

PARTICIPATION, COLLECTIVE IMPACT, AND YOUR INSTRUMENTAL SIGNIFICANCE


JULIA NEFSKY

1. Introduction

In past work over a series of papers, I have discussed what I have called ‘the problem of collective impact’, or sometimes following others ‘the inefficacy problem’ (Nefsky 2011, 2015, 2017, 2018, 2019, 2021). This is a problem that arises whenever there is some sort of choice that is such that (1) if enough people make that choice, serious harm could be avoided or reduced, and yet (2) it does not seem that one’s doing so, on any given occasion, will itself make a difference. We encounter this frequently in everyday decision-making. Consider, for example, how our collective consumer choices affect poverty and the exploitation of people in various parts of the world; yet, for many products, it is doubtful that one purchase more or less will make a difference. Things are not, it seems, going to go differently at any company, or for any worker, depending on whether I buy a given mass-produced T-shirt or cell phone. Similarly, while it’s true that if enough people refrain from driving, flying, consuming beef, etc., climate change and resulting harms will not be as severe, when we look at one such choice individually, it does not seem to make a difference. Will the harms of climate change be any less severe if I take public transit to work today instead of my car? It’s hard to believe that they would. And the problem is that, if this is right, it’s unclear what point there is in doing so. As concerned as I might be about climate change, or about the exploitative conditions of workers in developing countries, if my forgoing the car trip, or the purchase of the T-shirt or the iPhone, won’t make a difference for the better, what reason is there to do so?

It is important to be clear about exactly which problem I am focused on, because there are other related issues that can sometimes be conflated with it. The problem I am interested in here is specifically about how, in these sorts of contexts, bad outcomes that are at risk of occurring in the future can be prevented

Contact: Julia Nefsky <julia.nefsky@utoronto.ca>

 <https://orcid.org/0000-0000-0000-0000>

or avoided if enough people make certain sorts of choices rather than others. It seems like in that sort of context, morality should be able to tell each of us that there is good reason to make those choices. At least in the absence of strong countervailing considerations, morality should be able to steer us away from together causing, or failing to prevent, serious avoidable harm. The problem is one of explaining how morality can do that, when it appears that an individual act of the relevant sort won't itself make a difference.

It will be helpful to have an example to focus on. There are many important real-world cases we could use. But because my interest here is in this moral-philosophical problem wherever it arises, as opposed to one particular instantiation of it, I think it is best to use an imaginary case. Let's take the following example (which I've also discussed elsewhere) from Parfit (1984):

Drops of Water: Ten thousand men are in the desert, suffering from intensely painful thirst. We are ten thousand people near the desert, and each of us has a pint of water. We cannot go into the desert ourselves, but we can pour our pints into a water cart. The cart will be driven into the desert, and any water in it will be evenly distributed amongst the men. If most or all of us add our pints, this will relieve the men's suffering. But while that is so, any individual addition of a pint does not seem to make a difference. Adding your pint only allows each man to drink an extra ten-thousandth of a pint of water. This is no more than a single drop, and one drop more or less seems to be too small an amount to make a difference to how they feel. As a result, for any individual, it's unclear what point there is in donating one's pint of water. It's true that these people are suffering and a large amount of water in this cart would help them. But if adding my pint of water won't make any difference to how well or poorly this goes, then what point is there in adding it?¹

Consider, for a moment—from a simple, commonsense perspective—why it is that the 'it won't make a difference' claim seems to undermine there being a reason to add one's pint. If adding your water won't make a difference, this means that things will go the same for the suffering men whether or not you do so. The extent to which their suffering is relieved will be the same, give or take your water in the cart. This seems to mean that adding it would be a useless, merely superfluous thing to do, as far as the goal of relieving their suffering is concerned. That is why it does not seem there is a point in doing so. Why give this away if it won't actually help the situation at all?

