

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

GUNNAR BJÖRNSSON, STOCKHOLM UNIVERSITY

1. The activist's perspective

Issues like climate change, massive inequality, or threatening fascism pose *collective* practical problems: problems that will only be resolved if a great many people contribute. People who make substantial efforts to help resolve such problems often take what I will call “the activist's perspective,” thinking that:

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

COLLECTIVE OBLIGATION: In virtue of COLLECTIVE CAPABILITY, we have a moral obligation to 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 as best they can, at least when there is a high enough chance that it will be significantly resolved.

The first two elements of the activist's perspective attribute capability and obligations to some relevant group of agents. 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 CAPABILITY 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: *we have to* resolve these issues and would be to blame if we failed. 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.

2. Two problems of individual incapability and group agency

The activist's perspective seems both prevalent and important in motivating collective action. But it can also seem theoretically deeply problematic. The first problem has to do with the sheer scale of the collective issues in question:

INDIVIDUAL INCAPABILITY: 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 INCAPABILITY, it might also seem that

NO INDIVIDUAL OBLIGATION: Because of INDIVIDUAL INCAPABILITY, 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 activism: the activist's efforts are meant to effectively 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 obligations.

The AGENCY REQUIREMENT does not rule out all attributions of obligations to groups. Many have argued that groups can constitute bona fide moral agents: correctly organized, they can have their 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; Hess 2014; Björnsson and Hess 2017; for an overview, see Tollefsen 2015). It is also 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 (see e.g. Collins 2013; Lawford-Smith 2015; for discussion of the problem and proposed solutions, see Björnsson Forthcoming). Moreover, even if the groups in question can be understood as agents or potential agents, this looks unlikely to help make sense of the activist's perspective. Activists are concerned

with how to reach out and engage other *individuals* in these groups, have in mind how difficult it is for *individuals* to contribute, and think about what can be achieved if *individuals* can be motivated by the moral urgency of the issue at hand. 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 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 one global master plan or procedure to which all or even most agents sign on. It is just not part of this picture that the group is or will become a moral agent in its own right, much less that it already satisfies standard requirements of moral agency.¹

The task of this paper is to explain how, in spite of problems of individual incapability and group agency, COLLECTIVE OBLIGATION and RESULTING INDIVIDUAL OBLIGATION are coherent and sensible, in particular with respect to climate obligations. In other work, I have suggested that we should understand the relevant kind of collective obligation as *shared* obligation (Björnsson 2014a, Forthcoming; cf. Aas 2015; Schwenkenbecher 2014). Though the groups in question are not themselves moral agents, they can straightforwardly have obligations because their members are such agents—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 INCAPABILITY.

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. I provide no defense for these assumptions apart from appealing to some initial intuitive support or suggestive examples, 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 would be helped by a clear general understanding of what moral obligations are. Unfortunately for our purposes, there are several notions of moral obligation, and philosophers variously appeal to *objective*, *prospective*, and *subjective* obligations, and to *all-things-considered*, *prima facie*, and *pro tanto* obligations. On top of that, talk of what one *must* do is famously context dependent.² 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

¹ For one recent book-length treatment of climate obligations that attributes important such obligations to non-agential groups, see Cripps 2013.

² The *locus classicus* is Kratzer 1977; for my own understanding of the distinction between *must* and *ought*, see Björnsson and Shanklin 2014.

analysis of collective obligations that I propose based on this will then have to be expanded to cover other species of moral obligation.

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 moral ideals *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 demanded.) Other demands must be satisfied *on pain of making one a fitting target of sanctions or attempts to force one to comply that are more hands-on than 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.³) Partly for convenience, and partly because I think these to be an important part of common sense as well as the activist's perspective, my focus here will be on demands the violations of which normally ground blame. In this section, focus is on individual obligations:

INDIVIDUAL OBLIGATION & BLAME: X has a moral obligation to ϕ if and only if X would be morally to blame for not ϕ -ing.⁴

To better understand this notion of obligation, we will be helped by an account of when one is to blame for not ϕ -ing. Ideally, such an account should help us 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: ordinary actions and omissions as well as various cognitive and emotional reactions or lack thereof.

Such general facts, I have suggested, can be accounted for if we follow an 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).

³ 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.

⁴ This notion of obligation leaves no room for failures to discharge an obligation that are excused due to stress, duress, or disease. If such conditions leave X blameless for not ϕ -ing, then they also render it false that X *had to* ϕ in the relevant sense. For notions that allow for excuses, see the end of this section and n. 17.

As this tradition has made clear, paradigmatic cases of blaming involve thinking that someone did something bad because she didn't care to the right extent about morally relevant considerations. Conversely, a bad action is typically excused when the agent did it because of external forces or because she was ignorant of crucial information, rather than because of a lack of concern with or respect for the values involved.

