

On individual and shared obligations: in defense of the activist's perspective

GUNNAR BJÖRNSSON, STOCKHOLM UNIVERSITY

1. The activist's perspective

To resolve issues like climate change, massive inequality, widespread racism, or threatening fascism, a great many people must contribute. People who make substantial efforts to help resolve such collective practical problems often take what I will call “the activist’s perspective,” thinking that:

COLLECTIVE EFFICACY: Together, we can significantly resolve the issue (at an acceptable cost).

COLLECTIVE OBLIGATION: In virtue of COLLECTIVE EFFICACY, we have a moral obligation to do significantly resolve the issue.

RESULTING INDIVIDUAL OBLIGATIONS: The considerations that ground this collective obligation also ground individual *prima facie* moral obligations to help resolve the issue, at least when there is a real chance that it will be significantly resolved.

The first two elements of the activist’s perspective attribute efficacy and obligations to some relevant group of agents. (Typically, the group is one comparatively well placed to make a difference. In the case of climate change, for example, it might include all affluent contemporaries; in the case of racism members of the racist society in question.) Importantly, these attributions seem to be *non-distributive*: they do not imply that each individual member can significantly resolve the issue or has an obligation to do so. By contrast, the third element does attribute obligations to the individuals in the group.

COLLECTIVE EFFICACY is exactly the sort of fact that motivates activism (see e.g. Van Zomeren, Postmes, and Spears 2008), likely in part because it is seen as giving rise to collective and individual obligations. Admittedly, explicit ideas about “obligations” might play little role in some activists’ thinking about collective practical problems. But consider failures to resolve such problems: to significantly decrease the risk or extent of climate catastrophe, or prevent a fascist takeover of a constitutional democracy, say. My impression is that the typical activist thinks that many such failures would be failures to live up to what is morally required of us: if we can resolve these issues, *we have to*. Similarly, the perspective involves the sense that one would be at least somewhat to blame if one did nothing to help in spite of being aware of the issue, particularly well placed to contribute to its resolution at a low cost, and not already occupied with similarly important tasks.

Though prevalent and potentially crucial in motivating activism, the activist’s perspective can seem deeply philosophically problematic.

2. Problems of group agency and individual inefficacy

The first problem has to do with the sheer scale of the collective issues in question:

INDIVIDUAL INEFFECTICACY: With few exceptions, no action by an individual human agent can significantly affect whether the issue is resolved.

For example, however successful we will be in preventing or reducing the threat of global climate catastrophe, the extent of that success is unlikely to have been much affected by *my* actions. The same seems true for most moral agents, including most that the activist's perspective is concerned with. In light of INDIVIDUAL INEFFECTICACY, it might also seem that

NO INDIVIDUAL OBLIGATION: Because of INDIVIDUAL INEFFECTICACY, the importance of resolving the issue at hand and our collective capacity to do so cannot ground individual obligations to act in ways aimed at achieving resolution (see e.g. Sinnott-Armstrong 2005; Nefsky 2011; but cf. Nefsky 2016).

This directly contradicts RESULTING INDIVIDUAL OBLIGATIONS. Of course, individuals might have other reasons to contribute: perhaps contributing might show solidarity with others who contribute, or be a way to take on a fair share of the burdens involved (for discussion, see e.g. Brennan and Sayre-McCord 2015). But this does not seem to fully capture the obvious teleological elements of the activist's efforts to contribute to the collective end.

The second problem with the activist's perspective concerns COLLECTIVE OBLIGATION. As already noted, it should not be understood as implying that any individual in the group has an obligation to resolve the issue, as such resolution is beyond the capacity of any or most individual agents. But the claim also seems deeply problematic if it attributes obligations to *the group*. So understood, it seems to conflict with the appealing thought that

AGENCY REQUIREMENT: Only moral agents are subject to moral demands.

The AGENCY REQUIREMENT does not rule out all attributions of obligations to groups. Many have argued that groups of agents can be organized such that they together constitute a collective moral agent with its own beliefs, goals, rational point of view, and capacity for self-regulation, free choice, moral thinking, and moral motivation (see e.g. French 1984; List and Pettit 2011; Björnsson and Hess 2017; Hess 2014). Moreover, it is clear that massive problems have been resolved when people have organized in such complex ways, not least by forming states. However, it is not clear how a group can have obligations when, like the groups of COLLECTIVE OBLIGATION, it has not yet organized in this way.

Perhaps the AGENCY REQUIREMENT should be rejected. Requirements on obligation bearers have been formulated largely based on intuitions about individuals, with scant attention to group cases. In light of this, it makes sense to similarly start from the pre-theoretic sense that unstructured groups can have obligations under the right circumstances, and then proceed to ask what these conditions are, and whether there is a unifying set of conditions for

individual and group cases.¹ Still, since it might turn out that there is no non-arbitrary set of such conditions for group cases, and since standard requirements of individual moral agency seem non-arbitrary, this assumption is defeasible. To put the worry to rest, we need a plausible account of the relevant conditions.

One reason to think that a unifying account might be available is that what is standardly required for moral agency are *capacities*, along with opportunities to exercise such capacities: capacities to respond to moral reasons by forming goals and plans and to reflectively regulate one's responsiveness to reasons, say. Moreover, there is surely a sense in which many groups of people have the capacity to do this, exactly because they can get organized in ways allowing them to do so. Given this, it is tempting to simply assume that these groups are moral agents, subject to moral obligations.²

The capacity response is not without problems, however. The extent to which it can be morally demanded that an agent responds to certain reasons and self-regulates presumably depends in part on whether and if so how easily the agent can do so. This in turn depends on whether the agent would do so under some relevant range of circumstances. For self-organizing, massively homeostatic systems like human beings operating in fairly stable environments, we have ways of roughly identifying such a range, based on the agent's normal organization and normal circumstances and our familiarity with difficulties involved in adjusting one's responsiveness to various kinds of reasons. But attributions of difficulty of self-regulated reasons-responsiveness becomes considerably more obscure when we turn to the sort of yet-to-be-organized groups that concern us. This is not to deny that there are intelligible questions to ask about how difficult it would be for a group to do something. For example, we might ask how hard it would be, on average, for members of a group to do something if all members lived up to what could reasonably be demanded of them. But the most obvious such questions do not seem to concern the capacity of the group as an agent of its own.

Moreover, even a successful defense of the capacity approach to group agency looks unlikely to help make sense of the activist's perspective. The groups with which COLLECTIVE EFFICACY and COLLECTIVE OBLIGATION are concerned are not themselves understood as agents, but exactly as pluralities of agents. Activists are concerned with how to reach out and engage individual agents, have in mind how difficult it is for individual agents to contribute, and think about what can be achieved if individuals can be motivated by moral urgency of the issue at hand. Moreover, it is not central to the perspective that all relevant agents organize around some one plan or accept the same decision-making procedure. Climate catastrophe

¹ The contemporary discussion often starts with (Held 1970). For a recent argument, see e.g. (Schwenkenbecher 2014, 63–65). Wringe (2016) makes an explicit appeal to “moral phenomenology.” For one recent book-length treatment of the specific case of climate change, see (Cripps 2013).