1. Slightly adjusted from Parfit (1984: 76), following my presentation in other work.

With this in mind, we can divide solutions to the problem into two categories: ‘instrumental’ and ‘non-instrumental’.² ‘Instrumental’ approaches attempt to show that it is not actually true that the individual act is instrumentally merely superfluous. They aim to show, for instance, that adding your pint of water actually does, or could, help the situation. Your doing so could make real progress toward the goal of relieving the men’s suffering; it could have an influence on the outcome.³ The common way people have tried to do this is to argue against the ‘it won’t make a difference’ claim: to try to show that it is a mistake to say that things will go the same in the relevant respects whether or not you act in the relevant way.⁴ There is also (as I have tried to show elsewhere) a different way of taking an instrumental approach: one could argue that playing a nonsuperfluous, instrumental part in determining what happens can come apart from individually making a difference. The general idea of this second strategy is to try to show that in collective impact contexts, an act can make progress toward a change in outcome even if it will not by itself make the difference between one outcome and another (Nefsky 2017).⁵ Both of these strategies attempt to show that, contrary to initial appearances, an act of the relevant sort does matter for reasons having to do with the instrumental role it could play—with its potential to have an influence on the outcome.⁶

‘Non-instrumental’ approaches, on the other hand, do not try to challenge the apparent instrumental superfluity of the individual act. Instead, they point to other sorts of reasons for action—reasons that are not supposed to be about trying to influence the outcome. Some, for instance, appeal to expressive reasons,

2. I divide up the landscape of views differently in other work. This is just the division that makes sense for my purposes here.

3. As you can see, I give various different characterizations of the notion of instrumental nonsuperfluity, and so of the instrumental-versus-non-instrumental-reasons distinction. This is intentional. None of the characterizations is meant to function as *the* definition, in the sense of providing a single authoritative statement of the notion. Instead, these different characterizations are meant to help the reader glom onto what I take to be a simple, ordinary notion—one of having some instrumental usefulness with respect to a certain goal or outcome. Note that it would be problematic for me to try to provide necessary and sufficient conditions for instrumental usefulness, or nonsuperfluity, in this essay, because I think it is—importantly—an open question what exactly it takes for a contribution to be instrumentally nonsuperfluous. In Nefsky (2017), I argue that, while we often assume that difference-making potential is required, this is not correct. But I do not want to take a stand on this in the present essay.

4. This is typically done by arguing that there is always a chance of individually making a difference. Examples include Norcross 2004, Kagan 2011, Singer 1980, and Broome 2019.

5. There I develop a particular account of how this works, but there could be other approaches (with different accounts) that take this same general strategy.

6. These two strategies are not mutually exclusive. One could pursue or make use of both of them. My focus in past work has been on bringing out the second, and questioning the extent to which we can rely on or be fully satisfied with the first, but I think that both could be important in understanding many real-world cases.

arguing that what an act expresses matters morally for its own sake (e.g., Lane 2018; Kutz 2000: ch. 6; Benn 1979; Brennan and Lomasky 1989, 1993: 186–89). There have been various attempts at non-instrumental solutions, including—among others—appeals to the moral significance of participation in a group and appeals to collective obligations (e.g., Baatz 2014; McPherson 2016; Brennan 2009; Cullity 2000; Jamieson 2007; Sandler 2010; Strang 1960).

In some of my past work I argue that non-instrumental approaches do not work by themselves. A necessary and central part of addressing the problem of collective impact, I argue, is showing that individual acts of the sort in question are not instrumentally merely superfluous. I do not deny that there can be expressive reasons, participation-based reasons, or any number of other non-instrumental reasons for action. What I argue is that such reasons cannot by themselves satisfactorily address the problem. A core instrumental solution is needed; with it, these other considerations can build on top. In their paper ‘Participation and Superfluity’, however, Wieland and van Oeveren attempt to give a non-instrumental solution that avoids my argument (Wieland and van Oeveren 2020). In the present paper, I will take a look at their proposal. I will argue that it does not actually succeed in avoiding my point; the same essential issues plague their view too. This is worth discussing for its own sake; their proposal is an interesting, initially appealing one, worthy of consideration. But my aim in examining their view is broader and more constructive: to further support and illuminate my thesis about the need for a core instrumental solution to the problem of collective impact. I will also, in the final section of the paper, identify what I think is the main source of resistance to this thesis: a mistaken conflation of the problem of collective impact with a different problem. The hope is that with greater clarity on these points, we will be able to build a better, more accurate understanding of individual morality in collective impact contexts.