Though the general idea is powerful in itself, some precisifications and qualifications are in place.

First, there are various ways in which one might understand the *will* that determines one's 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 the object in that regard and to invest resources (cognition, action, material and social means) to promote it in light of these factors.

On this dispositional account of caring, whether and in what way a certain degree of caring has cognitive or behavioral upshots depends on information and resources available to the agent as well as competing concerns. It should also be noted that this understanding of the will in terms of caring can be combined with different ideas about what objects one needs to care about for one's will to be good, 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).

Second, not any *non-ideal* caring is bad or lacking in a way that grounds blame. We are not blameworthy for every action 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. Most obviously, indignation seems misplaced when directed at an agent who lacks the capacities required for caring about certain things—certain core elements of social and moral cognition, say. More subtly, such capacities come in degrees, and the degree of caring that is demanded might vary accordingly. Furthermore, that degree will depend on the importance of the object of care, and on one's relation to that object (for an overview of relevant relationships, see Björnsson and Brülde 2017). Quality of will accounts can accept all this but differ with respect to whether they take standards for caring to be grounded in some further considerations, such as rational agency, hypothetical contracts, or costs and consequences of imposing such standards.

Third, not any event expressive of substandard caring is a relevant object of blame. The object of blame must *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. However, just as we can be required to care about a variety of objects, a quality of will account can accept a variety of objects of blame, including not only clear cases of harm, rights violations, or suboptimal outcomes, but also attempts to cause harm and failures to notice or be motivated by some morally relevant fact.

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 2014a, 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 several such instances, sometimes on part of several people.⁵

The resulting version of the quality of will account is that

INDIVIDUAL 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 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 & BLAME, it follows that X has a moral obligation to ϕ if and only if X's not ϕ -ing would be bad and explained, in a normal way, by X's substandard caring. 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 .⁶ If we say that X's ϕ -ing is "morally important" exactly when X's not ϕ -ing would be bad, we get:

INDIVIDUAL OBLIGATION: X has a moral obligation to ϕ if and only if X's ϕ -ing would be (i) morally important and (ii) ensured, in a normal way, by X's caring as can be morally demanded (Björnsson 2014a; Björnsson and Brülde 2017).

⁵ Caring about a value to a certain extent can be understood as a form of skill in promoting that value to the right extent, a skill upheld in part by noticing how well one's caring does promote the value in question. Following Ruth Millikan (1984), we can then take as normal explanations those that are of the kinds that typically explain the reproduction of such skills. (Given that skills are honed in response to concrete lessons provided by good and bad concrete upshots of levels of skill, we can also say that one is blameworthy for something bad when it is due to one's substandard caring in such a way that it provides a concrete lesson against that substandard caring; see Björnsson (forthcoming).)

⁶ That X's caring appropriately would *ensure* X's ϕ -ing does not mean that if X cared appropriately, X would ϕ *because* of this caring. (The presence of a life guard might ensure that no one drowns without actually preventing anyone from drowning; correspondingly, INDIVIDUAL OBLIGATION allows that one does what one morally must do for non-moral reasons.) Rather, it means that (i) X ϕ s in all relevant possibilities where X cares appropriately and (ii) in at least some such possibility, X ϕ -ing is explained by X's caring appropriately. That it is ensured *in a normal way* further requires that this explanation is relevantly normal.

Apart from capturing the connection between obligation (of the relevant sort) and blame, INDIVIDUAL OBLIGATION gives content to the AGENCY REQUIREMENT: for an individual to be subject to moral obligations, it must be subject to moral standards of caring, which presumably requires the capacity to care about morally important matters to the relevant degrees. Relatedly, it 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, in the sense that caring appropriately would ensure, in a normal way, that one ϕ s.

In addition, INDIVIDUAL OBLIGATION leaves room for the full range 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.⁷

This concludes my introduction of the notion of obligation that will be central to our discussion. I stress again that this notion is one among many. Specifically, what INDIVIDUAL OBLIGATION aims to capture is a notion of all-things-considered obligations that depend on available information and current levels of non-culpable confusion, stress, or duress, rather than prima facie, pro tanto, or objective obligations. To further bring out what sort of notion I have and do not have in mind, I offer the following rough analyses of some neighboring notions for contrast:

PRIMA FACIE: X has a *prima facie* obligation to ϕ when considerations obtain that would characteristically ground a regular (all-things-considered) obligation in the absence of contrary considerations.