² Relatedly, Bill Wringe (2016, 13–14, 2010, 221–24) suggests that a group's potential to become an agent might be enough for obligations and Tracy Isaacs (2011, ch. 5) suggests that groups of the right kind can have “putative” collective obligations that can serve to guide the actions of moral agents in much the way that ordinary obligations do, i.e. one that could underpin RESULTING INDIVIDUAL OBLIGATIONS.

prevention, for instance, will realistically happen through a wide variety of actions motivated by the threats of climate change, actions which, although often locally coordinated, will not necessarily be guided by some global master plan or procedure to which all or even most agents sign on.³

The task of this paper is to explain how, in spite of problems of group agency and individual inefficacy, COLLECTIVE OBLIGATION and RESULTING INDIVIDUAL OBLIGATION are coherent and sensible, focusing on the case of threatening climate catastrophe. In previous work, I have suggested that we should understand the relevant kind of collective obligation as *shared* obligation (Björnsson 2014). Because the subject of the obligation in question is a plurality of individuals each satisfying the conditions of moral agency, it is non-mysterious that disorganized groups that do not together constitute unified agents can nevertheless have obligations—or so I have argued. In sections 3 and 4, I begin by briefly recounting and motivating a general understanding of individual obligations before explaining how it naturally yields an account of shared obligations. In sections 5 and 6, I then explain how this understanding of obligations lets us make sense of the activist’s perspective, in spite of INDIVIDUAL INEFFICACY.

My ambition is to show that the activist’s perspective on our climate related obligations is coherent and sensible in spite of the apparent problems, not necessarily that it is correct. For this reason, I take a number of things for granted, most obviously the reality of anthropogenic climate change and the moral importance of the threat it poses. In sections 5 and 6, I rely on further normative assumptions, sometimes based on intuitions about cases. I provide no defense for these assumptions apart from some initial intuitive support, as one can agree that the perspective is coherent and sensible without ultimately accepting its normative content.

3. Obligations

Can collectives that do not themselves constitute moral agents have moral obligations? To answer this question, we are helped by a clear general understanding of what moral obligations are. Unfortunately for our purposes, there is no one sharp notion of moral obligation, and philosophers variously appeal to *objective* and *subjective* obligations, and to *all-things-considered*, *prima facie*, and *pro tanto* obligations. In what follows, I will therefore propose an explication of one sort of phenomenon that we have in mind when talking of obligations in moral contexts. The analysis of collective obligations that I propose based on this will then have to be adjusted to cover other kinds of moral obligations.

³ Another response to the AGENCY REQUIREMENT understands attributions of obligations to some non-agential relevant groups distributively, as telling members “to take steps to transform the group” into an agent that can perform the relevant action (Collins 2013, 233; cf. Lawford-Smith 2015). Apart from failing to account for collective obligations that do not necessarily involve the creation of a unified group agent, this proposal also postulates an ambiguity in attributions of obligations, treating some as tacitly attributing obligations with a different object than the explicit one. Even if ambiguities should not be antecedently ruled out, a more unified account would be preferable.

Abstractly, we can say that our obligations are what can be properly demanded of us, with moral and legal obligations being what can be properly morally and legally demanded of us, respectively. What can be morally demanded is a matter of what can be demanded on moral grounds, but demands come in different forms, where failures to satisfy the demand have different consequences. Trivially and generically, one must satisfy any moral ideal *on pain of being morally non-ideal*. Certain more specific moral demands must be satisfied *on pain of making one a fitting target of moral blame in the form of guilt (in the first-personal case) or indignation*. (To the extent that such demands permit actions that are morally non-ideal, they allow for supererogatory actions: actions that are morally better than what is morally required.) Other demands still must be satisfied *on pain of making one a fitting target of more hands-on sanctions or attempts to force us to comply than what is provided by basic expressions of negative reactive attitudes*. (To the extent that such demands allow for acts that merit blame, they allow for suberogatory actions: actions that are permitted but still blameworthy.)⁴ For convenience, my focus here will be on demands the violations of which ground blame.

In understanding obligations as demands the violations of which ground blame, we can distinguish *unconditional* from various kinds of *conditional* grounding relations:

UNCONDITIONAL GROUNDING: If X has an obligation to ϕ , X would be morally to blame for not ϕ -ing.

CONDITIONAL GROUNDING: If X has an obligation to ϕ , X would be morally to blame for not ϕ -ing *if ...* X could ϕ / X were, or should be, aware of the relevant moral considerations / X had sufficient control over ϕ -ing / there were no further countervailing moral considerations /

Some of these different kinds of relations correspond at least roughly to familiar distinctions between obligations that are subjective (relative to the agent's evidence about his circumstances) and objective (relative to the actual circumstances) as well as between all-things-considered and prima facie obligations: X has an [*objective* / *prima facie*] obligation to ϕ if and only if X would be to blame for not ϕ -ing if [X knew all relevant facts about the situation / X could ϕ and there were no stronger, countervailing, moral considerations].

Conditional grounding allows for failures to discharge obligations that are excused and so blameless; unconditional does not.⁵ Since excuses play no role in the discussion that follows, I

⁴ For the category of the suberogatory, see (Driver 1992). There are more possibilities: demands that must be satisfied on pain of (i) being a fitting target of moral protest, where this need not involve moral indignation (Smith 2013), (ii) not living up to a standard that should be actively encouraged, where encouragement is different from threats of blame or protest should one fall short, or (iii) not doing what an ideal moral advisor, or any morally reasonable person, would want one to do in the circumstances.

⁵ Michael Zimmerman (2008, 10) rejects an unconditional connection between obligations and blameworthiness based on the possibility of blameless wrongdoing, or failure to discharge obligations, where the agent's circumstances and cognitive or other shortcomings excuse. But while some notions of obligation let the two come apart, not all do: one can naturally deny that someone did morally wrong when she caused a risk for others exactly because she did so blamelessly, perhaps simply failing to combine information available to her.

will here simplify matters by focusing on a notion of moral obligation as an unconditional moral demand:⁶

OBLIGATION: X is morally required to ϕ if and only if X would be to blame for not ϕ -ing.

To further spell out what moral obligations of the relevant sort are, we will be helped by an account of when one is to blame for not ϕ -ing. Ideally, such an account should explain a number of general facts about moral obligations. One is that someone is required to do something only if doing it is in some relevant sense under her control (“ought implies can”). Another is that obligations take a variety of objects, at least as far as ordinary moral thinking goes. We can of course be to blame for actions, but also for omissions, (lack of) cognitive and emotional reactions, and states of (in)attention.

To account for these general facts, I follow an increasingly influential tradition of taking blameworthiness to be grounded in the agent’s “quality of will”:

QUALITY OF WILL: One is blameworthy insofar as the object of blame – what one is blamed for – is an “expression” of a bad quality of will (e.g. Strawson 1962; Arpaly 2006; McKenna 2012).

As this tradition has made clear, paradigmatic cases of blaming involve thinking that something bad happened because the agent did not care enough about morally relevant considerations. Conversely, we typically do not blame someone for an action if we find out that they did it because of external forces or because they were ignorant of crucial information rather than because of a lack of concern with or respect for the values involved.

Some precisifications and qualifications are in place, though.

First, there are various ways in which one might understand *quality of will*. I take it that blameworthiness is grounded in failures to *care* about certain things in the right way, where

CARING: To care about some object in some regard is to be disposed to notice factors relevant to how well it goes with it in that regard and to invest resources to promote it.

Quality of will accounts can accept this but differ in what objects one needs to care about for one’s will to be good: total wellbeing, justice, or one’s own happiness, relations, character, or avoidance of certain actions or intentions, to mention some possibilities.

Second, not any *non-ideal* caring is bad or lacking in a way that grounds blame. We are not blameworthy for actions falling short of ideal virtue. Rather, the caring must fall below some relevant standard, with normative force to back up demands that agents satisfy it and blame when their behavior expresses substandard caring. Moreover, standards do not apply uniformly to all agents or at all times. Indignation seems misplaced when directed at an agent

⁶ Those who find this unintuitive can substitute for OBLIGATION a version that conditionalizes blame on relevant kind of control and understanding. For a version that conditionalizes blame on (instrumentally) rational control, see n. 14.

who lacks the capacities required for caring about certain things—certain core elements of social and moral cognition, say—and to the extent that such capacities come in degrees, the relevant standards might vary accordingly. Furthermore, the degree of caring demanded by the standards will vary depending on the object of care, and might vary depending on one's relation to that object (for an overview of relevant relationships, see Björnsson and Brülde 2017).