2. The Superfluity and Disconnect Problems

From a bird’s-eye perspective, what my past discussions of non-instrumental approaches have argued is that they face a dilemma: either the reason for action identified implicitly relies on there being a solution of an instrumental sort *or* the reason identified is not really addressing what is at issue.⁷

Let me give an example of how this goes. Since Wieland and van Oeveren give a participation-based proposal, it will be useful to look at a different such view as our example. According to participation views, even though your

7. These discussions occur across Nefsky (2015; 2018) and also in a smaller way in Nefsky (2019).

addition of water won't make a difference to the suffering men, you should add it because this would (or could) make you a participant in a group that together prevents their continued suffering. Such a view needs an account of what it takes to be a participant in the group, and this needs to be a morally significant, reason-giving kind of participation. Chris Kutz offers such an account. On his view, one is a participant in a collective action if one acts with a 'participatory intention', an intention to do your part in a joint project (Kutz 2000: ch. 3). So, adding your pint makes you a participant in the group that aids the suffering men because in doing so you are intentionally doing your part in the joint project of alleviating their suffering. By intentionally doing your part, you are an 'inclusive' author of the beneficial outcome; you are one of those who can say, 'We did it'. This, the claims goes, is something you have moral reason to do, even if your action will not itself make a difference to the outcome (Kutz 2000: ch. 4).⁸

But here's the problem: suppose I am debating whether to add my pint to the cart, and I'm wondering, 'Why do this if it won't make any difference?' If you reply, 'You should do it because then you would be intentionally doing your part in a joint project of relieving suffering', this doesn't make sense. My worry is that it seems to be an entirely useless thing to do as far as that project is concerned. This is precisely why I can't see why there is a point in adding it. Doing so does not seem to actually help the project at all. So, why should I think of it as 'doing my part' in that shared project? The point, in other words, is that the reason identified (that you would be intentionally *doing your part* in this beneficial joint project) is only going to seem to apply if I am thinking that adding my pint to the cart would help the project, and not thinking of it as a merely superfluous, useless thing to do. This account, therefore, presupposes a prior solution to the 'it won't make a difference' challenge—a solution of the instrumental sort—rather than providing an independent answer to it. I call this the *superfluity problem*: the reason for action identified only works if we assume that the act is *not* instrumentally merely superfluous. Thus, rather than providing its own solution, it's an account of a reason that we have that presupposes a prior solution (Nefsky 2015).

A proponent of the view, however, might say I am making a mistake in assuming an instrumental reading of 'doing one's part' in a joint project. Perhaps, as Kutz sometimes suggests, one can count as doing one's part in a non-instrumental sense; for instance, in an expressive sense. Adding my water to the cart could be understood as intentionally doing my part in virtue of the

8. Note that there is more to Kutz's picture than what I am including here. This is not meant to be a full representation of Kutz's views on the topic, but rather an example that I'm using to illustrate the dilemma for non-instrumental approaches.

concern or solidarity it expresses, even if is entirely superfluous, instrumentally speaking.

This would avoid the superfluity problem, but only at the cost of falling on the second problem, which we can call the *disconnect problem*. If I can count as a participant in the group merely by expressing concern or solidarity, then I can respond to the participation-based reason by holding up a banner with a concerned or supportive message (instead of giving my pint). And if we are saying that it is only via what it expresses that adding my pint counts as participatory, then holding up the banner seems just as good (or better) a way of participating, and thus just as good a way of responding to the participation-based reason in question. This proposal, therefore, does not address what is at issue. The problem of collective impact concerns why each of us has moral reason specifically to do the sort of thing that if enough of us do it, this will prevent or avoid causing the serious avoidable harm. A reason that can be responded to just as well by doing something else instead cannot satisfactorily address this (Nefsky 2015).

