

Participation and Collective Harm

Rutger van Oeveren

Jan Willem Wieland

Abstract

Many global problems nowadays are collective action problems. It is only because many people act in some way that certain problems arise (e.g., climate harms). But when can one be said to participate in such groups in the first place? As Derek Parfit asked, “When we appeal to what groups together do, whom should we count as members of these groups?” As Julia Nefsky has shown, this problem has proven to be notoriously tricky, and without any adequate solution. We present a solution. The proposed account not only overcomes a number of problems for participation-based views, according to which it is morally problematic to participate in such groups at all, but also helps make precise a prominent alternative view, according to which one should help bring about better outcomes.

Keywords:

Collective harm, participation, helping, superfluity problem, disconnect problem

Journal of Moral Philosophy

1. Introduction

There are situations in which a set of actions together make a difference to some morally relevant outcome, but individually do not. One need think only about large swaths of consumer choices, such as choices regarding energy, transportation, clothing, and food, with outcomes such as global warming and the exploitation of sweatshop workers (cf. Lichtenberg 2010). It is easy to think in such cases that we have a duty, or at least very good reason, to refrain from joining such collective actions. However, if an individual person makes no relevant difference, one may wonder whether there is any moral reason whatsoever to act in one way or another. Thus, a question arises as to the moral status of a large group of common, everyday actions, actions that we tend to think of as morally problematic. This is the collective action problem.

Let us consider the following analogy:¹

Hot Day

It is a very hot summer day. Thousands of people have gathered around a small lake. It is so hot that everyone wants to cool down in the water. Yet, if too many people go into the lake, the water level of one part of the lake will rise to an unsafe level, the people in it will not have the space to rescue themselves and will drown.

In terms of this case, the question is: why refrain from cooling down in the lake if you make no difference to whether or not anyone drowns?

These and like cases might have no thresholds: no sharp cut-off for the size of the group before which the relevant outcome fails to come about, and at or after which the outcome will come about (cf. Nefsky 2017).² It presumably does not take an *exact* number of people to raise the water level sufficiently to drown people (i.e., there is no integer n such that they are drowned after n swimmers, yet not after $n-1$). It all depends on how

¹ Inspired by similar cases in Glover (1975), Parfit (1984, ch. 3), and Barry and Øverland (2016, ch. 11).

² More precisely, either there is no sharp boundary at or after which there is *any* difference in harm (or benefit), or there is no sharp boundary at or after which there is *any perceptible* difference in harm *and only perceptible differences in harm are morally relevant* (cf. Gunnemyr 2022). We will remain agnostic about whether only *perceptible* differences in harm matter morally, and simply assume that in many of the given cases, there is no such sharp boundary for the morally relevant outcome. If you think that this is incorrect, i.e., if you think that there are imperceptible differences in harm and benefit *and* think that such differences matter morally, you may read the rest of this paper as arguing for a conditional thesis.

much each contributes to the water level, and on other unpredictable factors such as the wind, the current, or the soil. If this is the case, no individual could possibly make a difference to that outcome. This makes the problem even harder. For, in such cases, we cannot straightforwardly appeal to chances that you make a difference and expected value (cf. Kagan 2011, though see Hedden 2020).³

One promising solution – originally due to Parfit (1984, ch. 3) – is in terms of *participation*. According to participation accounts, an agent S has a reason to refrain from doing X if by doing X, S participates in a group that harms others. By going into the lake, you participate in the group that drowns others, and this gives you reason not to do it.⁴

However, as Julia Nefsky has shown, such solutions face additional difficulties: the so-called Superfluity Problem and the Disconnect Problem. In terms of participation, the Superfluity Problem is basically the question: if my going into the water is superfluous and makes no difference, then why should it count as “participating in a group that drowns others”? As Nefsky puts it: “It is unclear ... why I should count as part of the group that harms if what I do makes no difference with respect to that harm” (2015, 249).

The Disconnect Problem is the problem “why each of us has moral reason specifically to do the sort of thing that if enough of us do it, will prevent or avoid causing the serious avoidable harm” (Nefsky 2023, 43). Suppose that I could avoid participating in the group that drowns others by holding up a banner saying, “DON’T LET THEM DROWN.” Then, the reason why I should participate would not be such that, if everyone acts on it, these people would not be harmed. For if we all merely held banners *rather than refrained from going into the lake*, the people would still drown. In general, then, solutions to the collective action problem need to identify a reason “specifically to do the sort of thing that if enough of us do it, will prevent or avoid causing the serious avoidable harm.”

This paper’s ambition is to solve these two problems due to Nefsky. Our account tells us *when*, in general, people participate in bringing about collective harms (or collective failures to benefit), and so tells us *when* people are acting wrongly (that is, if participation theorists à la Parfit are right that participation in itself is at least pro tanto morally wrong⁵).

³ Nor to the act being a “necessary element of a sufficient set” (NESS). In short, for any sufficient set, an action cannot make a difference as to whether the set is sufficient or not; hence, it cannot be a necessary element of a sufficient set (cf. Wieland and Van Oeveren 2020, 175-176). The same problem afflicts membership criteria initially considered in Parfit (1984, ch. 3).

⁴ Moreover, the proposal is: S has a reason to do X if by doing X, S participates in a group that *benefits* others. Whether collective harm and collective benefit cases are indeed symmetrical in this way is questionable, though we will set that aside here and focus on collective harm cases.

⁵ In this paper, we are sympathetic to Parfit’s position here. This leaves it open that, depending on one’s normative perspective, *certain types* of participation will appear especially problematic (or

We will also show that the conditions we propose below can be inserted into further, existing solutions to the collective action problem. For, it is intended to provide a necessary condition for what it means to participate in collective actions in the first place. Think, for example, of Kutz’s account based on the notion of a “participatory intention.” In order to participate, you must act with an intention to do your part (as Kutz proposes), but also, we will argue, participate in our sense. Or think of Nefsky’s account based on the notions of “helping” and “being instrumentally significant.” One way of understanding her account (in response to worries that we will discuss below), is exactly in terms of participation in our sense.⁶ We will return to these accounts in due course.

The paper is structured as follows. First, we describe the Superfluity Problem, as well as an updated version of this problem, and present our solution – and indeed our account of what it means to “participate.” Next, we address the Disconnect Problem. We add that Nefsky’s own helping-based account faces similar issues, and that our account can actually be used to accommodate them. Finally, we distinguish participation-based and helping-based reasons more clearly.

2. Superfluity Problem 2.0

Here is a first general statement of the participation account, call it *Naïve Participation* (considered by Parfit 1986, 848):

Naïve Participation

S has a participation-based reason not to do X iff: X is such that because enough others do that kind of act, they together cause bad outcome O*.⁷

The basic idea is this: even if you make no difference to O and can say “if I had acted differently, the outcome would not have been better,” you are still one of those who caused

commendable). It is up to the proponents of various normative theories to utilize the account provided in explaining what is wrong or right about (certain types of) participation.