Third, not any event expressive of substandard caring is a relevant object of blame. The object of blame needs to *itself* be morally bad, i.e. of a kind that we have moral reason to prevent. For example, an agent might not be to blame for an expression of substandard caring about his father during free association therapy. As with objects of caring, a quality of will account can allow for a variety of objects of blame, including not only clear cases of harm, rights violations, or suboptimal outcomes, but also failures to notice or be motivated by some morally relevant fact and attempts to do cause harm.

Finally, the object of blame needs to be explained *in a normal way* by the substandard caring, not through some random chain of events (Björnsson 2014, 114–15; Björnsson and Persson 2012, 330–31). The normal way in which substandard caring explains a bad event (action or outcome) is by explaining agents' failures to recognize or be motivated by the fact that their behavior would either contribute to this bad event or allow for other, preventable, events that do so. In the simplest case, the bad event is explained by a single instance of substandard caring. But many bad events, such as the deterioration of trust among people or failures to complete a task at a promised time, are upshots of a several such instances, sometimes on part of several people.

The resulting version of the quality of will account is that

BLAMEWORTHINESS: X is morally to blame for Y if and only if Y is morally bad and explained in a normal way by X's substandard caring (Björnsson 2011; Björnsson and Persson 2012; Björnsson 2017a, b).

As I have argued elsewhere, this proposal (together with a companion characterization of degrees of blameworthiness) accounts for everyday attributions of blameworthiness as well as the role of ordinary excuses and justifications in mitigating blameworthiness to various degrees (Björnsson 2017a, b).

When we conjoin BLAMEWORTHINESS with OBLIGATION, it follows that X is required to ϕ if and only if not ϕ -ing is bad and X's not ϕ -ing would be explained, in a normal way, by X's caring falling short of applicable standards. When is X's not ϕ -ing guaranteed to be so explained? Exactly when the satisfaction of these standards would ensure, in a normal way, that X ϕ . Hence:

OBLIGATIONS AS ENSURED BY CARING (OEC): X is morally required to ϕ if and only if X's not ϕ -ing would be morally bad and X's ϕ -ing would be ensured, in a normal way, by X's caring appropriately (Björnsson 2014; Björnsson and Brülde 2017).

Apart from capturing the connection between obligation (of the relevant sort) and blame, OEC straightforwardly accounts for the fact that one has an obligation to ϕ only if ϕ -ing is under one's control, or something one can do, as one must be in a position where caring appropriately ensures, in a normal way, that one ϕ . OEC also leaves room for a wide variety of objects of obligations. That one cares appropriately can ensure that one acts or refrains from certain actions, but also that one notices or ignores certain things, or has (or lacks) certain cognitive or emotional reactions. Moreover, what is ensured might be a series of coordinated actions guided by an explicit intention, such as the act of *walking over to the stereo and turning down the volume*, as well as a series of separate and individually motivated actions and omissions that has an overall effect, such as the act of *letting the neighbor sleep through the night* by avoiding activities that make much noise throughout the night. It is also worth noting that OEC allows for a variety of substantive views about the objects of appropriate caring, including both impersonal values (total happiness, justice, beauty) and agent-relative matters (not lying, not stealing, not killing, being a good human being, citizen, friend, parent, spouse, colleague).⁷

4. Essentially shared obligations

Return now to the difficulty of understanding obligations of groups that are not themselves moral agents. At a first glance, OEC might seem to be of little help, as it requires that obligation bearers are subject to moral standards of caring, and as it is natural to think that such standards only apply to moral agents. Indeed, we might want to identify moral agents with those who are subject to moral standards of caring. On closer inspection, however, we can see how OEC can be satisfied for some plural subjects:

First, and trivially, even if X is plural, X 's not ϕ -ing might be morally bad, i.e. the sort of thing that there are moral reasons to prevent: there might be moral reason to prevent a group of teenage yahoos from waking up your sweet elderly neighbor, or from abandoning a cat helplessly stuck in a tree after they chased it up there.

Second, groups of agents can satisfy or fail to satisfy standards of caring. The group of teenagers, for example, might care as can be appropriately demanded of them about the neighbor, or the cat. Of course, we are assuming that only fully fledged moral agents can satisfy or breach moral standards of caring and our concern here is with groups that are not themselves such agents. Consequently, that a group satisfies standards of caring must be understood distributively: it is simply a matter of each individual member satisfying standards applicable to her. But notice that the question of whether some individual cares appropriately about something is a correspondingly distributive matter. Since, plausibly, only an agent *at a time* can satisfy or breach the relevant standards, the thought that an agent cares appropriately

⁷ Notably, OEC is structurally similar to virtue ethical identifications of right actions with actions performed by the virtuous (see e.g. Hursthouse 1999, ch. 2), though it is concern with duty rather than rightness, and appeals to appropriate caring rather than (presumably more demanding) agential virtue. For a brief discussion of how OEC might handle some problems raised for such accounts, see Björnsson 2014, 116–17, n13.

over some stretch of time will have to be understood distributively, as the thought that she cares appropriately at each time during that period.⁸

Third, appropriate caring by a group of agents can ensure, in normal ways, a certain behavior from the group. If the teenagers care appropriately about the neighbor, this will ensure that each keeps calm in the proximity of her bedroom and that they thereby keep the noise they collectively make below a certain level. If they care appropriately about the cat, this might ensure that they organize to help it down. Notice that whereas coordination and mutual awareness might be essential in the second case, the first case is different. Here, as in countless other cases, appropriate caring avoids cumulative problems by separate efforts of a large number of individuals, efforts that need not be made in awareness of the other efforts. Again, this is directly parallel to a typical individual case. It is not the teenager's caring about the neighbor at one specific time that ensures that the neighbor gets to sleep, but her caring throughout the night and making efforts to keep the noise down at various times, without necessarily being aware at any of those times of her other efforts.⁹

In light of these considerations, it seems unproblematic and non-mysterious that the teenagers might have a moral obligation to help the cat down, or not keep the neighbor awake. It does not matter whether the teenagers together constitute a fully-fledged moral agent. Although the demands that ultimately ground obligations—demands on caring—always apply to moral agents at times, collective obligation of this sort are functions of such demands in the same way as individual obligations are. Moreover, these collective obligations imply capability in just the same way as individual obligations: if a group has an obligation to ϕ , the group's ϕ -ing would be ensured, in normal ways, by appropriate caring.¹⁰

⁸ I suspect that the last claim might need to be complicated, as standards of caring for individuals might be less than fully distributive. However, as my grounds for this suspicion apply equally to the group case, an individual's satisfaction of standards for caring in temporally extended cases and a group's satisfaction of standards would still be structurally analogous. I thus put this complication to the side.

⁹ One might worry in general whether the possibility of bringing about some important effect through a series of uncoordinated actions really provides the sort of control needed to ground obligation (Lawford-Smith 2015, 233–35; cf. Pinkert 2014; Schwenkenbecher 2014, 68–69). But notice that the effects by uncoordinated actions need not come about by chance: they are normal upshots of separate behaviors all guided by the value realized or promoted by the effect.

¹⁰ Though unstructured groups can be subjects of the OEC variety of moral obligations, only moral agents can be subject to a closely related variety:

OBLIGATIONS AS OBJECTS OF CARING (OOC): X has an obligation to ϕ if and only if it can be properly morally demanded that X cares about ϕ -ing.