PRO TANTO: X has a *pro tanto* obligation to ϕ if and only if X has a prima facie obligation to ϕ based on which it can be properly morally demanded that X cares about ϕ -ing.⁸

OBJECTIVE EXCUSE-FRIENDLY: X has an *objective excuse-friendly* moral obligation to ϕ if and only if X's ϕ -ing would be (i) morally important and (ii) ensured, in a normal way if X

⁷ Notably, INDIVIDUAL OBLIGATION 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 obligations rather than rightness, and appeals to appropriate caring rather than (presumably more demanding) agential virtue. For a brief discussion of how INDIVIDUAL OBLIGATION might handle some problems raised for such accounts, see Björnsson 2014a, 116–17, n13.

⁸ Pro tanto obligations are *normative responsibilities* to (not) perform certain actions, in the terminology of Björnsson and Brülde 2017.

(a) cared as can be morally demanded, (b) were governed in a normal way by her caring, and (c) had all morally relevant information.⁹

NON-OBJECTIVE EXCUSE-FRIENDLY: X has a *non-objective excuse-friendly* moral obligation to ϕ if and only if X's not ϕ -ing would be (i) morally important and (ii) ensured, in a normal way if X (a) cared as can be morally demanded and (b) were governed in a normal way by her caring.

If something that I say about obligations seems counterintuitive, I urge the reader to make sure that the intuitions concern the intended notion of obligation.

4. Irreducibly shared obligations

Return now to the difficulty of understanding obligations of groups that are not themselves moral agents. What I suggest is that such obligations can be adequately understood in terms of the following simple generalization of INDIVIDUAL OBLIGATION, treating the latter as an account of the special case involving a single individual:

OBLIGATION: One or more individuals have a moral obligation to ϕ if and only if their ϕ -ing would be (i) morally important and (ii) ensured, in a normal way, by the appropriate caring of them all (Björnsson 2014b, Forthcoming).¹⁰

No further adjustments seem to be required to capture obligations of non-agential groups, as we see when considering the various parts of the account.

First, the behavior of a plurality might obviously be morally bad, i.e. the sort of thing that there are moral reasons to prevent. For example, 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, just as there might be reason to prevent a single individual from doing those things.

Second, OBLIGATION allows for the attribution of obligations to groups that are not themselves moral agents in any qualified sense: what is required for group obligations is merely that individual members of the group are subject to moral standards of caring. (Though leaving

⁹ Unlike violations of the INDIVIDUAL OBLIGATION variety, an agent need not be blameworthy for violating her excuse-friendly obligations: she might fail to do so because she was not governed in normal ways by her caring or lacked relevant information.

Central notions of excuse-friendly obligations rely on a notion of normal governance that might best be understood in terms of Ruth Millikan's (1984) notion of a Normal explanation. Also see n. 5.

¹⁰ OBLIGATION attributes obligations non-distributively: obligations of groups are held together, not individually. For related views, see Schwenkenbecher 2018; Aas 2015; Wringe 2016; for discussion of these, see Björnsson Forthcoming. Corresponding to this account of obligations for both individual agents and groups is the following account of blameworthiness:

BLAMEWORTHINESS: One or more individuals are morally to blame for Y if and only if Y is morally bad and explained in a normal way by substandard caring among those individuals (Björnsson Forthcoming, 2011).

room for obligations of non-agential groups, OBLIGATION neatly explains why the AGENCY REQUIREMENT holds for the cases where it seems most clearly correct, namely for individual obligation bearers, if we equate moral agents with agents that are subject to moral standards of caring.)

Third, appropriate caring by all members of a group of agents can ensure, in normal ways, a certain behavior from the group. If all the teenagers care appropriately about the neighbor, this might straightforwardly ensure that each keeps reasonably quiet in the proximity of her bedroom and that they thereby keep the noise they collectively make below a certain level. If they all care appropriately about the cat, this might similarly ensure that they organize to help it down. Notice that whereas coordination and mutual awareness might be required in the second case, the first case is different. Here, as in countless other cases, appropriate caring avoids cumulative problems by separate efforts of numerous of individuals, efforts that need not be made in awareness of the other efforts. 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; it is 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.¹¹

If INDIVIDUAL OBLIGATION is correct and generalizes in this way, the teenagers' moral obligation to help the cat down, or not keep the neighbor awake seems perfectly intelligible, without assuming that the teenagers together constitute a fully-fledged moral agent. Although the standards that ultimately ground obligations—standards for 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 one or more instances of appropriate caring.¹²

Because OBLIGATION takes group obligations to consist in a number of agents *together* standing in a certain relation to ϕ -ing, I will describe these as “shared” obligations. Notice that such obligations are *irreducibly* 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

¹¹ 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 OBLIGATION is friendly to stronger control requirements. In n. 6, I said that for X's caring appropriately to ensure, in a normal way, that X ϕ , there has to be some relevant possibility in which such caring explains, in a normal way, that X ϕ . It might perhaps be argued that in order for a group's ϕ -ing to be explained in a normal way of the appropriate caring of its members, there must be some coordination between them, though perhaps not of a centralized or very sophisticated sort.