⁶ For a further account, see Andreou’s (2014) account based on the notion of a “doing in progress,” and the issue of *when* an individual action takes part in such a doing. There are still stronger forms of participation, that in ordinary parlance we would commonly denote with “cooperation” (cf. Salomone-Sehr 2022).

⁷ Or: to *do* X iff X is such that because enough others do that kind of act, they together cause *good* outcome O. Throughout the paper, “O” will refer to a good outcome, and “O*” to the corresponding bad one.

O. However, there are two problems with this account (cf. Wieland and Van Oeveren 2020, 176-177). First, acting in the same way is not *sufficient* for participation. Take the case of driving cars and global warming. If one drives a car on green energy, driving cars is insufficient for joining the group that causes climate harms. To this first problem, one might reply that much depends on the specificity of the description: green energy driving as opposed to petrol driving. No matter how many people drive cars on green energy, they together will not cause climate harms. In contrast, driving petrol-fueled cars *does* suffice for participation.

But there is a second problem: acting in the same way is not *necessary* for participation either. One can participate in the group that causes climate harms in many ways: some drive cars, some run polluting factories, others invest in these factories, others sell and buy products from them, still others keep lots of cows, others eat them, and so on. Even so, as we notice in our (2020), there is still something common to all these actions. That is, they all contribute to the same “underlying dimension”: the amount of greenhouse gases in the atmosphere, the amount of force exerted on the boulder, the amount of chicken sold in the supermarket, the amount of votes cast for a certain candidate, and so on. The proposal was:⁸

Contributory Participation

S has a participation-based reason not to do X iff:

- (i) X adds to an underlying dimension D; and
- (ii) because enough others add to D, D causes bad outcome O*.

For example, in Hot Day, you belong to the group of agents who together drown people when you: (i) contribute to the water level in the lake, and (ii) people drown because too many people make such a contribution.

Contributory Participation proceeds in two steps. First, S can make a difference to D. Even if one can make no difference to whether or not people drown in the lake, one can still make a difference to the water level. Second, D can make a difference to O as long as enough people add to D. Sufficient contributions to the water level can make a difference to whether or not people drown. This, we claim in (2020), solves the Superfluity Problem. That is, you cannot say that, since your act is superfluous, it is not part of the set of acts

⁸ We call it “Contributory Participation” because the account is about *contributing* to some underlying dimension. Again, a similar principle may hold for collective benefits rather than harms: S has a participation-based reason to *do* X iff: (i) X adds to an underlying dimension D; and (ii) because enough others also add to D, D causes *good* outcome O.

that together makes a difference for the worse. For, even when you make no difference to O, you can still add to D, and participate in this way.⁹

Let us add two points. Firstly, at the time of acting, different things might be true about bad outcome O*: it is guaranteed that O* will be brought about (because too many people are adding to D), or this is only possible. This distinction has some interesting implications that will be discussed towards the end of the paper. Secondly, we take the account (and subsequent versions) to be compatible with an objective or a subjective reading of those reasons – roughly: what participation-based reasons we have is determined by whether conditions (i) and (ii) actually obtain, or by whether we (ought to) believe or know that conditions (i) and (ii) obtain. For the most part, we will simply discuss these accounts as talking about objective reasons.

Regardless of such details, Nefsky argues that Contributory Participation faces a number of problems, starting with a variant of the Superfluity Problem: Superfluity Problem 2.0. She imagines the following collective benefit case:

Vending Machine

Three people, A, B, and C, are walking in a national park, when they come across two hikers who have been lost for several days in the backcountry, and who have finally stumbled back to the trails. They are starving – they have not eaten in days. Luckily, there is a vending machine nearby, selling granola bars for \$4 a piece. The vending machine accepts all coins and bills, but it does not give change. After putting money in, you press a button, and the number of granola bars this amount of money buys drop out. (So if you put in a \$10 bill and press the button, two granola bars will come out, just as if you had put in \$8.) The two starving people do not have any money on them. But A has a \$5 bill, B has a \$10 bill, and C has only a quarter. There is no one else around (2023, 46).

According to Contributory Participation, C has a participation-based reason to donate the quarter because: (i) donating it adds to the amount of money in the vending machine, and (ii) the starving hikers will be helped when A and B also add their money. Yet, in contrast to what Contributory Participation predicts, C should not put her quarter in. It is a total waste to add it, given what A and B can add.

⁹ Note that “adding to D” is meant as a technical notion. Some cases do not seem to be about adding, but about “subtracting.” In the case of overfishing, for example, fishermen subtract fish from the sea. This is a matter of terminology: when fishermen subtract from the amount of fish in the sea, they add to the amount of “caught fish.”

Nefsky poses the same problem in terms of a collective harm case (which can be seen as a variant of Hot Day):

Locks

There is a wide canal with a large region in which the water level is controlled by locks. Right now, the water level is low, around knee deep. But when the locks are open, the water level will rise high, well above head height. Austin Powers is currently, while the water is low, walking in the middle of the canal looking for a key that has fallen in and that is crucial for his mission. He does not know how to swim. Dr. Evil has anticipated this situation, and he has had his henchman, Mustafa, knock out the person who was manning the locks. Now Mustafa and Dr. Evil are opening the locks, in an effort to drown Powers. You are on the land nearby, and you see what is going on. You decide to slip into the water to swim to the middle of the canal where Powers is and try to rescue him. You don't know if you will make it in time, but you think it's worth a try (2023, 48).¹⁰

According to Contributory Participation, you have a participation-based reason not to jump in the water and rescue Powers because: (i) jumping into the water raises the water level, and (ii) Powers drowns because others – Mustafa and Dr. Evil – also raise the water level. Yet, in contrast to what Contributory Participation predicts, you do not participate in the group that is (potentially) drowning Austin Powers: “the claim that this instrumentally irrelevant, totally harmless increase to the water level makes you a participant in what Dr. Evil and Mustafa are doing seems false. At least, it does not track any familiar or normal notion of participation” (Nefsky 2023, 49). If this is right, you do not have a participation-based reason, not even pro tanto, to refrain from going into the water.

The intuition is that one's contribution (to the money in the machine, or the water level) is superfluous. But what does that mean? A quick fix might be to add the qualification, “unless O/O* can be realized (e.g., the hikers can be helped or Austin Powers can be drowned) even without one's contribution.” Unfortunately, one cannot appeal to such a qualification. For, this qualification holds in *all* collective action cases without thresholds, and where one's contribution is never necessary.