Thoughts about basic general obligations—obligations to keep promises, treat people with respect, not lie, not steal, not kill, etc.—and corresponding *pro tanto* obligations concerning specific actions or omissions are often well understood as obligations of this latter sort (or as *responsibilities* to care about (not) performing certain actions, in the terminology of Björnsson and Brülde 2017). Some resistance to the idea that unstructured groups can have obligations (e.g. that of Collins 2013) might be best

For an obligation like this to obtain is for a number of agents to stand in a certain relation to ϕ -ing *together*, not for one unified agent constituted by such a plurality to relate to ϕ -ing. I will thus call it a “shared” obligation.

Notice that such obligations are *essentially* shared. To see this, compare two cases where an individual teenager cannot help the cat down the tree because at least two people are required to erect the ladder and because there is no help available. In the first case there is no help available because the teenager is there alone; in the second case because everyone else present is unwilling to help. In neither case does she have a regular individual obligation to help the cat—she cannot—but in the second, she might share an obligation with the others to do so. The latter possibility also illustrates how shared obligations do not consist in regular individual obligations to achieve some end together with others: given the unwillingness of others, she has no individual obligation to help. Nonetheless, this account of collective obligations satisfies something close to the AGENCY REQUIREMENT, as shared obligations derive from demands that bona fide moral agents care in certain ways about various things.

Thus far I have argued that on an independently supported understanding of obligations, collective obligations of a certain kind—shared obligations—are entirely non-mysterious. But three questions remain to be answered, neither of which has an obvious answer:

Given that shared obligations depend on the motivation of appropriately caring individuals, how can such obligations emerge in spite of INDIVIDUAL INEFFICACY?

Given that shared obligations depend on the motivation of appropriately caring individuals, how can COLLECTIVE EFFICACY ground a shared obligation?

Given INDIVIDUAL INEFFICACY, how can the considerations grounding COLLECTIVE OBLIGATION also ground individual obligations to contribute?

I address these questions in the following two sections, focusing on the issue of threatening climate catastrophe.

5. COLLECTIVE OBLIGATION in spite of INDIVIDUAL INEFFICACY

Our account of shared obligations might seem to make straightforward sense of COLLECTIVE OBLIGATION for the case of threatening climate catastrophe. First, according to OEC, we have a moral obligation to significantly reduce that threat insofar as our caring appropriately would ensure, in a normal way, that we do. Second, it seems plausible that we can be properly required to care about things threatened by climate change, in particular the wellbeing of people in areas most likely to be hit the hardest by flooding, drought, social unrest, disease, and war. Of course, our production and use of energy will have to change radically in a generation, imposing significant costs on many people. But key technologies have developed swiftly even without costly measures, and more forceful policy changes, should such receive

understood as concerned with something like OOC rather than OEC obligations. I thank Stephanie Collins for discussion of this issue.

widespread support, would surely speed up the process. Given the magnitude of what is at stake, surely the required sacrifices can be demanded.¹¹

On reflection, though, INDIVIDUAL INEFFECTIVITY might seem to undermine this conclusion. Given how small an impact we can expect of an individual agent's actions on overall preventive success—given INDIVIDUAL INEFFECTIVITY—it is unclear why climate threat should prompt even moderately costly action among appropriately caring individuals. That question is the topic of this section.

An obvious first answer is that although we are extremely unlikely to make any major difference with respect to the grand outcome that is the object of COLLECTIVE OBLIGATION, the difference we can make is not thereby morally irrelevant. We might cut down on our carbon footprint, support emission reduction schemes to offset what remains, vote for politicians that support important climate legislation, support climate friendly technologies by our everyday choices, and help spread awareness of the risks and what needs to be done. Taking these steps, one might think, would likely make *some* contribution in positive direction. After all, whatever threat reduction is eventually achieved will in significant part be achieved by a multitude of such steps. Given appropriate caring, this could motivate enough action to ensure significant overall reduction.

I will argue that this is a sensible perspective, but it too might seem to exaggerate the role we can play as individuals. Consider the fact that many of the systems that we operate in have what we might call “buffer zones,” i.e. ranges within which differences in input values make no difference to the overall outcome. Voting is the most familiar example: unless an election is as close as can be, no one individual vote will make a difference as to who wins. The likelihood that my vote will make a difference in an election with millions of voters is thus typically vanishingly small, especially where reliable advance polling indicates a clear winner. But numerous natural and human systems display similar buffering, with corresponding effects on likelihoods: it might seem extremely unlikely that my emissions of greenhouse gases will cause a vanishingly small increase in temperature or sea level rise that in turn makes a difference whether an area is hit by drought or forest fires, whether someone dies from dehydration or overheating, or whether flood barriers resist the forces of waves. Moreover, it might seem that unless one's actions make some concrete identifiable difference with respect to some value—unless it harms an identifiable human being, say—, that value cannot make performing or abstaining from these actions morally obligatory (Sinnott-Armstrong 2005, 289–94).

In response to this sort of difficulty, it might be thought that focus on the effects or difference-making of individual actions is a mistake, as well as a bad fit with the activist's

¹¹ Costs imposed need to be measured against the costs of not taking action. In the case of climate catastrophe, the latter could be significant for those who are young today and their children. Exact estimates of the former from IPCC and others vary considerably depending on what assumptions are made, but estimates for quick and decisive global action suggest individual burdens falling far short of the sorts of costs countless people pay every year because of individual health problems or changes in labor market conditions that force them into periods of joblessness or require relocation. For one recent argument that costs are reasonable, see (Stern 2015).

perspective. Perhaps we should also care about not adding to sets of actions that together cause significant identifiable harm, and care about adding to sets of actions that together delivers significant identifiable goods. To capture this wider sense of causal relevance, let us say that an individual action (where omissions are understood as actions) “contributes” to some outcome if it is part of why the outcome came about. An action can thus contribute to an outcome without making a difference to it, as illustrated by cases where the outcome is the joint effect of numerous actions neither of which was necessary for the outcome. With this terminology, the suggestion might be put as follows:

MERE CONTRIBUTION: If one cares appropriately about some important outcome and thinks that actions by a group of people can significantly promote (obstruct) that outcome, one cares about contributing to its promotion (obstruction), even if one will not thereby make a significant difference to whether it is realized.¹²

Given MERE CONTRIBUTION, it is clear why individuals would care about not adding to the excess of greenhouse gases and about supporting pro-environmental policies. Depending on the strength of the required caring, this could conceivably ensure a significant reduction of climate threat.

Some would accept a principle like MERE CONTRIBUTION only for the special case where the contributor intends a collective harm as a means or an end (e.g. Sinnott-Armstrong 2005, 289–90). Perhaps they should generalize: as Chad Vance (forthcoming) has argued, a harm that is merely foreseen can be as serious as one that is intended (though Vance allows that the former might be more easily offset by some greater good). For cases involving collective benefits rather than harms, Julia Nefsky (2016; cf. Parfit 1984, 83) has recently defended a principle related to MERE CONTRIBUTION. However, she notes that without restriction such principles imply that a valuable outcome can provide an agent with reasons to contribute even when it is known that enough has already been done to secure that outcome. Finding this consequence implausible, she proposes that the principle applies only to contributions that satisfy

NON-SUPERFLUITY: An action is *non-superfluous* with respect to an outcome if and only if it is possible at the time of action that it will contribute to the outcome and also possible that the outcome fails to materialize for lack of such contributions (Nefsky 2016, §4).

Perhaps a similar restriction makes sense for cases of foreseen collective harm. If the actions that others have taken have already ensured the harm, it is unclear why concern with that harm should motivate an agent not to perform another action of the relevant kind. (The agent might of course have reason to avoid such actions under certain circumstances—perhaps doing so will have powerful symbolic effects, say—but then it is not the *mere contribution* to harm that provides such reason.)