¹² Though unstructured groups can be subjects of the OBLIGATION variety of moral obligations, only groups of which it can be properly demanded that they care about doing certain things are subject to the PRO TANTO variety, as defined above. Some resistance to the idea that unstructured groups can have obligations (e.g. that of Collins 2013) might be best understood as concerned with something like the latter rather than with OBLIGATION obligations. I thank Stephanie Collins for discussion of this issue.

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. (Björnsson 2014b, Forthcoming)

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

I have said that shared obligations to resolve collective problems depend on what appropriate caring among individual agents would ensure. But given that an appropriately caring individual is unable to affect whether such obligations are discharged—given INDIVIDUAL INCAPABILITY—why would she be motivated to contribute?

In answering this question, I will argue that appropriately caring individuals are motivated to improve the expected value of the outcome of the collective problem. But given this, how can we explain the idea—codified in COLLECTIVE OBLIGATION and RESULTING INDIVIDUAL OBLIGATIONS—that activists are motivated by *the group's ability* to resolve the issue?

Given that an individual is incapable of affecting whether a collective problem is resolved, how can we account for the activist's sense that they are motivated to *help* resolve it, and their efforts to help in *more effective* ways?

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

5. Collective obligations in spite of individual incapability

Our account of shared obligations might seem to make straightforward sense of COLLECTIVE OBLIGATION for the case of threatening climate catastrophe. First, according to OBLIGATION, 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 morally 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, disease, social unrest, 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 widespread support, would speed up the process even more. Given the magnitude of what is at stake, surely the required sacrifices can be demanded.¹³

¹³ Exact estimates of required sacrifices vary considerably with assumptions made. But estimates for quick and decisive global action suggest that individual burdens would fall far short of costs regularly

On reflection, though, INDIVIDUAL INCAPABILITY 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 INCAPABILITY—it is unclear why climate threat should prompt even moderately costly action among appropriately caring individuals. That issue 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 our obligation, the difference we can make is not thereby morally irrelevant. We might vote for politicians that support important climate legislation, favor climate friendly technologies when choosing energy company and retirement plan, decrease air travel and car use, support emission reduction schemes to offset our remaining carbon footprint, 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 indeed a sensible perspective, but it too might seem to exaggerate the role we can play as individuals. The first problem is 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 decide 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 the vanishingly small increase in temperature or sea level rise that might be caused by my emissions of greenhouse gases will make a difference to 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).

All this might suggest that focus on the effects or difference-making of individual actions is a mistake, as well as a bad fit with the activist's 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 “contributes” to some outcome if it is part of why the outcome came about. (Understand “action” so as to include omissions.) When a set of actions overdetermine an outcome, individual actions can thus contribute to the outcome without making a difference to it. With this terminology, the suggestion might be put as follows:

paid by countless people when changes in labor market conditions force relocation or periods of joblessness. For one recent argument that costs are reasonable, see Stern 2015.

MERE CONTRIBUTION: If one cares appropriately about some important outcome and thinks that the collective behavior of a group of people can significantly promote (obstruct) that outcome, one cares about (not) contributing to that behavior, even if one will not thereby significantly affect the likelihood that the outcome is realized.¹⁴

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.

For cases involving collective benefits rather than harms, Julia Nefsky (2016; cf. Parfit 1984, 83) has defended a principle related to MERE CONTRIBUTION. But 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, Nefsky suggests 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 the action will contribute to the outcome and 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 already taken have 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.)

For our purposes, adding a requirement of non-superfluity makes no difference to the overall relevance of MERE CONTRIBUTION, as climate related outcomes are still very much dependent on future actions. Why, though, should we think that mere contributions matter? The most straightforward answer is that this explains the rightness and wrongness of actions in a variety of cases. The most obvious cases might be ones where the contributor intends the collective harm, such as:

Car Push: A member of the opposition has been trapped in a car. Four people are pushing the car over the edge of a cliff, more than enough to get the work done. You cannot prevent the outcome or increase its likelihood, but you nevertheless join in, intending to contribute to resulting harm.¹⁵

¹⁴ For considerations that might seem to motivate MERE CONTRIBUTION, see Brennan and Sayre-McCord 2015, 53–54; Petersson 2004; Parfit 1984, ch. 3.

¹⁵ See Sinnott-Armstrong 2005, 289–90, who advocates a principle that arguably implies MERE CONTRIBUTION for intended harms.