Another suggestion would be to appeal to Kutz's (2000) account in terms of participatory intentions. After all, in Vending Machine, C does not intend to form a group with A and B and does not donate her quarter with the intention to do her part in helping the starving hikers. Indeed, she thinks that it is a total waste to add it given what A and B

¹⁰ Cf. Björnsson (2021, 265-266) for a similar case.

can add. Yet, we do not think that Kutz's account can be applied in this way. Particularly, the analysis does not work in cases where agents *do* have such participatory intentions (cf. Driver 2015, 72-73). Even if for some reason, C were to act with a participatory intention (i.e., to do her part in helping the hikers), then that would not make her a participant, or in any case, not give her reason to donate her quarter.¹¹

Alternatively, one may suggest that you do not count as a participant in these cases because your contribution to D and the contributions of the other agents are not *of the same type*. In Vending Machine, the starving hikers cannot be helped if A and B also put a quarter in. A and B do not perform the same kind of act as C, and so C does not seem to belong to their group. And in Locks, you do not act in the same way as Mustafa and Dr. Evil. For you jump into the water, while they open the locks. And so, again, you do not seem to belong to their group. Looking back at Naïve Participation, this seems to be exactly what that account would say about these cases. In this respect, Naïve Participation is on the right track: you do not belong to that group. Even so, Naïve Participation is problematic for the reason discussed, and we cannot appeal to it. What is needed is a more precise explanation of why you do not count as a participant in these cases.

3. Contributory Participation 2.0

One conclusion that one might draw from these cases is that participation-based considerations are just morally irrelevant. But that would be to throw out the baby with the bathwater. Instead, and contrary to Contributory Participation, we think that you are not really acting with others in the problem cases just discussed. You do not really count as a participant in some relevant group and are not one of those who together harm (or fail to benefit). When Parfit asks "Will my act be one of a set of acts that will together harm [or benefit] other people?" (1984, 86), the answer that we want to give in cases like Vending Machine and Locks is "No."

Our proposal is that participation-based reasons kick in *only when there is a group of agents who cannot make a difference individually, but can make a difference collectively*. The underlying idea is this: even when you cannot make a difference on your own, at the

¹¹ Kutz himself is well aware of the limitations of his account, especially when it comes to "unstructured" collective harms, where agents do not act with a participatory intention to contribute to those harms (2000, ch. 6). While such unstructured cases provide a counterexample to Kutz's participatory intention condition as being necessary for participation-based reasons, this version of Vending Machine shows that it is not sufficient either.

very least, the group that you join should still be able to do so. For if the group cannot even make a difference, then it does not matter whether or not you join it. It does not matter if you join a group that harms or helps no one. More precisely:¹²

Contributory Participation 2.0

S has a participation-based reason not to do X iff:

- (i) X adds to an underlying dimension D; and
- (ii) there are (or will be) enough potential contributors such that (a) their contributions to D together could make a difference to O*,¹³ but (b) individually cannot make a difference to O*.

For example, you have a participation-based reason to refrain from going into the lake in Hot Day because: (i) doing so raises the water level of the lake; and (ii) there are enough potential contributors who (a) together could raise the water level in such a way that people will drown, though (b) individually cannot make such a difference.

Back to Locks and Vending Machine: in these cases, there are no other agents who can collectively *but not individually* make a difference (for, they *can* make a difference individually), and so you do not participate in any such group (and so have no participation-based reasons to (not) act in a certain way given Contributory Participation 2.0). In Vending Machine, clause (ii)(b) fails because all other available agents (namely A and B) *can* make a difference to the starving hikers. Similarly, in Locks, you are not a participant because all other available agents (Dr. Evil and Mustafa) *do* make a difference to Powers' situation,

¹² This principle is meant to apply to all and only cases in which you individually cannot make a difference to O (either because it is a non-threshold case, or in cases such as Vending Machine). Presumably, one can also participate in groups when one *can* make a difference to O*. We take no stance on the conditions for potentially difference-making forms of participation. We state the view as a necessary and sufficient condition, though the necessary condition is most essential for current purposes. Whether Contributory Participation 2.0 is more than a participation criterion and indeed sufficient for providing participation-based reasons depends, as we said, on whether one accepts Parfit's view that participation itself is reason providing. The accounts of Kutz and Nefsky will require additional conditions: a participatory intention and instrumental significance, respectively (see Section 5 for Nefsky's account).

¹³ As we discuss below, there can also be participation-based reasons in "closed" cases where O* is certain to come about. (ii)(a) as stated here is somewhat ambiguous about closed cases, and should be understood as shorthand for "*as far as the aggregate size of their contributions to D is concerned, their contributions to D could make a difference to O*.*"

and there is no group of agents that you would join, who only together, but not individually, can make a difference.¹⁴

A note on the idea of *potential contributors* is in order. The potential contributors must exist in the *actual world*.¹⁵ It is not enough if in Vending Machine, there *could have been* many others with a quarter – while in fact they are not and will not be around – and that you would join that merely counterfactual group. What matters is whether such others in fact are or will be around, and whether you can form a group with them and together harm or fail to benefit others. (Note that our principle only requires that there is *some such* group, though this need not include *all* the agents in the given situation. We will discuss examples in due course.)

It is instructive to contrast our solution to Superfluity Problem 2.0 with Nefsky’s analysis of the issue. She divides solutions to the collective action problem into “instrumental” and “non-instrumental” approaches:

Instrumental approaches attempt to show that it is not actually true that the individual act is instrumentally merely superfluous. ... Non-instrumental approaches ... point to other sorts of reasons for action – reasons that are not supposed to be about trying to influence the outcome. (2023, 40)

Nefsky avers that “it is easy to smuggle in an intuition that a contribution along the underlying dimension does matter instrumentally.” This, according to her, is what explains why it is initially appealing to think of Contributory Participation as providing a moral reason. We would agree that *one* way to make sure Contributory Participation does not give the wrong verdict in Nefsky’s cases is indeed by “smuggling in” a condition that the contribution must matter instrumentally. In Vending Machine and Locks, one’s contribution does not seem to matter instrumentally. C will not help out by adding the quarter, and neither do you contribute to Dr. Evil’s plan by jumping into the water. (However, the problem runs deeper than Nefsky recognizes. As we will show later, Nefsky’s very own helping account faces the same problems.)

¹⁴ A variant of Locks is conceivable where you do join such a group; see Section 7.

¹⁵ Not just counterfactually, cf. Wieland (2022) and Nefsky (2023). One might also think that it matters whether the contributions result from agents or from other types of contributors. Additionally, one might require those contributions to be *acts*, or even *intentional* acts (cf. Fanciullo 2020). We take no stance on this here, though important is that the others in the group act on their own behalf, and are not to be taken as a single actor (as we will discuss later).