¹² For considerations that might seem to motivate MERE CONTRIBUTION, see Brennan and Sayre-McCord 2015, 53–54; Petersson 2004; Parfit 1984, ch. 3.

On further reflection, however, even non-superfluous mere contributions seem to count for nil. Consider:

Unhelpful Contribution: I see two people trying to save a drowning man from the violent waves onto the rocky shore. I can contribute to the effort, and so can another bystander, whom I recognize as an off-duty lifeguard. He seems unwilling to get involved, but if he did, the man would most likely be saved. I am less suited for the task, however, having little experience with this sort of situation. If I got involved, my contribution would ever so slightly decrease the chance of success or at the very least complicate the situation for the others.

In this case, my attempts to aid the rescue could very well be non-superfluous parts of the cause of a successful outcome: if the man was brought ashore alive it could be because we did it. Even so, it is unclear why an appropriate concern with the man's life and the success of the rescue should motivate any such attempts. In fact, no matter how marginal the negative effect would be of getting involved, it would seem enough to counteract whatever positive value would be added by my non-superfluous contribution. The most straightforward explanation for this is that such contributions add no value of their own (cf. Parfit 1984, 83; Dietz 2016, 974–75). Likewise for cases involving mere contributions to harm:

Unhelpful Abstinence: Watching people climb onto the lifeboats, I can see that one is filled over capacity, enough so that this might sink it if the waves get any higher. By climbing onboard myself, I would weigh it down further still. But unlike the others, I know how to balance and steer a lifeboat to best keep it afloat. On balance, if I climb on, it would be slightly less likely to sink.

If I get onboard, the lifeboat might sink because too many people climbed onboard, in which case I would have non-superfluously contributed to the wreck. However, if this contribution has any negative value in itself, it seems outweighed by even the slightest increase in risk that would come from not contributing. This suggests that contributing non-superfluously to bad outcomes has no inherent negative value.

I conclude that the basic thought behind MERE CONTRIBUTION fails to explain why people caring appropriately would on aggregate take enough action to significantly reduce climate threat.

A better response to the worry about buffer zones relies on the fact that even highly unlikely outcomes matter if they are significant enough (see e.g. Parfit 1984, 73–75; Hiller 2011; Kagan 2011; Singer 1980; Lawford-Smith 2016). Our decisions should arguably be guided by the effect of available alternatives on something like *expected value* (i.e. on the sum of the products of the value of each of the possible outcomes and their probability). Assuming that this is correct, suppose that the minute difference in greenhouse gases that I might affect has some non-zero likelihood of affecting the occurrence or extent of forest fires, droughts, or whether someone dies of dehydration or overheating, and so forth. Even if each such effect is extremely unlikely, there are plausibly a staggering number of possibilities for such effects,

each individually of considerable importance. Moreover, unlike in cases of elections with reliable forecasts, it seems that individuals have no way of knowing whether they find themselves in buffer zones with respect to the vast majority of these possible effects. Given this, the difference in expected value that the individual can make by affecting his own or others' greenhouse gas emissions might still be significant. Of course, in light of the vast complexities, any calculation of expected value will have to rely on reasonable heuristics. A first approximation might be to look at something like the projected average marginal contribution of a given amount of pollution in the range of sufficiently likely outcomes. We might then find that the expected negative climate related value of an afternoon worth of gas guzzling joy riding is (very roughly) equivalent to that of ruining someone else's afternoon (cf. Hiller 2011, 355–61). That should be reason enough to consider carbon offsetting or alternative sources of enjoyment.

Still, one might think that because of the complexity and magnitude of the systems at work, no concrete harmful or beneficial effects can be traced back to an individual's contribution such that this particular contribution can be said to be causally responsible for those effects (Maltais 2013; Sinnott-Armstrong 2005).

The problem here cannot be that we lack reason to think that an individual's marginal contribution could make certain massive determinate differences (*pace* Maltais 2013, 592–94), although in unknowable ways. After all, ordinary physical mechanisms are often sensitive to subtle changes, translating minute differences in initial conditions to massive differences downstream in ways that make many natural and social events virtually unpredictable.

But appeals to causal responsibility for harms encounter two other problems. The first is that the vast complexity of the system in question, where the pre-history of events will involve an abundance of causal factors, make it unclear on what grounds a given event should be traced back to a particular individual's actions rather than to some other factor (see Sinnott-Armstrong 2005, 289–91). The second problem is that the consequences of our actions might be metaphysically indeterminate rather than merely in principle individually unknowable. One possibility would be that the laws of nature are indeterministic. Another is that the alternatives available to an agent themselves might involve significant metaphysically indeterminacy. Consider the sort of counterfactuals that would pin a particular outcome to one individual's contribution: *if you offset your greenhouse gas emissions, the outcome (e.g. the flooding of someone's cellar) is avoided; if you don't, it isn't*. Arguably, the antecedents of such conditionals typically fall short of identifying an exact and comprehensive physical state of the world. For example, there are many concrete ways in which you could make true that you offset your greenhouse gas emissions, differing in various ways at both macro and micro physical levels, and it is unclear what would make it determinately true of one of these ways that if you did offsetting, you would do it in *that* way. But without such determinacy, any consequences flowing from chaotic systems—systems whose outputs are highly sensitive to minute input differences—will be correspondingly indeterminate, even assuming that the universe operates on deterministic principles.

If only the likelihood of concrete, identifiable harmful or beneficial consequences of action were relevant for decision-making, indeterminacy of consequences could thus

undermine duties to reduce one's own carbon footprint or to otherwise contribute to prevention. On reflection, however, such consequences are just not required for decision-making relevance, as illustrated by:

Carcinogen: It is in your power to determine whether the population in a town of 35,000 people will be exposed to a carcinogen. Among the people who have the gene in virtue of which they risk developing and dying from a certain form of lung cancer, exposure increases the average risk of contraction from 50% to 52%, depending in each individual case on numerous factors and indeterministic processes. As the expected number of townspeople with the gene is 200, exposure increases the expected number of lung cancer deaths in the town by 4.

In *Carcinogen*, it is not clear that exposing people to the carcinogen would in concrete ways harm anyone identifiable. For if you did, we cannot say of any of the people who died from lung cancer that they *would* have been better off if you had decided against exposure, only that they might have been. Similarly, if you didn't, we could not say of anyone not developing lung cancer that they would have been worse off if you had allowed exposure, only that they might have been. Nevertheless, the difference in expected value between the alternatives seems to provide an obligation to protect people from exposure absent very strong contrary reasons. In terms of OEC: Acting in ways that increase the number of expected deaths is morally bad, and if you cared appropriately about the values and risks at hand, this would ensure, in normal ways, that you decide against exposure. Consequently, you have an obligation not to expose people to the carcinogen. Moreover, although this obligation is grounded in the harmfulness of the carcinogen, it does not depend on any requirement to avoid identifiable harm—unless increasing the risk of harm itself constitutes a harm.¹³

¹³ Whether increasing the risk that someone is harmed or lowering someone's expected wellbeing itself constitutes *harm*, or perhaps more weakly *injustice*, might matter for the strength of the resulting individual obligations. If avoiding carbon emissions is merely a matter of improving the state of the world, one might think that there are more effective allocations of the resources one would spend: malaria prevention, say. But if it is instead a matter of avoiding the violation of negative duties not to harm or commit injustice, it might trump positive duties to do good (see Broome 2012, ch. 4–5). I do think that when an action increases expected total harm without compensation or justification, it raises to the level of injustice. But for the issue that we are currently considering—what individuals would do on the assumption that *all* cared appropriately—this might not matter much: the marginal expected value of support for what currently yields most good per dollar would likely quickly go down with enough contributions, leaving enough resources to cut net carbon emissions.