As Chad Vance (2017) makes plausible, however, contributing to a harm that is merely foreseen can be as serious as contributing to one that is intended: consider a version of *Car Push* where you join in to get some exercise, not caring about the harm that results from their action. In light of this, it might seem that we should accept mere contribution for the case of (non-superfluous) collective harms generally. If we care appropriately about climate harms, we will be disinclined to significantly contribute to those harms just to reap minor personal benefits, and inclined to make some personal sacrifice to contribute to climate solutions.

While cases like *Car Push* suggest that mere contributions matter, other cases suggest that mere contributions count for nil, even when they are non-superfluous. Consider:

Unhelpful Contribution: Approaching the rocky shore, I see two people trying to save a drowning man from the violent waves, with uncertain success. A bystander who I recognize as an off-duty lifeguard is clearly conflicted about getting involved, screaming that the man in the water bankrupted his family. If the bystander were to join the effort it would almost certainly be successful. I could try to convince him, but this would likely increase his reluctance. I could also contribute to the rescue effort myself, but given my physical limitations and lack of experience, that wouldn't affect the chance of success. Instead, I'm considering making things easier for another man, about to make it ashore on his own.

In this case, my contribution to the rescue effort could very well be non-superfluous: if I get involved and the person is saved, it could be because of our efforts, and if he isn't saved, it could be because too few people got involved. Still, it seems that however marginally I could improve the final stretch of the other man's struggle to get out of the water, it would be enough to trump any reason I would have to contribute to the ongoing saving effort. And the most straightforward explanation for this seems to be that such contributions add no value of their own (cf. Parfit 1984, 83; Dietz 2016, 974–75). The problem is also illustrated by cases involving mere contributions to harm:

Unhelpful Abstention: In large amounts, Medicine minimizes otherwise deadly seizures in certain patients. But some of my colleagues at the town clinic where I work as a nurse have used small amounts of Medicine to effectively help patients with stubborn headaches. As supplies are running low, a group of doctors at the clinic will together decide whether to continue this practice. If they decide to, the one seizure patient in town will be on her own until stocks are replenished; otherwise, supplies will be sufficient for her by some margin. I am aware of all this as I'm considering whether to use a small dose to help that very patient with a headache caused by her recent Medicine-minimized seizure.

Suppose that I use Medicine to relieve the patient's headache. Then I might non-superfluously have contributed to depriving her of protection against deadly seizures: if she loses the protection, it will be because we have handed out Medicine to relieve headaches, though enough of us might still refrain from doing so, in which case she keeps it. But if my action doesn't increase the risk that she loses her protection, it seems that any degree of headache relief I can offer her outweighs whatever reasons I have not to non-superfluously contribute to depriving

her of that protection. Again, this suggest that caring appropriately about moral values does not imply caring about whether one non-superfluously contributes to good or bad outcomes.

Intuitively, cases like *Car Push* suggest that intended or foreseen non-superfluous mere contributions matter, while cases like *Unhelpful Contribution* and *Unhelpful Abstinence* suggest the opposite. But this apparent conflict can be straightforwardly explained, I will now suggest, with reference to the following assumption:

INSTINCTUAL ATTRACTION/AVERSION: If one cares about some value, one will be instinctually attracted to (averse to) saliently contributing to what promotes (impedes) that value, i.e. attracted (averse) in ways not relying on a reflective assessment of how so contributing promotes (impedes) that value compared to alternative actions.¹⁶

To say that one is instinctually attracted (or averse) to salient contributions to what promotes (or impedes) one's values is not to say that one will always try (not) to so contribute. But if one rejects the option one is attracted by or takes the option that one is averse to because one thereby hopes to promote some of one's other values, one will understand this as in tension with one's commitment to the first value.

Importantly, INSTINCTUAL ATTRACTION/AVERSION makes good sense given some basic constraints on practical reasoning. At almost any moment, there are potentially infinite different things that we could do if we decided to, but our capacity to consider and carefully compare these alternatives is limited. Because of this, we need pre-reflective grounds for considering some alternatives rather than others. Part of the solution to this problem is to focus on salient kinds of options that tend to contribute to what we care about and discarding salient kinds that have the opposite tendency. Notably, options that saliently contribute to what is likely to help and harm, respectively, form exactly such kinds: what saliently contributes to what has a good chance of helping or harming itself has a pretty good chance of helping or harming, respectively. It would thus make sense to have mechanisms that pre-reflectively push us towards considering the former kind of options and away from the latter. And given such mechanisms, we would be relevantly attracted to the former and averse to the latter.