Our point is that we *need* not appeal to instrumental significance. In fact, our Contributory Participation account (2020) was never about “influencing the outcome,” and neither is Contributory Participation 2.0. (So-called “closed” cases, that we will address later, provide good test cases: cases in which it is certain that the bad outcome will come about, and so cannot be influenced any longer.) Put differently, the account does not assume that the action has to play an instrumental role, thereby allowing for the possibility of a non-instrumental participation-based approach to the collective action problem.¹⁶

4. Disconnect Problem

There is a second problem waiting: the Disconnect Problem. Recall that the challenge is to provide a consideration “specifically to do the sort of thing that if enough of us do it, will prevent or avoid causing the serious avoidable harm.” This problem is best illustrated with Parfit’s initial case:

Drops of Water

Thousands of thirsty men are lying out in the desert, suffering from intense thirst. An equal number of people have a pint of water. They could pour these pints into a water cart. This cart will be driven to the desert, and the water shared equally among the thirsty men. Each person that donates a pint enables each wounded man to drink only slightly more water, at most a drop. Yet a single drop will not benefit even a very thirsty man. (cf. 1984, 76)

In this case, Contributory Participation as well as its upgrade 2.0 give the intuitively correct verdict: you have reason to donate your pint. After all, donating the pint adds to the amount of water in the cart, and together the group can alleviate the suffering (i.e., while no one could do so individually). However, Nefsky asks us to imagine the following variation on the case:

Single Drop

¹⁶ A case in point: Hill (1979). On Hill’s account, participation can have symbolic value: it can make a statement that you are opposed to bad outcome O* without also trying to influence it. Even so, accounts such as Hill’s have special difficulties with the upcoming Disconnect Problem (i.e., why make a statement through (not) contributing to D rather than by holding a banner?), and will not do without a solution to the latter (see Nefsky 2018).

You only add a drop to the water cart, while others add their full pint.

Adding a drop of water adds to the underlying dimension (the amount of water in the cart). But Contributory Participation, as well as Contributory Participation 2.0, seem to allow you to be “off the hook” by simply donating a single drop. After all, you now count as a participant just as much as had you donated your whole pint! But obviously, donating a drop is not the type of action such that, if all other available contributors do it, the suffering will be alleviated.

It is certainly true that for “X” (in the formulation of the Contributory Participation accounts) we can fill in “adding a drop” or “adding drop *n*.” There seems to be no *a priori* reason why we should avoid individuating the action space this way. As Nefsky admits, Contributory Participation does provide a consideration to add a drop (“drop 1”). Suppose you add drop 1. You must now ask: does drop 2 fulfill the conditions of Contributory Participation as well? The answer is “yes.” Adding drop 2 adds to underlying dimension D just as much as drop 1, and if there are enough contributions to D, D causes outcome O. The same goes for every drop in your pint. The Contributory Participation accounts do not let you off the hook so easily; they do not give you reason to participate less by gerrymandering the action space.¹⁷ To put it in a slogan: what matters is *to participate, not to become a participant*¹⁸ (Similar judgments are also plausible in analogous intrapersonal cases. If the account gives you a reason to study for 5 minutes, it gives you a reason to study for another 5 minutes, provided that both conditions still hold.)

Alternatively, one might think that this case shows that, in contrast to what Contributory Participation accounts predict, you do not have any reason to add the drop. After all, one might think, a single drop is too insignificant. To put it in Nefsky’s terms, it

¹⁷ One might suggest the same reply in defense of Kutz’s account: the latter does not provide a reason to hold up a banner *or* to donate a pint, but provides a reason to hold up a banner *and* donate a pint (i.e., if one does these things with a participatory intention). But even so, Kutz’s account would still face the following challenge: how does one explain that *merely* donating a pint is preferable to *merely* holding up a banner? There seems to be nothing in Kutz’s view that could account for this difference. Contributory Participation does not face this problem: for it does not recognize a reason to hold up a banner. In fact, we think that holding up a banner is not a way of participating, that is, *as long as it has no (direct or indirect) effect on D or O*.

¹⁸ Similar things may not be true in all participation cases. E.g., where the goal is to reach a high number of union members, what matters is being part of the group, not the amount of participation. But this is because the narrowly defined goal *reaching a high number of union members* is not something that you can participate in to a greater or lesser degree: participation is binary in such cases.

is not the type of action that, if enough people perform it, could make a difference to the outcome. But that is simply incorrect. If enough people add enough drops, the suffering of the men will be relieved.¹⁹

According to Nefsky, however, there is a more fundamental problem for Contributory Participation accounts as well. Consider:

Rocking the Cart

We are in Drops of Water. Instead of adding your pint, you rock the water cart back and forth on the mountain.

As Nefsky agrees, this case does not provide a direct counterexample to Contributory Participation. After all, Contributory Participation (as well as Contributory Participation 2.0) gets the right result: you lack a reason to rock the cart. Rocking the cart back and forth does not count as *contributing*, it does not add to the underlying dimension. But, Nefsky asks, why should a merely superfluous addition to the underlying dimension count as participating, but not such other instrumentally superfluous ways of “inserting oneself into the causal pathway” (e.g., by rocking the cart back and forth)?²⁰

But notice that the Contributory Participation accounts mirror natural and everyday usage of the term “participation” here. Suppose you send a check to Oxfam and then reverse it (you may imagine the check having cleared already). Clearly, you have not participated in the group that promoted Oxfam’s cause. Or, staying closer to the examples, suppose you donate a pint and then scoop a pint from the cart and drink it. Again, clearly, you have not participated in the group that helped alleviate suffering. If the suffering people put out a statement saying, “Many thanks to all the participants in alleviating our suffering,” they do not mean *you*. Crucially, then, there are two ways to “insert oneself into the causal pathway”: as *participant* (i.e., you add to D) and as *non-participant* (i.e., you make no net difference to D).²¹

¹⁹ Note that the account still says that *if* you only possess a drop (and not a whole pint), you have a reason to add just your drop. But this seems just right to us. Contrast Naïve Participation: the people cannot be helped if everyone donates just a drop (assuming that all drops taken together can make no difference), and so you are not a participant according to Naïve Participation if we take the relevant action description to be “adding a drop” (and not, e.g., “adding all the water you have”).

²⁰ It is not entirely clear what “inserting oneself into the causal pathway” amounts to, but for present purposes, we may adopt an intuitive understanding of what is at stake here.

²¹ In *Rocking the Cart*, D could be: the distance crossed between the cart and the thirsty citizens. One adds nothing to it by rocking it forth *and back*.

5. Helping

So far, this paper has been a defense of participation-based ideas. Next, we will show that Nefsky’s helping-based account faces very similar problems *and* can borrow our very same solution. We state her account as follows:²²

Helping

S has a helping-based reason to do X iff:

- (i) S’s Xing can be part of what causes good outcome O (or of what prevents bad outcome O*); and
- (ii) it is possible that O will fail to come about due, at least in part,²³ to a lack of Xing.

Helping faces a similar problem as the Naïve Participation account (discussed in Section 2), namely what actions fall under “X” (cf. Asker 2023)? Avoidance of (a certain level of) climate change may fail to come about due to an insufficient number of sustainable businesses, or due to an insufficient number of new forests planted, or due to an insufficient number of investments in green technologies, and so on. But suppose we take for “X,” “driving a Volkswagen Golf with license plate *p* for 15 minutes every Thanksgiving” (where *p* is a member of the set of all license plates). There seems to be no independent reason not to cut up the action space this way.