Even if increased risk for harm or lowered expected value for individuals can be relevant for action, one might think that such relevance requires *substantially* increased risk or lowered expected value for those individuals, even if involving a vast number of individuals. (Cf. Julia Nefsky's [2011] argument that minor differences in the distribution of goods might be taken not to make a difference as to whether the distribution is fair or not.) Contrary to this, I believe that variations on cases like *Carcinogen* help bring out that small differences in the dimensions that ground large relevant differences are themselves morally relevant, albeit typically negligible unless affecting a large number of people. Cf. discussion of EFFECTIVE CONTRIBUTION in section 6.

Suppose, as *Carcinogen* suggests, that we can be required to care about the expected extent of harm. Suppose further that, as previously suggested, individual contributions to climate catastrophe prevention can significantly affect expected climate related harm. It follows that we can have individual obligations to contribute even given INDIVIDUAL INEFFECTIVITY. Moreover, it suggests that we—affluent people living today, say—share an obligation to significantly reduce the threat of climate catastrophe. If all of us care as can be reasonably required of us, we will care about our marginal contribution to such threat reduction, and such caring would likely ensure significant threat reduction.

Of course, I have not yet said anything about how strongly individuals should care about such effects, or how much action could be expected. The caring required might depend on individuals' capacity to contribute, their degree of (retrospective) responsibility for the situation, and the extent to which they have disproportionately benefitted from the system that has given rise to the threat (for a catalogue of possible factors, see Björnsson and Brülde 2017). How much action will be ensured given a certain level of caring will in turn depend on the costs involved and competing considerations. Though it is not part of the current argument, it is entirely possible that for those of us who belong in the global 1%, caring appropriately will dispose us to considerably lower our material standards of living. But setting this possibility aside, it is easy to see how uniform moderate caring could ensure significant effects. Little or no sacrifice of personal wellbeing is involved in staying somewhat informed or helping raise awareness and spread relevant information, on one's own or together with others. Similarly, little sacrifice is involved in ranking climate policy as particularly important when casting votes in close elections, offsetting one's carbon emissions, or choosing the climate friendly alternative when options are otherwise roughly equivalent. For most affluent people, performing such actions out of caring about threat reduction would impose limited material hardships and might well involve positive net effects on happiness characteristic of involvement in altruistic causes. Given this, it is not wildly implausible that if everyone cared appropriately, this would ensure that a vast number of individual actions and widespread support from pro-environmental policies would greatly speed up the world's transition to carbon neutrality, thus significantly reducing climate threat. And from this, together with OEC, it would follow that we have a shared obligation to significantly reduce climate threat.¹⁴

¹⁴ I've suggested that people who care appropriately would take individual steps that would be jointly sufficient for significant threat reduction (typically through various forms of collective actions). But this might be resisted on the ground that even people who care about the threatening effects of climate change might fail to draw the conclusion that individuals have considerable reason to contribute. After all, as we have seen, a number of philosophers reject that conclusion. If sufficiently many draw the same conclusion, appropriate caring might not ensure enough action to significantly prevent catastrophic climate change, thus seemingly undermining the idea that prevention is a shared obligation. One response to this worry is to say that the relevant standards for caring require caring about the expected value of one's actions, independent of identifiable concrete harm or benefits resulting from these actions. On this view, those failing to contribute because they take their inaction to cause no identifiable harm and not caring about expected value would be morally to blame. Another response accepts that the worry might indeed undermine shared obligations of the OEC

6. Remaining problems of COLLECTIVE EFFICACY and INDIVIDUAL INEFFICACY

Thus far we have seen how the activist's attributions of collective obligations can be understood as concerned with an OEC variety of *shared* obligations. We have also seen reasons to think that we have such a shared obligation in the climate case: if we all cared appropriately, this could well ensure, in normal ways, that the climate threat is significantly resolved. But questions remain.

First, to make full sense of the activist's perspective, we need to explain how the shared obligation is relevantly grounded in COLLECTIVE EFFICACY, as postulated by COLLECTIVE OBLIGATION. Importantly, OEC itself guarantees that we *can* discharge any shared obligation, in the sense that we would discharge it if we cared as can be properly demanded of us. But as I understand the activist's perspective, the grounding claim involves something more: we share an obligation to resolve the issue because we are (or should be) aware of our collective capacity to resolve it. By contrast, the considerations of efficacy mentioned in the previous section concerned only the individual's capacity to marginally affect expected climate value. More thus needs to be said.

Second, according to RESULTING INDIVIDUAL OBLIGATION, considerations that ground our shared climate obligation also ground individual obligations to help discharge it. The problem here concerns the activist's understanding of what it is to *help* resolve an issue. Activists don't generally think that a given individual's actions can affect whether the issue is resolved, or think that this is what grounds the individual's obligations: such a resolution is too vague and too grand a matter, especially if an individual's actions might have no determinate climate related upshots.¹⁵ Moreover, as illustrated by *Unhelpful Contribution* and *Unhelpful Abstention*, the importance of some large-scale outcome does not in any direct way ground reasons to make mere contributions to it (if it is good) or abstain (if it is bad). What is needed is some other understanding of how individual actions can more or less effectively help bring about some important outcome, without relying on the assumption that they increase its likelihood to various degrees.

In this final section, I explain how concerns about important large-scale outcomes and considerations of collective efficacy with respect to such outcomes can intelligibly enter into ordinary individual practical reasoning and make intelligible individual efforts to efficiently help realize these outcomes. Such concerns and considerations, I suggest, play natural roles in what might be the most common mode of practical reasoning: w I will call *issue-oriented*

variety and instead say that it is moral demands on caring together with *teleological rationality* that require giving weight to expected value. Based on this, one could then say we have a shared obligation under the following, OEC-related, notion of obligation:

OEC*: X is morally required to ϕ if and only if X's not ϕ -ing would be morally bad and X's ϕ -ing would be ensured, in a normal way, by X's caring appropriately *and being teleologically rational*.

¹⁵ Given that agents can marginally affect climate related expected value, it might be hard to deny that they can marginally affect the likelihood of resolution given various complete precisifications of what counts as a "resolution." But I do not take that thought to be a natural part of the activist's perspective.

reasoning. Simplifying somewhat, such reasoning paradigmatically proceeds through the following steps:

ISSUE-ORIENTED REASONING:

Identification of an issue, i.e. some good or bad possibility whose realization might be helped or hindered.

Identification of one or more possible means by which that realization might be helped or hindered. A means may be non-agential: the threatening forest fire might be put out by rain; the tyrant's reign ended by cardiac arrest. But it will typically involve an agent, who may or may not be the agent engaged in reasoning, or a group of agents, which may or may not include the reasoner.¹⁶ Means might also be more or less abstractly understood, ranging from the broadest outline of what needs to happen to help or hinder the possibility in question to plans painstakingly detailing what would be done by everyone and everything involved.

Evaluation of means with respect to effectiveness, likelihood of realization (perhaps conditional upon the active involvement of the reasoner) and various cost and benefits. Costs saliently include resources used, and both costs and benefits can include moral considerations, such as unfairness or fairness in the distribution of costs or in how people are used, and whether the actions involved go beyond or is in accordance with what can be reasonably demanded of people. In the case of climate threat reduction, the fairness or justice of means might depend on assessments of the individual agents' capacity help effectively as well her prior reliance on more than her fair share of Earth's limited capacity to harmlessly absorb greenhouse gas emissions (see e.g. Baatz 2014; cf. the extensive discussion of "polluter pays" and "beneficiary pays" principles in environmental economics, law, and ethics). Evaluations might lead to rejection or refinement of means, and the identification of the best possible or only acceptable means. If moral and other costs of available means do not stand in proportion to the target outcome—to keep the risk of the direst catastrophe scenarios under 1%, say—or if the likelihood of success is insignificant, the target might be raised or lowered and the costs recalculated.