Looking at various non-instrumental approaches, I have argued that they face one or the other of these two problems: the superfluity problem or (what I am now calling) the disconnect problem.⁹ How this goes differs for different views. This is especially true for the disconnect problem; there are various ways a proposal can disconnect from the issue. But while there is considerable variation in the details, the overall point is the same: we cannot satisfactorily address what is at issue by embracing the supposed instrumental superfluity of the individual act and turning our attention to other sorts of moral considerations. We need, instead, to understand why adding your pint to the cart could help the cause of relieving the suffering. So, we should all be invested in the pursuit of instrumental approaches to the problem. Very importantly, this does *not* mean that the concepts and moral considerations invoked by non-instrumental proposals (e.g., participation, complicity, fairness, collective obligation, etc.) won't play important roles in a full account of individual morality in collective impact contexts. The point is that a crucial step in developing such an account is challenging the idea that the individual choice is instrumentally irrelevant.

3. Simple Participation

Wieland and van Oeveren agree with me that various non-instrumental accounts face one or the other of these problems.¹⁰ But they propose a new version of a

9. Including appeals to collective obligation and fairness, Kantian principles, virtue and vice, expressive reasons, and the idea that it is wrong to benefit from wrongdoing.

10. Though they present the second problem more narrowly than I conceive of it. What they state as the second problem is, in my view, just one way the problem can instantiate itself.

participation view that they think avoids them. Their proposal makes use of a distinction from Shelly Kagan between the ‘underlying dimension’ and the ‘morally significant outcome of concern’. The underlying dimension is the stuff in the world that the individual acts contribute to which, when there are enough such contributions, causes changes with respect to the morally significant dimension that we are concerned about. For instance, in the case of climate change, we can say that the underlying dimension is greenhouse gas quantities in the atmosphere, and the morally significant dimension is harm caused by climate change. In *Drops of Water*, the underlying dimension is water in the cart, and the morally significant dimension is the men in the desert’s suffering. When we say that adding your pint ‘won’t make a difference’ what we mean is that it won’t make a difference along that specific morally significant dimension. We don’t mean that it doesn’t make a difference to anything at all. And one thing it does or might make a difference to is the underlying dimension. It is just that this change seems to be too small to make a difference along the morally significant dimension in question.

Wieland and van Oeveren’s proposal is that in virtue of contributing to the relevant underlying dimension, you count as a participant in the group that together makes a difference along the morally significant dimension, and that this kind of participation gives you reason for action. Unlike Kutz’s account, one need not have any specific intentions (or other specific mental states) to count as a participant; mere contribution to the relevant underlying dimension is sufficient for participation. Let’s call this view *Simple Participation*. Simple Participation has two parts:

1. **The moral principle:** ‘S has reason to do X if by doing X S participates in a group that benefits others, and S has reason to refrain from doing X if by doing X S participates in a group that harms others’ (Wieland and van Oeveren 2020: 172).
2. **The account of participation:** S participates in a group that causes O if S does an act X ‘which is such that: (i) X adds to an underlying dimension D, and (ii) because enough others add to D, D causes O’ (Wieland and van Oeveren 2020: 177).

What if you are not sure whether by adding your pint you will be part of a group that benefits others, because you are not sure whether enough others will add theirs too? Wieland and van Oeveren certainly want to say that you still have reason to add your pint, but their account does not seem to say so. They address this issue by saying that when they talk about ‘participating in a group’ they mean to include ‘possible groups’. So even if you don’t know whether enough others will contribute, by adding your pint you are at least part of a possible

group that benefits others. The reason for action, they say, is there whether the group is actual or merely possible (Wieland and van Oeveren 2020: 181).