This presents a challenge for Nefsky’s account. Clearly, climate change mitigation cannot fail to come about due *solely* to that type of action – not even all actions of that type taken together could make a difference to climate harm. Now, can it fail to come about due, *at least in part*, to that type of action? Here the view faces a dilemma: either we should interpret condition (ii) as requiring that instances of Xing together can make a difference to the outcome, or not. *Ex hypothesi*, on the first horn, condition (ii) is not fulfilled: not

²² Based on the following definition of what it means to help: “Suppose your act of Xing could be part of what causes outcome Y. In this case, your act of Xing is non-superfluous and so could help to bring about Y iff, at the time at which you X, it is possible that Y will fail to come about due, at least in part, to a lack of Xing” (Nefsky 2017, 2753).

²³ “At least in part” because, for example, the thirsty men in Drops of Water may fail to be helped not only if not enough people donate their pint, but also if the cart is not driven to them fast enough.

even all acts of the type “driving a Volkswagen Golf with license plate p for 15 minutes every Thanksgiving” together could make a difference to climate harm.

Perhaps, though, we should read “at least in part” as being consistent with condition (ii) not implying that instances of Xing together could make a difference to the outcome. However, the intuitive appeal of the helping-based account seems to depend on the idea that, even if not alone, at least *together* we can make a difference. If even that much is not true, then it becomes unclear why one would have a helping-based reason to act in the first place. If even *together* we cannot make a difference, then it seems that one is simply not helping at all!²⁴

The second horn of the dilemma would also appear to imply a rather uncharitable reading of Nefsky. She seems to have in mind the first interpretation of clause (ii) when she explains: your act “is non-superfluous because, at the time at which you do it, there’s a real risk that the outcome will fail to come about as a result of not enough people acting in exactly that way” (2017, 2753). This indicates that “at least in part” is consistent with condition (ii) being read in a way such that it requires instances of Xing to potentially make a difference to the outcome.

In sum, if we take for “X,” “driving a Volkswagen Golf with license plate p for 15 minutes every Thanksgiving,” condition (ii) of Helping is not fulfilled. But gerrymandering the action space in this way is something that we can do for many everyday actions for which we think there is good moral reason to perform them: investing 10 dollars in sustainable energy on a specific date, voting in polling station such and such at time such and such.²⁵ In sum, for many everyday actions, the fulfillment of clause (ii) may well depend on the specificity of the action description.

Luckily, these problems are avoided if we update the account along the lines of Contributory Participation:

²⁴ Alternatively, one might put the point as follows. There is only reason to perform an individual act when there is a (group-level) reason to perform the *group* of acts of which the single act is a part. But on this interpretation, the Superfluity Problem appears to simply reiterate. To paraphrase Nefsky (2015, 264), “it’s not clear that there can be any reason to [perform a group of Xings], which connects appropriately to outcome Y, if [this group] cannot make any difference with respect to Y. If [the group of Xings] would be merely superfluous with respect to the outcome in question, it’s not clear that we can get any relevant story going as to why [the group has] reason to do it.”

²⁵ Perhaps similar gerrymandering is not plausible with big, one-time investments in fossil fuel, or with actions such as the taking off of a plane. A single big investment in fossil fuel may actually make a difference, and a single take-off might emit enough greenhouse gases to potentially make a difference to climate harms.

Helping 1.1

S has a helping-based reason to do X iff:

- (i) X adds to an underlying dimension D; and
- (ii) it is possible that O will come about, and it is possible that O will fail to come about due, at least in part, to insufficient contributions to D.

For example, we have a helping-based reason to reduce our emissions because (i) doing so reduces the amount of greenhouse gases in the atmosphere; and (ii) it is possible that climate harms will not be prevented due, at least in part, to insufficient such reductions. (Again, “adding” here is a technical notion: if one reduces the amount of greenhouse gases, one “adds” to the amount of reduced greenhouse gases. And people might do this with various types of actions.)

Unfortunately, the problems do not stop here. Since there is no preferred action description, it is not clear that the helping account (neither the original version nor Helping 1.1) can deal with Vending Machine and Locks.²⁶ In Vending Machine, take for “X,” “contributing money.” Conditions (i) and (ii) are fulfilled on this action description: (i) “[contributing money] can be a part of what causes O” (Helping 1.1: “[contributing money] will add to [the amount of money in the machine]”) and (ii) “it is possible that [the dispense of the granola bars] will fail to come about due, at least in part, to a lack of [contributions of money].” According to Helping and Helping 1.1, then, C should put her quarter in. But according to Nefsky herself – and this seems the correct verdict – it is *not* the case that C should put her quarter in.

Nefsky would have to argue that “contributing money” is not the relevant level of description, but that we should focus on *how much* money that they contribute. Thus, in Vending Machine, though contributing a quarter can be a part of what causes O, it is *not* possible that O will fail to come about due, at least in part, to a lack of contributing quarters. There are just not enough others around with such quarters. Even so, it is not clear why we should stick with the latter action description rather than the former. And allowing both action descriptions will not help Nefsky either: as soon as we (also) allow the action description “contributing money,” the account incorrectly recognizes a reason to add the quarter.

Luckily, the account can be fixed by taking on board Contributory Participation 2.0:

²⁶ Maria Lasonen-Aarnio voiced a similar worry for Nefsky’s view in her comments during the Small Acts Big Harms workshop in Helsinki.

Helping 2.0

S has a helping-based reason to do X iff:

- (i) X adds to an underlying dimension D; and
- (ii) there are enough potential contributors such that (a) their contributions to D together could make a difference to O, but (b) individually cannot make a difference to O; and
- (iii) it is possible that O will come about, and it is possible that O will fail to come about due, at least in part, to insufficient contributions *specified under (i) and (ii)*.

This account is the same as Contributory Participation 2.0, but with a clause – (iii) – added to it. This clause is meant to rule out “closed” cases, i.e., cases where the good outcome O is no longer possible and one can no longer help bringing it about, or where the bad outcome O* is no longer possible and one can no longer help preventing it. Notice that this clause cannot simply read “it is possible that O will come about” and “it is possible that O will fail to come about.” For, consider:

Mr. Rich Guy

There are two carts, where one cart could only be filled by a large group of people (as in Drops of Water), and the other entirely by a rich person. The collective cart is full, but it is still up in the air whether the rich person is going to donate and whether the carts will be sent off.

In Mr. Rich Guy, you cannot help anymore, and your act has no instrumental significance. It should not just be up in the air if O will come about for it to be possible to help. It should also be up in the air whether the group of which you would become a part by contributing can still make a difference to the outcome. If not, there is no helping-based reason to add your pint. Hence, there is no reason to add your pint to the cart that is already full. Helping 2.0 captures all this: in this case, the outcome cannot fail to come about due to a lack of contributions to D that *collectively can* make a difference but *not individually*. The alleviation of suffering can only fail to come about in this case due to a lack of contributions to D that can make a difference *individually* (i.e., due to the rich person).

Somewhat surprisingly, then, it turns out that helping (as understood along the lines of Helping 2.0) is related to participation in an important way: it can be seen as a *type* of participation (as understood along the lines of Contributory Participation 2.0). Indeed, Helping 2.0 *just is* Contributory Participation 2.0, but with clause (iii) added to it. Clearly,

this result is not trivial – it is not pretheoretically obvious that we should think of helping as a type of participation. Moreover, cases like Mr. Rich Guy illustrate that the concern with who is in and out of the group based on the type of contribution that they make is doubly important for helping accounts: once to determine whether S is part of a group of agents who only together can make a difference, and once to determine whether it is still possible that O will or will fail to come about at least partly due to how the members of this group act (rather than merely due to how other agents, such as Mr. Rich Guy, act).