Identification of way that the individual can further a certain means, by making it more likely or improving on it. Such ways can be more or less abstractly understood.

Evaluation of ways of promoting means, with respect to the same range of considerations as evaluation of means; can lead to rejection or refinements of certain ways. Can also lead to reevaluation of means, as prior evaluations might have relied on now rejected assumption about reasoner's input.

Individual decision about what to do about issue, if anything.

¹⁶ As should be clear from this outline, the sort of reasoning I have in mind here is neither "we-mode" nor "I-mode" reasoning of the sort outlined by (Tuomela (2007)), though it might conceivably involve either.

We can easily see how both considerations of collective efficacy and judgments of shared obligations have a natural home in such reasoning. When appropriately caring individuals become aware of the sort of issue that the activist's perspective concerns, they will tend to engage in issue-oriented reasoning. Because of the scale of the issue, potentially effective means characteristically involve great many people, making considerations of collective efficacy highly relevant. At this stage, it might also matter whether reasonable demands on the agents involved would ensure realization of the means—whether the realization is a shared obligation—as this might affect whether the means can likely be realized in a morally acceptable fashion.

To see how individuals who have taken this perspective can also rightly see their own actions as more or less effective and helpful contributions to such collective means, we need two more general ideas. The first, which has already been alluded to, is that:

HELPFUL CONTRIBUTION: A contribution to an important outcome can be helpful both by increasing the likelihood of the outcome and by improving the means by which the outcome is reached.

Correspondingly, appropriately caring agents will be motivated to contribute not only in ways that make a valuable outcome more likely, but also in ways that make reaching it easier, faster, cheaper, safer, involving a fairer distribution of efforts, and so forth. In line with this, we have already noted that an individual agent can make a difference to how fairly the threat of climate catastrophe is resolved, by not using more than her fair share of a vital common resource and by not unjustifiably reducing the expected wellbeing of others.

How, though, can the individual contribution be rationally directed at and tailored to efficiently help achieve a certain end when the contribution cannot significantly affect the likelihood that it is achieved? The answer is provided by something like the following plausible principle:

EFFECTIVE CONTRIBUTION: An action or omission is an effective contribution to an outcome to the degree that it increases the extent to which the outcome is realized.

Some outcomes are perhaps naturally understood as being either fully realized or not at all: preventing a fascist politician from getting a cabinet position, say, or becoming pregnant. For such outcomes, a contribution is either fully effective, or not at all. Other outcomes, however, are naturally understood as involving the realization of *enough* of some underlying dimension that can be realized to degrees. Some such outcomes are probabilistic: your goal might be to make it unlikely that the fascist gets a cabinet position or that you or your partner becomes pregnant. If so, a contribution is effective to the extent that it affects the likelihood in question. Others are non-probabilistic. You might strive to be well-educated, fair-minded, or wealthy, and contributions to those ends are more effective the more they add to your degree of education, fair-mindedness, or wealth. Or you might be working for a decisive election win, or a cancer treatment that prevents or reverses tumor growth, in which case a contribution might be more effective the more votes it adds to the victory or the more it prevents or reverses

growth, respectively. Given this understanding of effectiveness, the *efficiency* of a contribution can simply be understood in terms of the costs involved in achieving a certain degree of effectiveness.

Importantly, this understanding of degrees of effectiveness and efficiency applies even when, because of the vagueness of the end, no individual action will decide whether the outcome is realized: whether you will now be well educated, or whether the victory was now decisive. For this reason, EFFECTIVE CONTRIBUTION also makes clear how contributions towards significantly reducing climate threat can be more or less effective and efficient depending on how they affect expected climate related value.

We have now seen how considerations of collective efficacy can ground both a shared obligation to significantly reduce climate threat and individual obligations to contribute to such reduction, in spite of INDIVIDUAL INEFFECTIVITY. Appropriately caring individuals who become aware of an issue like this will consider means of resolution that involve groups of people who can resolve the issue. Such individuals will also be motivated to perform actions that are helpful—efficient and fair—contributions to threat resolution even if they do not see their own contribution as increasing its likelihood.

What is missing for a complete account of RESULTING INDIVIDUAL OBLIGATION are reasons to think that if an individual in the relevant group (the affluent, say) would care appropriately, she would be significantly motivated to make helpful contributions, grounding an individual *prima facie* obligation to do so. In section 5, I argued that if we *all* cared appropriately, enough of us would perform actions increasing climate related expected value for these actions to significantly resolve the threat of climate catastrophe. It seems painfully obvious, however, that many people care very little and are unwilling to contribute, and one might worry that this undermines what the appropriately caring individual can be expected to do. The worry concerns both efficiency and fairness. If others are not doing their part, it is both unclear whether doing my part will be a comparatively efficient way of contributing to a resolution and unclear how things are made fairer if I incur the cost of doing my part.¹⁷

On reflection, though, matters are not quite that bleak. Not only is there still a good chance that enough individuals will each act in ways that marginally increase climate related expected value and together significantly reduce climate threat, judging by progress over the last few years. Much has also been done to provide cognitive and material infrastructures that make helpful contributions both convenient and inexpensive, at least for the affluent among us. There are increasingly efficient renewable energy sources, improved low-emission technologies, and detailed sophisticated scientific assessments of both risks and the effects of various policies. There are also policies and services in place that make it easy for individuals to choose climate friendlier options, vote for and otherwise support climate-friendly politicians and policies, or inexpensively offset their carbon footprint. Efficiency should thus not in general be a worry. Considerations relevant to fairness also do not seem significantly affected by current reluctance to contribute. That others are using or benefitting from more than their

¹⁷ For general concerns about obligations under conditions of futility, see Budolfson unpublished manuscript; cf. Sinnott-Armstrong 2005. The argument of section 5 rejected the idea that individual contributions to climate threat resolution are futile given that others care appropriately.

fair share of a common resource without trying to compensate others does not obviously mean that I should not care about not going beyond my fair share. It thus makes good sense to assume that, in spite of fairly widespread tendencies not to take the threat of climate catastrophe seriously, and in spite of INDIVIDUAL INEFFECTICACY, we do have individual *prima facie* obligations to help achieve resolution, obligations grounded in the same facts as our shared obligation. In light of this, RESULTING INDIVIDUAL OBLIGATION seems perfectly reasonable.

This concludes my attempt to make sense of central aspects of the activist's perspective, with a specific focus on our climate obligations. In sections 3 and 4, I explained why there need not be anything inherently mysterious about shared obligations. Such obligations have the same basic structure as individual obligations and do not require that the group in question is itself a *bona fide* moral agent. In section 5, I argued that we likely have a shared obligation to significantly resolve the threat of climate catastrophe, as many of us can significantly affect climate related expected value. In this section, finally, I have explained how this shared obligation would be based in part on the recognition that we are collectively able to significantly resolve the issue, and how individuals can have obligations to help achieve climate threat resolution even though they cannot affect its likelihood. Both considerations of collective efficacy and attempts to help achieve a collective solution to a large-scale problem have natural homes in ordinary issue-oriented practical reasoning. If the arguments are correct, the activist's perspective thus remains reasonable in spite of the size and complex nature of the threat at hand and the largely disorganized group of agents to which it attributes both shared and individual obligations.

Acknowledgments

This chapter has benefited greatly from comments from audiences at seminars at Umeå University and the University of Gothenburg, the Metaphysics and Collectivity workshop at Lund University, the Trondheim ISP-FIDE PhD forum, the Philosophy and Climate Change workshop at Princeton University, Collective Intentionality X, and MANCEPT 2016. I'm particularly grateful for written comments from Ragnar Francén, Brian Epstein, Anne Schwenkenbecher, Stephanie Collins, Olof Leffler, Stefano Cossara, Daniel Telech, Holly Lawford-Smith, Erik Brandstedt, Mark Budolfsson and Tristram McPherson.