A more natural way of handling this, I think, would be to adjust the moral principle to say that you have reason to do something if doing it *might* make you a participant in a group that benefits others. That is, I think they should appeal to the *possibility* of being part of a group that benefits, instead of talking about having reason to participate in a ‘merely possible group’. Perhaps this is just a difference in phrasing rather than content. But there is at least one substantive difference between what I am suggesting and what they say. On their view you have the participation-based reason even if you are sure there won’t be enough others in the group, and so even if the group is merely counterfactually possible. So, you have the reason to add your pint to the cart even if you know no one else will contribute, and thus even if there is no possibility of being part of a group that actually benefits the suffering people. This cannot be right. It implies that I have moral reason not to walk down a quiet street in my quiet neighborhood because this would make me part of the merely counterfactually possible group that is such that because everyone in the city walked down this street at the same time there was a stampede and people were crushed. ‘Participation’ in this merely counterfactually possible group surely does not give me any moral reason not to walk down the street.

This problem can be easily remedied by excluding merely counterfactually possible groups. And then the proposal could be put just as well in the terms I specified: that there is reason to act (or to refrain) if doing so might make you a participant in a group that benefits (or harms).¹¹ With that revision, this is a nice proposal worthy of careful consideration. There is something very intuitive to the idea that simply by contributing to the relevant underlying dimension, you are, or might be, taking part in a group that together brings about the outcome of concern, and that this is all we need to address the problem. And Wieland and van Oeveren suggest that it avoids both of the issues that I raise for other non-instrumental approaches.

But does it avoid both of the issues? I agree that it avoids the superfluity problem. All that is needed for Simple Participation’s consideration to apply is that you have added to the underlying dimension. So, its application does not depend on an implicit assumption that you do in fact do something useful toward bringing about a better outcome. But I am going to argue that it faces a very similar problem, which we can call *superfluity problem 2.0*: the consideration it identifies is not actually a plausible candidate for a morally significant reason for action, and if it seems like a moral-reason-giving consideration this

11. In a new paper, Wieland (2022) admits that this is implausible and amends his view accordingly to exclude merely counterfactually possible groups.

is because we are implicitly assuming that a contribution along the underlying dimension is not instrumentally superfluous. So, just as with views that face the original superfluity problem, this view relies for its plausibility on implicitly presupposing a prior solution of an instrumental sort. I am also going to argue that, even if I am wrong about that, the view still has trouble with the disconnect problem. So, either way, the proposal does not work. And the general point I was trying to make about the need for and centrality of an instrumental solution remains.

4. Superfluity Problem 2.0

To see superfluity problem 2.0, what we need is an example in which we can't be smuggling in an implicit assumption that your contribution matters instrumentally. That is, we need an example of an action that satisfies their account of participation but that is uncontroversially superfluous. In collective impact contexts our intuitions about whether an act is useful or merely superfluous are muddy. So, to get a clear test case, we need one that does not have the collective impact structure.

Here is an example:

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 haven't eaten in days. Luckily there is a vending machine nearby, selling granola bars for four dollars apiece. 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 ten-dollar bill and press the button, two granola bars will come out, just as if you had put in eight dollars.) The two starving people do not have any money on them. But A has a five-dollar bill, B has a ten-dollar bill, and C has only a quarter. There is no one else around.¹² Assuming no other complicating factors, A and B should put in their money: each of their contributions would make a difference to the number of granola bars the hikers would get. Either contribution alone would provide some food, and together they will give the hikers three granola bars.

12. Assume that the people in the example know that this is how the machine works.

Should C put their quarter in before the button is pressed? No matter what A and B do, C's quarter won't make any difference at all to the quantity of food. And, as long as we don't make any additional stipulations beyond what is in the example, there is no other sense in which adding the quarter would progress the cause of helping these hikers. It would be an entirely useless thing to do, as far as helping the hikers is concerned.