INSTINCTIVE ATTRACTION/AVERSION says that we can be attracted or averse to mere contributions. But it would also explain why a simple principle like MERE CONTRIBUTION is too strong. The reason is that the instinctual attraction or aversion will weaken when the action is no longer saliently seen as contributing to promoting or impeding the object of caring. Most obviously, this will happen when the harm is already seen as a done deal, in line with Nefsky's observation. But it will also happen when the very value that grounds the attraction or aversion prompts consideration of options that might better promote that same value: such options will not be seen as in tension with one's commitment to the value in question. This, I suggest, is why the possibility of contributing to the saving effort in *Unhelpful Contribution* seems to carry no weight against the option of helping the man who is saving himself: the latter option is superior in relation to the very same concern that would make contributing attractive, namely the

¹⁶ On this view, what one is averse to is being implicated in the collective harm in a particular way. Elsewhere (Björnsson forthcoming), I argue that being implicated in such a way grounds reactive attitudes like indignation and guilt over that harm.

concern for the people in the water. It is also why the possibility of contributing to the shortage of Medicine in *Unhelpful Abstention* seems to carry no weight against the option of using Medicine to relieve the seizure patient's headache: here the latter option is superior in relation to the same concern that would make abstaining from using Medicine attractive, namely concern for the seizure patient.

Return now to our initial conundrum. For OBLIGATION to yield the result that we have an obligation to significantly reduce the threat of climate catastrophe, it needs to be the case that if we each individually cared appropriately about the relevant values, this would ensure such a reduction. The problem was to explain how appropriate caring could ensure sufficient action on part of individuals to ensure the outcome given that no individual can affect whether the outcome takes place or not. Though MERE CONTRIBUTION seemed implausible, INSTINCTUAL ATTRACTION/AVERSION has a better chance of filling that gap. Moreover, it identifies the relevant motivation as based on instrumental concern with the outcomes at stake, thus capturing an important teleological element of the activist's perspective.

On reflection, though, it is not clear that this source of motivation is sufficient. One immediate problem is that the strength of the relevant attraction and aversion relies on the salience of the contribution. Because individual contributions to climate-related harms are indirect and diffuse, the obligations will likely be weak, leaving us shy of a full explanation. Another problem arises for those who are interested in promoting the good overall and are seeing climate-related benefits as elements of that good. Given that perspective, if a mere contribution to climate harm ever so slightly improves the situation with respect to the overall good, it will not be seen as in tension with commitment to that good: the situation will be analogous to that of *Unhelpful Abstention*. Because of these limitations of our attraction and aversion to certain mere contributions, more needs to be said to ground COLLECTIVE OBLIGATION in relation to climate catastrophe prevention.

The literature contains the basic elements of what I think is the right complementary response. Reflecting on worries about buffer zones, many have pointed out 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). To balance the significance of possible outcomes with their likelihood, our decisions should arguably be based on something like the relative expected values of available alternatives (i.e. by the sums of the products of the value of each of the possible outcomes and their probability for each alternative). 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 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.¹⁷

Still, one might think that because of the vast 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 (Sinnott-Armstrong 2005, 289–91). Perhaps causation of climate related harm and benefit is an emergent phenomenon that cannot be pinned on the quickly dispersed contributions of any single individual (Kingston and Sinnott-Armstrong 2018, 175–76). And perhaps the climate-related consequences of our actions are metaphysically indeterminate rather than merely in principle individually unknowable. If so, then even if the expected degree of climate-related harm is a positive function of the totality of individual contributions, perhaps there is not only low but *zero* chance that one individual's contributions will actually have any definite harmful or beneficial climate-related consequences.

If only the comparative likelihood of concrete, identifiable harmful or beneficial consequences of action were relevant for decision-making, this could undermine duties to reduce one's own carbon footprint or to otherwise contribute to prevention. On reflection, though, 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 OBLIGATION: Acting in ways that increase the number of expected deaths is morally bad, and

¹⁷ Some take that these and similar assumptions about expected value to be insufficiently motivated (see e.g. Maltais 2013, 593–4; Kingston and Sinnott-Armstrong 2018, 178–81). However, since my ambition is to make the activist perspective intelligible rather than showing that it is correct, I will not further defend them here.

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 general 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.¹⁸

One might agree with me about a case like *Carcinogen*, but object that in the case of climate change, the effect that a normal individual's actions has on the expected value of any one individual will be much lower. If the negative expected value of an afternoon's gas guzzling is the equivalent of a ruined afternoon, for example, the decrease in expected value per individual will not even be one billionth of that. This, one might think, makes these effects on individuals not only very small, but normatively completely insignificant.¹⁹ Something close to this thought is plausibly correct: we cannot be morally required to take one of these, on its own, into account in our deliberation. On reflection, though, even individually morally insignificant differences add up. To see this, consider

Wheel of Fortune: Consider two possible states of a world of 10 billion inhabitants. For each, we can represent the relative expected value of an inhabitant's remaining life—their “expectations”—as a position on a circle. At 12 o'clock we find the individual whose expectations are the highest, in the next position, counting clockwise, the individual whose expectations are second highest, and so forth, with the 10 billionth individual, whose expectations are the lowest, just next to the first, closing the circle. (For simplicity, suppose that there are no ties.) The difference in expectations between the first and the last individual is huge, but the difference between any other two adjacent individuals is plausibly morally negligible when taken on its own.