6. Discussion

Before delving into the exact difference between helping and participating, let us consider a series of cases that shed further light on the core proposal that is common to both Contributory Participation 2.0 and Helping 2.0 and distinguishes them from the original versions (i.e., their clause (ii)).

Massive Torturer

There are 1,000 switches on a torture machine that is connected to one victim. Pulling a single switch does not make any (morally relevant) difference to the victim's pain, though after pulling too many switches the victim will be in extreme pain. One massive torturer pulls 900 switches on a torture machine, while 100 others, including you, can pull only one.²⁷

In this case, Contributory Participation 2.0 and Helping 2.0 recognize a reason not to pull. For there is a group of agents who (a) together make a difference to the victim's pain (assuming, that is, that 100 switches do raise the electric current sufficiently in this respect), but (b) individually make no such difference. Important to notice is that clause (ii) leaves it open whether there are *other* agents (e.g., the massive torturer) who *can* make a difference to D.

Massive Torturer II

One massive torturer pulls 999 switches, and you can pull or not pull the last one.

Clearly, you do not have a difference-making reason, since you make no difference to the victim's pain level. Neither do you have a helping-based reason, nor a participation-based

²⁷ Adapted from Parfit's "Harmless Torturers" (1984, 80).

reason, because you do not participate in any group (i.e., of agents who individually cannot make a difference). Clause (ii) fails.

One might think it is still problematic to pull the last switch. If so, our intuition should be explained by other, non-participation-based reasons, such as reasons based on considerations of vice or the expression of bad attitudes. Perhaps one risks the cultivation of the wrong kinds of habits (vices) when one pulls the last switch. Or, perhaps by pulling the last switch, one expresses one's indifference or disrespect towards the victims.²⁸

Massive Torturer III

One massive torturer pulls 995 switches. Five others, including you, can pull only one.

Let us assume, for a moment, that in Massive Torturer III, clause (a) fails: the five contributions together make no difference to the victim's pain. If this is so, there is no relevant group in which you participate, and no participation-based reason not to pull the switch. Arguably, however, there is a vague distinction between cases where a group can and cannot make a difference – where we read “vague” not, e.g., in an epistemicist sense, but as implying the absence of a threshold. But this is no issue for the current account. Sometimes, it is similarly vague whether an *individual's* contribution can make a difference. As a consequence, it is vague whether there is a *difference-making* reason for that individual to act in some way. Here we have the same problem: it is vague whether a *group's* contribution can make a difference. As a consequence, it is vague whether there are *participation-based* reasons for the individuals in the group. But that is exactly what we should expect, for both difference-making reasons (when it is vague whether an *individual's* contribution can make a difference) and participation-based reasons (when it is vague whether a *group* can make a difference).²⁹

²⁸ Revisit a further variant by Parfit: “One morning, only one of the torturers turns up for work. It happens to be true that, through natural causes, each of the victims is already suffering fairly severe pain. This pain is about as bad as it would be after the switches had been turned five hundred times. Knowing this fact, the Single Torturer presses the button that turns the switch once on all of the machines. The effect is the same as in the days when all the torturers act” (1984, 81). Here, the same analysis can be given as in Massive Torturer II: clause (ii) fails (there is no group of agents), and so there is no participation-based reason not to pull the switch – though, again, there may be an alternative story for why it is problematic to do this.

²⁹ To be clear, this is not a point about second-order vagueness. Suppose that the situation is as follows: at switches 1–990, the victims do not suffer maximally, at switches 991–998, it is vague whether they suffer maximally, and at switches 999–1,000, the victims suffer maximally. In this

It is important to clarify in more detail what we mean by the ability to make a difference. Consider the following case:

Certain Torture

As before, there are 1,000 switches on a torture machine that is connected to one victim. After anyone (out of a 1,000 potential contributors) who wants to pull a switch does so, Evil Gates will pull the rest of the switches. So, whether or not others participate, it is certain that the victim will suffer to the maximum extent (cf. Nefsky 2019, 9).

Condition (ii) states that “there are (or will be) enough potential contributors such that (a) their contributions to D together could make a difference to O*.” Given Evil Gates’ presence, there is a clear sense in which the group of 1,000 potential contributors cannot make a difference as to whether the outcome will come about. After all, Evil Gates will step in if any of the 1,000 potential contributors fails to pull a switch. Thus read, Contributory Participation 2.0 fails to recognize a participation-based reason not to pull in Certain Torture. But that seems wrong. The fact that Evil Gates would step in does not mean that you can participate in a group that brings about extreme suffering.

Notice, however, that something else still holds true: the group of potential contributors could have made a difference *had Evil Gates not been present*. What this suggests is that the group of potential contributors need not be able to make a difference *given a full specification of all the circumstances*. Instead, it must have the capacity to do so in some other sense. What is of relevance here is the size of the aggregate contributions. In particular, as far as the *aggregate size* of their contributions to D is concerned (that is, the aggregate size of the contributions to D by those who cannot individually make a difference to O*), their contributions to D could make a difference to O*. It does not matter, then, whether it is already settled that O* will come about.

A further problem case to consider:

The Train

The case is as in Drops of Water. The cart is still empty. However, this time, there is also a train with people with enough pints to fill the cart completely, who would all donate their pints if they arrive in time. The train, however, may or may not

setup, it is vague whether the turning of the last five switches makes a difference, without there being any second-order vagueness as to where the vagueness starts. When applying this to subjective reasons, talk of vagueness may be replaced with talk about uncertainty.

arrive in time. You are standing next to the cart, alone, and also have a pint. But you have to leave before the train arrives and before the cart is driven into the desert. You therefore have to decide before the train arrives whether or not you will add your single pint.

The problem is this: Contributory Participation 2.0 and Helping 2.0 seem to say that you have a reason to add your pint. There is, after all, a group of agents who can only collectively make a difference, and you can participate. Notice that the same goes for Nefsky's original Helping account: your adding your pint can be a part of what causes the relief of suffering, and relief of suffering may come about or fail to come about due to an (in)sufficient number of people adding their pint. But just as in Locks and Vending Machine, it does not seem plausible that you are part of the group on the train, and it does not seem that you have any reason to add your pint. Whether or not the train arrives on time, your pint is a total waste.