According to Simple Participation, though, there is moral reason to add the quarter before the button is pressed. The reason is that by adding the quarter, C would add to the amount of money (the underlying dimension), and this would make C count as a participant in the group. If both A and B put in their money, then if C puts in their quarter, the input will be \$15.25, rather than just fifteen dollars. True, this doesn't do any good at all. It is an entirely useless change along the underlying dimension. But it does add to the underlying dimension nonetheless, and this, Wieland and van Oeveren claim, would make C count as a participant in the group that benefits the starving hikers. This, they say, is a moral reason to add the quarter.

This is implausible. This does not seem to be a consideration that speaks morally in favor of adding the quarter. Perhaps someone might like to count themselves a participant in this way, but this does not seem like something morality would tell them to do.

To be clear, I don't want to rule out that there could be some other moral reason for C to add the quarter; there could be, depending on further details about the case. What is under consideration here is whether the simple fact that contributing to the underlying dimension and thereby, in that sense, making oneself a participant in the group is itself a morally significant reason to do it. The answer to that seems to me to be a clear 'no'.¹³

13. An anonymous reviewer asked why the underlying dimension needs to be money in the machine rather than food. A and B make a difference to the amount of food the starving hikers receive, whereas C does not. Why not say, then, that the underlying dimension is food? If we say that then donating a quarter would not add to the underlying dimension, and so Simple Participation would not have to say there is this reason to do so. But this does not work as a way out of the issue. Money in the machine (prior to the button being pressed) fits the notion of underlying dimension; there is no non-question-begging reason to say otherwise. And if Simple Participation is to have any chance at answering the problem of collective impact, it needs to be able to count stuff 'at that level' as the underlying dimension. In many collective impact contexts, there is something that people directly contribute to that affects some further thing in the world, which impacts the morally significant dimension. For instance, people's greenhouse gas emissions affect global temperatures, which has an impact on human and animal suffering (and various other morally significant outcomes). Simple Participation needs to count contribution to the first 'layer' of this as contribution to the underlying dimension; that is, contribution to greenhouse gas quantities needs to be enough. If they say, instead, that only a contribution that makes a difference to global temperatures counts as contributing to the underlying dimension, this clearly won't be useful in solving the problem. Finally, note that we can easily come up with a collective impact case in which C is directly manipulating the quantity of food but in an amount too small to make a difference to

If you're not sure, it might help to consider the following: Imagine telling some other person who has no money, D, that he should ask if he can quickly hold the granola bars before passing them to the hungry hikers so that he can count as a participant in what A and B are doing. It is true that Simple Participation does not make holding the granola bars count as participating. But why not? There doesn't seem to be any morally significant difference between D holding the granola bars and C adding the unhelpful quarter. Both of these insert oneself into the causal pathway, but in an entirely superfluous, unhelpful way. It seems arbitrary to count the superfluous addition of a quarter, and not the superfluous holding of the granola bars, as participatory. And it seems clear that there is no moral reason coming from the sense in which doing so would make one a participant to uselessly take a turn holding the granola bars. Perhaps one might *like* to be a participant in this way, but this is not something morality would tell us to do.¹⁴

To make the point even clearer, let's consider a different example—this time one in which the group would be causing harm rather than preventing it. Again here, what we need is a case in which—unlike in collective impact contexts—our intuitions about whether the act in question is instrumentally superfluous are not muddy.

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 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.

the relief of the hunger (e.g., A and B each give a cup of rice, while C can only give a hundredth of a grain of rice). There is no denying here that C is making a difference with respect to the same underlying dimension as A and B (food), and yet C's contribution would be useless.

14. If we say there is another reason to hold the granola bars (e.g., an expressive reason), this is a different proposal. It is not relevant to assessing whether or not Simple Participation is identifying a morally significant reason-giving consideration. I discuss the expressive proposal in Nefsky (2018), and below I will discuss trying to combine Simple Participation with the expressive proposal.

