individuals’ Climate-Related Obligations”. In: Clare Hayward and Dominic Roser (Eds.): *Climate Justice in an Idealised World*. Oxford: Oxford University Press.
- Lewis, David (1989): “The Punishment That Leaves Something to Chance”. In: *Philosophy and Public Affairs* 18. No. 1, p. 53–67.
- Miller, David (2007): *National Responsibility and Global Justice*. Oxford: Oxford University Press.
- Monbiot, George (2014): “The Pricing of Everything”. *The Guardian*, 24th July 2014. Online at www.monbiot.com/2014/07/24/the-pricing-of-everything/.

- Nussbaum, Martha (2007) *Frontiers of Justice*. Cambridge, MA: Harvard University Press.
- Page, Ed (2012): "Give It Up for Climate Change: A Defence of the Beneficiary Pays Principle". In: *International Theory* 4. p. 300–330.
- Pattison, James (2008): "Whose Responsibility to Protect? The Duties of Humanitarian Intervention". In: *Journal of Military Ethics* 7. No. 4, p. 262–283.
- Pettit, Philip and David Schweikard (2006): "Joint Actions and Group Agents". In: *Philosophy of the Social Sciences* 36. p. 18–39.
- Pinkert, Felix (2014): "What we Together Can (Be Required To) Do". In: *Midwest Studies in Philosophy* XXXVIII, p. 187–202.
- Pinkert, Felix (forthcoming): "What If I Can Not Make A Difference (And Know It)". In: *Ethics*.
- Preda, Adina (2012): "Group Rights and Group Agency". In: *Journal of Moral Philosophy* 9. p. 229–254.
- Sartorio, Carolina (2004): "How To Be Responsible For Something Without Causing It". In: *Philosophical Perspectives* 18. p. 315–336.
- Schwenkenbecher, Anne (2013): "Joint Duties and Global Moral Obligations". In: *Ratio* 26. No. 3, 1–19.
- Schwenkenbecher, Anne (2014): "Joint Moral Duties". In: *Midwest Studies in Philosophy* XXXVIII. p. 58–75.
- Singer, Peter (2009): *The Life You Can Save*. New York: Random House.
- Sinnott-Armstrong, Walter (2005) "It's Not My Fault: Global Warming and Individual Moral Obligations". In: *Advances in the Economics of Environmental Research* 5. p. 293–315.
- Smith, Holly (2010a): "Subjective Rightness". In: *Social Philosophy and Policy* 27. No. 2, p. 64–110.
- Smith, Holly (2010b): "The Moral Clout of Reasonable Beliefs". In: Mark Timmons (Ed.): *Oxford Studies in Normative Ethics*, Vol. 1 Oxford: Oxford University Press.
- United Nations (2009): Implementing the Responsibility to Protect: Report of the Secretary General. 12th January. Online at http://www.un.org/en/ga/search/view_doc.asp?symbol=A/63/677 accessed 11th April 2014.
- Vranas, Peter (2007): "I Ought Therefore I Can". In: *Philosophical Studies* 136. No. 2, p. 167–216.
- Ward Smith, James (1961): "Impossibility and Morals". In: *Mind* LXX. No. 279, p. 362–375.
- Woodard, Christopher (2007): *Reasons, Patterns and Cooperation*. New York: Routledge.
- Wringe, Bill (2005): "Needs, Rights, and Collective Obligations". In: *Royal Institute of Philosophy Supplement* 80. p. 187–206.
- Wringe, Bill (2010): "Global Obligations and the Agency Objection". In: *Ratio* XXIII. p. 217–231.
- Wringe, Bill (2014): "From Global Collective Obligations to Institutional Obligations". In: *Midwest Studies in Philosophy* XXXVIII. p.171–186.
- Zimmerman, Michael (2002): "Taking Luck Seriously". In: *The Journal of Philosophy* 98. No. 11, p. 553–576.