The two states of this world, A and B, differ in this regard: In B, the ordering is turned one step compared to A such that each individual has the expectations that the individual in the next position had in A.

¹⁸ Whether an increased risk of harm itself constitutes harm, or perhaps more weakly injustice, might matter for the strength of the resulting individual obligations (see e.g. Broome 2012, ch. 4–5). Perhaps, if the effects of my carbon footprint on expected value constitute neither harm nor injustice and 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. Now, because marginal effects on expectations of large numbers of people add up (as I briefly argue below), and because the poor will bear most of the negative effects of climate change, I do think that sustaining an uncompensated-for carbon footprint typical of citizens of rich countries typically does constitute 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 significantly cut net carbon emissions.

¹⁹ Cf. 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.

Imagine that we have to decide whether A or B is realized and have no further information about the two states. Then both of the following seem undeniable:

NO REASON: Since A and B are qualitatively identical with respect to all known morally relevant properties, we do not, *on balance*, have reason to choose one over the other.

BETTER OFF: Since the person worst off in A is so much better off in B, we do have *a* reason to choose the latter over the former.

Why is BETTER OFF not a conclusive reason to choose B? The most straightforward answer is that the individually insignificant differences add up, without loss. Whatever claim the person who is worst off in A would have on you to choose B, it can be thus outweighed by the fact the everyone else has their expectations lowered ever so slightly. (For the argument on which *Wheel of Fortune* is modelled and further discussion, see Zach Barnett 2018.)

Suppose, as *Carcinogen* suggests, that we can be required to care about effects on expected value even when concrete effects cannot be pinned on us. Suppose further, as *Wheel of Fortune* suggests, that the effects of our actions on individual expectations add up. And suppose, as previously suggested, that 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 INCAPABILITY. It also plausibly follows 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, especially in conjunction with INSTINCTUAL ATTRACTION/AVERSION.

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 towards carbon neutrality, thus significantly reducing climate threat. And from this, together with OBLIGATION, it would follow that we have a shared obligation to significantly reduce climate threat.²⁰

6. Remaining problems of collective capability and individual incapability

Thus far we have seen how the activist's attributions of collective obligations can be understood as concerned with *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 our *collective* ability to resolve the issue can ground our shared obligation (as postulated by COLLECTIVE OBLIGATION), and thus indirectly ground our individual obligations to help (as postulated by RESULTING INDIVIDUAL OBLIGATION). Importantly, OBLIGATION 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, it involves something more: our collective capacity to resolve an issue is part of what should motivate us to contribute. By contrast, the considerations of capability mentioned in the previous section were concerned primarily with the *individual's* capacity to marginally affect expected climate value, contribute to climate related benefits, or avoid contributing to climate related harms. More thus needs to be said.

Second, according to the activist's perspective (as stipulated by RESULTING INDIVIDUAL OBLIGATION), our shared climate obligation to resolve the issue grounds individual (prima facie) obligations to help resolve it as best one can. The problem here is that 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

²⁰ I've suggested that appropriately caring individuals would take steps jointly sufficient for significant threat reduction. But might not individuals who care about climate related harms fail to draw the conclusion that they have considerable reason to contribute? After all, a number of philosophers reject that conclusion. And if sufficiently many draw the same conclusion, might appropriate caring fail to ensure significantly prevent catastrophic climate change, thus 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 OBLIGATION 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, OBLIGATION-related, excuse-friendly, notion of obligation:

OBLIGATION*: One or more individuals have a moral obligation to ϕ if and only if their not ϕ -ing would be morally bad and their ϕ -ing would be ensured, in a normal way, by the appropriate caring and *teleologically rationally* of them all.

matter, especially if an individual's actions might have no determinate climate related upshots.²¹ A first answer, in line with INSTINCTUAL ATTRACTION/AVERSION, would be that to help is to contribute to a resolution. But this fails to account for the activist's sense that individual contributions can be more or less *effective* in helping to resolve the grand issue. Again, more needs to be said.

In this final section, I explain how concerns about important large-scale outcomes and considerations of collective capability 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: what I will call *issue-oriented* 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 or group of agents, which may or may not be identified and may or may not involve the agent engaged in reasoning.²² 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 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 to 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

²¹ 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.

²² As should be clear from this outline, the sort of reasoning I have in mind here is neither "we-mode" nor "I-mode" reasoning (in the terms of Tuomela 2007), though it might conceivably involve either.