Bibliography

- Arpaly, Nomy. 2006. *Merit, Meaning, and Human Bondage: An Essay on Free Will*. Princeton, NJ: Princeton University Press.
- Baatz, Christian. 2014. "Climate Change and Individual Duties to Reduce GHG Emissions." *Ethics, Policy & Environment* 17 (1):1-19. doi: 10.1080/21550085.2014.885406.
- Björnsson, Gunnar. 2011. "Joint responsibility without individual control: Applying the Explanation Hypothesis." In *Moral Responsibility: beyond free will and determinism*, edited by Nicole Vincent, Ibo van de Poel and Jeroen van den Hoven, 181–99. Dordrecht: Springer.
- Björnsson, Gunnar. 2014. "Essentially Shared Obligations." *Midwest Studies in Philosophy* 38 (1):103–20.
- Björnsson, Gunnar. 2017a. "Explaining (away) the epistemic condition on moral responsibility." In *Responsibility—The Epistemic Condition*, edited by Philip Robichaud and Jan Willem Wieland, 146–62. New York: Oxford University Press.
- Björnsson, Gunnar. 2017b. "Explaining Away Epistemic Skepticism about Culpability." In *Oxford Studies in Agency and Responsibility*, edited by David Shoemaker, 141–64. Oxford University Press.
- Björnsson, Gunnar, and Bengt Brülde. 2017. "Normative Responsibilities: Structure and Sources." In *Parental Responsibility in the Context of Neuroscience and Genetics*, edited by Kristien Hens, Daniela Cutas and Dorothee Horstkötter, 13-33. Cham: Springer International Publishing.
- Björnsson, Gunnar, and Kendy Hess. 2017. "Corporate Crocodile Tears? On the Reactive Attitudes of Corporate Agents." *Philosophy and Phenomenological Research* 94 (2):273–98. doi: 10.1111/phpr.12260.
- Björnsson, Gunnar, and Karl Persson. 2012. "The Explanatory Component of Moral Responsibility." *Notûs* 46 (2):326–54. doi: 10.1111/j.1468-0068.2010.00813.x.
- Brennan, Geoffrey, and Geoffrey Sayre-McCord. 2015. "Voting and Causal Responsibility." *Oxford Studies in Political Philosophy* 1:36–59.
- Broome, John. 2012. *Climate Matters*. W. W. Norton.
- Budolfson, Mark Bryant. unpublished manuscript. Collective Action, Climate Change, and the Ethical Significance of Futility.
- Collins, Stephanie. 2013. "Collectives' Duties and Collectivization Duties." *Australasian Journal of Philosophy* 91 (2):231–48. doi: 10.1080/00048402.2012.717533.
- Cripps, Elizabeth. 2013. *Climate change and the moral agent: individual duties in an interdependent world*. Oxford University Press.
- Dietz, Alexander. 2016. "What We Together Ought to Do*." *Ethics* 126 (4):955-82.
- Driver, Julia. 1992. "The suberogatory." *Australasian Journal of Philosophy* 70 (3):286-95. doi: 10.1080/00048409212345181.
- French, Peter A. 1984. *Collective and Corporate Responsibility*. New York: Columbia University Press.
- Held, Virginia. 1970. "Can a Random Collection of Individuals be Morally Responsible?" *The Journal of Philosophy* 67 (14):471-81.
- Hess, Kendy M. 2014. "The free will of corporations (and other collectives)." *Philosophical Studies* 168 (1):241-60. doi: 10.1007/s11098-013-0128-4.
- Hiller, Avram. 2011. "Climate Change and Individual Responsibility." *The Monist* 94 (3):349–68.
- Hursthouse, Rosalind. 1999. *On Virtue Ethics*. Vol. 2: Oxford University Press.
- Isaacs, Tracy. 2011. *Moral Responsibility in Collective Contexts*. Oxford: Oxford University Press.
- Kagan, Shelly. 2011. "Do I Make a Difference?" *Philosophy & Public Affairs* 39 (2):105-41. doi: 10.1111/j.1088-4963.2011.01203.x.
- Lawford-Smith, Holly. 2015. What 'We'? In *Journal of Social Ontology*.
- Lawford-Smith, Holly. 2016. "Difference-Making and Individuals' Climate-Related Obligations." In *Climate Justice in a Non-Ideal World*, edited by Clare Hayward and Dominic Roser, 64–82. Oxford: Oxford University Press.
- List, Christian, and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Maltais, Aaron. 2013. "Radically Non-Ideal Climate Politics and the Obligation to at Least Vote Green." *Environmental Values* 22 (5):589-608. doi: 10.3197/096327113X13745164553798.
- McKenna, Michael. 2012. *Conversation and Responsibility*. New York: Oxford University Press. Book manuscript.

- Nefsky, Julia. 2011. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan." *Philosophy & Public Affairs* 39 (4):364-95. doi: 10.1111/j.1088-4963.2012.01209.x.
- Nefsky, Julia. 2016. "How you can help, without making a difference." *Philosophical Studies*:1-25. doi: 10.1007/s11098-016-0808-y.
- Parfit, Derek. 1984. *Reasons and Persons*: Oxford: Oxford U. P.
- Petersson, Björn. 2004. "The Second Mistake in Moral Mathematics is not about the Worth of Mere Participation." *Utilitas* 16 (3):288-315.
- Pinkert, Felix. 2014. "What We Together Can (Be Required to) Do." *Midwest Studies In Philosophy* 38 (1):187-202. doi: 10.1111/misp.12023.
- Schwenkenbecher, Anne. 2014. "Joint Moral Duties." *Midwest Studies In Philosophy* 38 (1):58-74. doi: 10.1111/misp.12016.
- Singer, Peter. 1980. "Utilitarianism and vegetarianism." *Philosophy and Public Affairs* 9 (4):325-37.
- Sinnott-Armstrong, Walter. 2005. "It's Not *My* Fault: Global Warming and Individual Moral Obligations." In *Perspectives on Climate Change: Science, Economics, Politics, Ethics*, edited by Walter Sinnott-Armstrong and Richard B. Howarth, 285–307. Amsterdam: Elsevier.
- Smith, Angela. 2013. "Moral Blame and Moral Protest." In *Blame: Its Nature and Norms*, edited by D. Justin Coates and Neal A. Tognazzini, 27–48. Oxford University Press.
- Stern, Nicholas. 2015. *Why are we waiting?: the logic, urgency, and promise of tackling climate change*: Mit Press.
- Strawson, Peter F. 1962. "Freedom and Resentment." *Proceedings of the British Academy* 48:187–211.
- Tuomela, Raimo. 2007. *The philosophy of sociality: The shared point of view*: Oxford University Press.
- Van Zomeren, Martijn, Tom Postmes, and Russell Spears. 2008. "Toward an integrative social identity model of collective action: a quantitative research synthesis of three socio-psychological perspectives." *Psychological bulletin* 134 (4):504.
- Vance, Chad. forthcoming. "Climate Change, Individual Emissions, and Foreseeing Harm." *Journal of Moral Philosophy*.
- Wringe, Bill. 2010. "Global Obligations and the Agency Objection." *Ratio* 23 (2):217-31. doi: 10.1111/j.1467-9329.2010.00462.x.
- Wringe, Bill. 2016. "Collective Obligations: Their Existence, Their Explanatory Power, and Their Supervenience on the Obligations of Individuals." *European Journal of Philosophy* 24 (2):472-97. doi: 10.1111/ejop.12076.
- Zimmerman, Michael J. 2008. *Living with Uncertainty*: Cambridge U. P.