A number of potential replies present themselves. Reply 1: if O comes about, this is not explained by all contributions to D being enough to bring about O, but rather by the train arriving in time and the subsequent contributions by the train travelers being enough to bring about O. If it arrives, then *everyone* will simply add their pint. One way of accounting for this thought would be to read clause (ii)(a) as follows: "their contributions to D together *explain* O's coming about."³⁰

Reply 2: We need to allow for the possibility that multiple agents can make a single contribution. The problem with this reply is that it is not clear that this *is* a single contribution. To make the worry about this reply more graphic, think of this as a case in which all train travelers individually decide (not) to add their pints, but you know, because you have knowledge about this particular group, that they will all decide the same.³¹

Reply 3: Given the case description, there is no "live option" that some are going to contribute and some not: either they all contribute, or they all will not. Hence, for purposes of practical deliberation, we may (or perhaps should) treat the group of people on the train as *one big* difference-making contributor. And if it is rationally permissible (or required) to do so, then as far as Contributory Participation 2.0 and Helping 2.0 are concerned, I should conceive of my contribution as not being a part of the contribution by the group of people on the train. If so, perhaps condition (ii) – "there are enough potential contributors such that (a) their contributions to D together could make a difference to O,

³⁰ For such an explanatory clause, cf. Björnsson (2021).

³¹ This scenario need not be too outlandish. E.g., you may know that their individual decisions are conditional on a specific condition, of which you know that it is fulfilled.

but (b) individually cannot make a difference to O” – is not fulfilled, and Contributory Participation 2.0 and Helping 2.0 do not give me a reason to add my pint.

This final reply – that we should treat the group of people on the train as one big difference-making contributor – raises questions about the nature of practical rationality. But even if it is true that I should treat all the people on the train as one big difference-making contributor, it does not yet follow that Contributory Participation 2.0 and Helping 2.0 will not (erroneously) give me reason to add my pint. This will depend on what grounds the relevant reasons in The Train: can the reasons that Contributory Participation 2.0 and Helping 2.0 give rise to only be grounded by how contributions are in fact individuated or are they possibly grounded by how I should individuate them *for purposes of practical deliberation*? In other words, can the fulfillment of condition (ii) depend entirely on how I may (or should) individuate contributions for purposes of practical deliberation? This is a problem above our paygrade.

7. Participating vs. Helping

The two accounts – Contributory Participation 2.0 and Helping 2.0 – are very similar (except for one clause). What is the difference? Essentially, the helping account cares about “making progress,” namely towards the good outcome, and towards preventing the corresponding bad one. As Nefsky puts it: “even if your act cannot itself make the difference between one outcome and another, it can make non-trivial progress toward a better outcome, and this can give you reason to do it” (2017: 2764). The same does not apply to the participation account. Even when the bad outcome is already certain, and one can no longer make any progress towards the good outcome, one may still have participation-based reasons not to overdetermine the bad outcome.

This yields a central difference between the two accounts, which can be characterized in terms of “open” versus “closed” cases. According to Nefsky (2017: 2752-53), one can only help if it is still up in the air whether the good outcome O will be brought about (and the corresponding bad one avoided), and one cannot help anymore when it is already certain that O will not be realized or already certain that O will be realized. That is, there exist no helping-based reasons in closed cases, but only in open cases. In contrast, according to participation accounts, there can still be reasons not to participate *even when*

*it is already certain that the bad outcome will come about.*³² For, the idea is, one should not participate in groups that bring about collective harms.³³ To illustrate:

Certain Torture II

As before, there are 1,000 switches on a torture machine that is connected to one victim. You know that 999 torturers are going to pull their switch, and it is certain that the victim will be in excruciating pain. You still have to decide whether or not to pull yours, but you know, as always, that this is not going to make the pain any worse.³⁴

Pulling your switch is wrong. Yet the helping account fails to give you a reason not to do so. After all, the excruciating pain is already guaranteed, and there is no way your action can change that. In contrast to the helping accounts, Contributory Participation 2.0 does recognize a reason not to pull. Both conditions (i) and (ii) are fulfilled: your pulling a switch would add to the underlying dimension (of electric current) and there are enough other potential contributors such that their contributions together could make a difference to the suffering of the victim (but not individually).

Compare the following collective benefit case:

Certain Water

Just like Drops of Water, except for one big difference: Bill Gates has set up a perfectly reliable machine that will – after anyone who wants to has added their pints to the cart – fill up the cart the rest of the way with water. So, whether or not people add their pints, the cart will be totally filled and the suffering will be relieved to the maximal extent (Nefsky 2023, 54).³⁵

If Contributory Participation 2.0 applies to Certain Water, it would seem to recognize a reason to add your pint in this case as well. Both conditions (i) and (ii) are fulfilled: your

³² They may even be strongest in such situations. Generally, participation-based and helping-based reasons may come in different degrees in strength – an issue that we will not address here.

³³ Cf. Talbot (2018) and Wieland (2022). Pinkert (2015), too, offers an account of closed cases. According to him, you need not cooperate in the actual situation when you know that others do not cooperate, but you should still be willing to cooperate, i.e., cooperate in counterfactual situations where (enough) others act differently.

³⁴ Slightly adjusted from Van Oeveren, ms.

³⁵ There are real-world examples like this; cf. Budolfson and Spears (2019).

adding a pint would add to the underlying dimension, and there are enough other potential contributors such that their contributions together, but not individually, have the ability (in the sense indicated in Section 6) to make a difference to the suffering of the people in the desert. But this seems to flout our moral judgments: there seems to be no reason to add your pint in *Certain Water*.

Stated more generally, there seems to be no reason to participate in groups that bring about collective benefits when the outcome is already guaranteed to occur. We therefore suggest *Contributory Participation 2.0* as an account that applies to harm cases only. Perhaps there are participation-based reasons in collective benefit cases, but they come with more stringent conditions (and maybe such reasons are well captured by *Helping 2.0*).³⁶ These are issues for future consideration. The central difference between helping and participation accounts is that the former, but not the latter, require that it is possible that the bad outcome will not be brought about.

Consider a final case:³⁷

Many Little Locks

There is a wide canal with a large region in which the water level is controlled by a great many locks. Right now, the water level is low, around knee deep. But when enough locks are open, the water level will rise above head height, but the amount of water released by any particular lock is minute. Austin Powers is currently, while the water is low, walking in the middle of the canal looking for a key that has fallen in and that is crucial for his mission. He does not know how to swim. Dr. Evil has anticipated this situation, and he has had his henchman, Mustafa, knock out the person who was manning the locks. Now Dr. Evil instructs droves of their minions to open the locks, in an effort to drown Powers. You are on the land nearby, and you see what is going on. You decide to slip into the water to swim to the middle of the canal where Powers is and try to rescue him.

Of particular interest to us is a version of the case in which, if you decide to go into the water, success of the rescue mission is guaranteed. If it is, it would seem that there is no reason not to wade in, not even *pro tanto*. Yet in this case, both conditions of *Contributory Participation 2.0* are fulfilled: (i) is fulfilled because you add to the underlying dimension (the water level in this case), and (ii) is fulfilled because there are enough potential

³⁶ This restriction of *Contributory Participation 2.0* also avoids the wrong result in *The Train*.

³⁷ Thanks to an anonymous reviewer for presenting us with this case.

contributors such that (as far as their aggregate size is concerned) their contributions together could make a difference to O^* , but not individually.