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 ways that the individual can further a certain means, by making it more likely to be implemented, raising its likelihood of success, or otherwise 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 the issue, if anything.

We can easily see how both considerations of collective capability 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 capability 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 to decrease the likelihood that a 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 large-scale vagueness of the outcome, 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 capability can ground both a shared obligation to significantly reduce climate threat and individual obligations to contribute to such reduction, in spite of INDIVIDUAL INCAPABILITY. 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.

Let me end this section by considering an objection to these last claims. It seems painfully obvious 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.²³

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

²³ 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.

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 fair share of a common resource without trying to compensate others does not obviously mean that I should be willing to use more than 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 our individual inability to affect whether the threat is averted or not, we do have individual *prima facie* obligations to help achieve resolution, obligations grounded in the same facts as our shared obligation.

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 explained how individuals can have obligations to help achieve climate threat resolution even though they cannot affect its likelihood. Both considerations of collective capability 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, Olle Blomberg, Stefano Cossara, Daniel Telech, Holly Lawford-Smith, Erik Brandstedt, Mattias Gunnemyr, Olle Risberg, Richard Endörfer, Mark Budolfson and Tristram McPherson, Michael Morek, Adrian Russian, Coby Gibson, Maxwell Frye, and Ira Richardson. Work on this paper has been funded by the Swedish Research Council.

References

- Aas, Sean. 2015. "Distributing Collective Obligation." *Journal of Ethics and Social Philosophy* 9 (3):1-23.
- 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.
- Barnett, Zach. 2018. "No free lunch: The significance of tiny contributions." *Analysis* 78 (1):3–13.
- 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. 2014a. "Essentially Shared Obligations." *Midwest Studies in Philosophy* 38 (1):103–20.
- Björnsson, Gunnar. 2014b. "Incompatibilism and 'Bypassed' Agency." In *Surrounding Free Will*, edited by Alfred Mele, 95–122. New York: Oxford University Press.
- 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. 2020. "Collective responsibility and collective obligations without collective agents." In *The Routledge Handbook of Collective Responsibility*, edited by Saba Bazargan-Forward and Deborah Tollefsen, 127–41. New York: Routledge.
- Björnsson, G (forthcoming). "Being Implicated: On the Fittingness of Guilt and Indignation over Outcomes," *Philosophical Studies*.
- 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.
- Björnsson, Gunnar, and Karl Persson. 2012. "The Explanatory Component of Moral Responsibility." *Noûs* 46 (2):326–54.
- Björnsson, Gunnar, and Robert Shanklin. 2014. "'Must', 'Ought' and the Structure of Standards." In *Deontic Logic and Normative Systems*, edited by Fabrizio Cariani, Davide Grossi, Joke Meheus and Xavier Parent, 33-48. Springer International Publishing.
- 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.
- Cripps, Elizabeth. 2013. *Climate change and the moral agent*: 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.
- French, Peter A. 1984. *Collective and Corporate Responsibility*. New York: Columbia University Press.
- Hess, Kendy M. 2014. "The free will of corporations (and other collectives)." *Philosophical Studies* 168 (1):241-60.

- 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.
- Kagan, Shelly. 2011. "Do I Make a Difference?" *Philosophy & Public Affairs* 39 (2):105–41.
- Kingston, Ewan, and Walter Sinnott-Armstrong. 2018. "What's Wrong with Joyguzzling?" *Ethical Theory and Moral Practice* 21 (1):169–86
- Kratzer, A. 1977. "What 'must' and 'can' must and can mean." *Linguistics and Philosophy* 1(3):337–55
- Lawford-Smith, Holly. 2015. "What 'We'?" In *Journal of Social Ontology* 1(2): 225–49.
- 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.
- McKenna, Michael. 2012. *Conversation and Responsibility*. New York: Oxford University Press.
- Millikan, Ruth Garrett 1984. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.
- Nefsky, Julia. 2011. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan." *Philosophy & Public Affairs* 39 (4):364–95.
- Nefsky, Julia. 2016. "How you can help, without making a difference." *Philosophical Studies*:1–25.
- 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.
- Schwenkenbecher, Anne. 2014. "Joint Moral Duties." *Midwest Studies In Philosophy* 38 (1):58–74.
- Schwenkenbecher, Anne. 2019. "Collective Moral Obligations: 'We-Reasoning' and the Perspective of the Deliberating Agent." *The Monist* 102 (2):151–71.
- 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.
- Tollefsen, Deborah Perron. 2015. *Groups as agents*: John Wiley & Sons.
- 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. 2017. "Climate Change, Individual Emissions, and Foreseeing Harm." *Journal of Moral Philosophy* 14(5):562–84.
- 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.