One plausible suggestion is that the bad outcome O^* must at least be possible in order for there to be participation-based reasons not to add to D. The basic idea of participation accounts is that you should not participate in groups that bring about harm, but if the group cannot bring about harm (because the harm is circumvented), then you cannot join this group, and there are no participation-based reasons to act in one way or another. Since we assume that O^* is no longer possible the moment that you wade into the water to save Powers – after all, success of the mission is guaranteed – there are no participation-based reasons not to wade into the water.³⁸ Thus, whereas helping accounts require that it be possible both that the good outcome come about and that it fail to come about, Contributory Participation 2.0 would possibly only require that it be possible that the bad outcome come about. Again, this is an issue for future consideration.

8. Conclusion

It is time to take stock. Parfit wrote: “In my appeal to the effects of what groups together do, I did not fully explain whom we should count as members of these groups. This gap needs to be filled” (1986, 848). This paper’s ambition was to do exactly that. Firstly, we provided a more sophisticated account of “participation” – Contributory Participation 2.0 – premised on the idea that even when you cannot make a difference on your own, at the very least, the group that you join should still be able to do so. We argued that this requirement makes the account resistant to the recurring Superfluity and Disconnect problems. In fact, our account can be inserted in *further* accounts, such as Helping 2.0. Indeed, helping could be seen as a special type of participation: *participation in open cases*. Especially, and in contrast to what Nefsky assumed, our account does not “smuggle in” the assumption that one’s action does matter instrumentally, and so can equally be inserted in non-instrumental solutions to the collective action problem (such as Kutz’s account based on participatory intentions). To be sure, we did not show that our account is better than

³⁸ If the case is *not* closed, and there is a chance that you will fail in your mission, we suggest reading the case as Locks (Section 2), namely when the minions are no real agents, but fully obedient to Dr. Evil, and are merely an extension of him; or as Hot Day (Section 1), when the minions act on their own behalf and decide for themselves if they wish to cooperate with Dr. Evil’s plan; or as The Train (Section 6), when it is Mustafa’s task to communicate Dr. Evil’s plans to the minions, and it is up in the air if he can do that successfully.

alternative accounts³⁹ – an issue that merits further attention – yet non-instrumental solutions deserve our attention as well, and our aim was to lay the ground for them.

References

- Andreou, C. 2014. The Good, the Bad, and the Trivial. *Philosophical Studies* 169: 209–25.
- Asker, A. S. 2023. The Problem of Collective Impact: Why Helping Doesn't Do the Trick. *Philosophical Studies* 180: 2377–97.
- Barry, C. and Øverland, G. 2016. *Responding to Global Poverty. Harm, Responsibility, and Agency*. Cambridge: CUP.
- Björnsson, G. 2021. On Individual and Shared Obligations: In Defense of the Activist's Perspective. In Budolfson, M. et al., *Philosophy and Climate Change*, ch. 11. Oxford: OUP.
- Budolfson, M. B. and Spears, D. 2019. The Hidden Zero Problem. In Hilary, G. et al., *Effective Altruism. Philosophical Issues*, ch. 12. Oxford: OUP.
- Driver, J. 2015. Individual Consumption and Moral Complicity. In Bramble, B. et al., *The Moral Complexities of Eating Meat*, ch. 4. Oxford: OUP.
- Fanciullo, J. 2020. What is the Point of Helping? *Philosophical Studies* 177: 1487–1500.
- Glover, J. 1975. It Makes No Difference Whether or Not I Do It. *Proceedings of the Aristotelian Society* 49: 171–90.
- Gunnemyr, M. 2022. Making a Vague Difference: Kagan, Nefsky and the Sorites Paradox. *Inquiry*.
- Gunnemyr, M. and Touborg, C. T. 2023. Reasons for Action: Making a Difference to the Security of Outcomes. *Philosophical Studies* 180: 333–62.
- Hedden, B. 2020. Consequentialism and Collective Action. *Ethics* 130: 530–54.
- Hill, T. E. 1979. Symbolic Protest and Calculated Silence. *Philosophy and Public Affairs* 9: 83–102.
- Kagan, S. 2011. Do I Make a Difference? *Philosophy and Public Affairs* 39: 105–41.
- Kutz, C. 2000. *Complicity. Ethics and Law for a Collective Age*. Cambridge: CUP.
- Lee, S. 2022. Collective Actions, Individual Reasons, and the Metaphysics of Consequence. *Ethics* 133: 72–105.

³⁹ Such as the causation-based accounts by Lee (2022) and Gunnemyr and Touborg (2023). Interestingly, there could be causation-based reasons not to go into the water in Locks to save Powers, while, as we argued, there are no such participation-based reasons.

- Lichtenberg, J. 2010. Negative Duties, Positive Duties, and the New Harms. *Ethics* 120: 557–78.
- Nefsky, J. 2015. Fairness, Participation, and the Real Problem of Collective Harm. *Oxford Studies in Normative Ethics* 5: 245–71.
- Nefsky, J. 2017. How You Can Help, Without Making a Difference. *Philosophical Studies* 174: 2743–67.
- Nefsky, J. 2018. Consumer Choice and Collective Impact. In Barnhill, A. et al., *The Oxford Handbook of Food Ethics*, ch. 12. Oxford: OUP.
- Nefsky, J. 2019. Collective Harm and the Inefficacy Problem. *Philosophy Compass* 14: 1–17.
- Nefsky, J. 2023. Participation, Collective Impact, and Your Instrumental Significance. *Journal of Practical Ethics* 11: 38–58.
- Parfit, D. 1984. *Reasons and Persons*. Oxford: OUP.
- Parfit, D. 1986. Comments. *Ethics* 96: 832–72.
- Pinkert, F. 2016. What If I Cannot Make a Difference (and Know It). *Ethics* 125: 971–98.
- Salomone-Sehr, J. 2022. Cooperation: With or Without Shared Intentions. *Ethics* 132: 414–44.
- Talbot, B. 2018. Collective Action Problems and Conflicting Obligations. *Philosophical Studies* 175: 2239–61.
- Van Oeveren, ms. Causation, Omission, and the Problem of Collective Action.
- Wieland, J. W. and Van Oeveren, R. 2020. Participation and Superfluity. *Journal of Moral Philosophy* 17: 163–87.
- Wieland, J. W. 2022. Participation and Degrees. *Utilitas* 34: 39–56.

Acknowledgements

The main idea in this paper (CP2) is Rutger’s, the arguments are structured by both authors, and Rutger took the initiative to write most of it. We would like to thank, first of all, Julia Nefsky for her work on this topic and the challenges she developed for our approach. Thanks also to: Nathan Engel-Hawbecker, Mattias Gunnemyr, Benjamin Mullins, a reviewer of the journal, and audiences at Reading University (British Society for Ethical Theory), Vrije Universiteit Amsterdam (Debunking Inefficacy Reasoning), Bowling Green State University (Applied Philosophy Graduate Workshop), OZSW/VU Amsterdam (Corona Ethics), University of Helsinki (Small Acts Big Harms).