Notice that in going into the water, you are adding to the underlying dimension (the water level) that is such that *because* enough others are contributing to it, a harmful outcome will or might occur. Adding your body mass to this part of the canal technically increases the water level, but only in a tiny, imperceptible way. We can stipulate that this miniscule, imperceptible increase in water level could not possibly make any difference at all for the worse. It has no possibility of having any effect (perceptible or otherwise) on how badly things go for Powers. It is an increase in the water level that is totally instrumentally irrelevant, as far as what Dr. Evil and Mustafa are doing. It couldn't help their project in any way.

According to Simple Participation, though, because your going in the water does in fact contribute, albeit totally insignificantly, to the underlying dimension, this would make you a participant in what Dr. Evil and Mustafa are doing. It makes you a participant in this group that is causing, or might cause, serious harm. This gives you moral reason not to do it. Wieland and van Oeveren can say that this reason is outweighed by the reason you have to go in—namely, that there is a chance that you will save Powers's life. But even if it can be outweighed, Simple Participation does say that the fact that you will be adding to the underlying dimension makes you a participant in the group of people harming Powers, and that this is a morally significant consideration against going in.

This is surely not right. First, 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. They can call this 'participation' if they would like, but this example highlights that this is a fully stipulative use of the term, rather than capturing a commonsense notion. The reason it seemed like a more ordinary notion of participation in Vending Machine was because of the intentions we were imagining C would have if they did add their quarter. C would be doing so—we were likely imagining—with the intention of counting as a participant. But their account of participation was not supposed to depend on the agent having such intentions; they want it to apply in cases like buying products whose mass consumption has harmful effects, even if the agent has no participatory intentions when they do so. In Locks, where there is no temptation to think that the agent has a participatory intention, there is nothing intuitive or commonsense about the claim that adding to the underlying dimension in an instrumentally irrelevant way is enough to make one a participant in the group.

The question then is whether, even though it is not an ordinary notion of participation, it is picking out a morally significant form of involvement in a group. Is this sort of involvement in what Dr. Evil and Mustafa are doing—adding to the same underlying dimension but only in an entirely instrumentally irrelevant way—a kind of involvement that gives you moral reason *not* to go into

the water? Surely not. If it was possible that entering the water could contribute instrumentally—that it could make things worse or in some other way push things further along toward a worse outcome—this would be different. I can certainly see saying then that, regardless of your intentions, entering the water counts as a morally significant form of involvement in what they are doing and gives you at least some overridable reason not to do so. But insofar as the increase in water level that occurs as a result of your entry *won't* have any influence at all on the outcome, and *won't* help progress Dr. Evil and Mustafa's cause in any other way, there is no plausibility to the claim that this is a morally significant form of involvement in their harmful activities.

My conclusion so far is that the consideration Simple Participation identifies is not in itself a moral reason to do something. Why did it look like a moral reason (if it did) when we first heard it as a proposed solution to the problem of collective impact? It looked that way because in collective impact contexts, it is easy to smuggle in an intuition that a contribution along the underlying dimension *does* matter instrumentally. I think that intuition is correct, and I argue for this elsewhere (Nefsky 2017). But regardless of whether or not it's correct, I think we often have it. While sometimes we are pulled by the 'it won't make a difference' thought to think that refraining from driving, or buying another fast-fashion T-shirt, etc., would be useless, much of the time many of us think of these choices as mattering instrumentally, as progressing things in a certain direction. This intuition sneaks its way in when we are evaluating proposals that are not supposed to depend on its truth. If we do smuggle it in, then, of course, it seems like a morally significant way of participating. Contributing in a way that makes progress toward a better outcome is a morally significant thing to do.

Before moving on, I want to consider another possibility. In a recent follow-up paper, Wieland suggests a different way of thinking about the role of participation in a solution to the problem of collective impact. Wieland's claim there is not that participation in their sense is itself reason-giving, but rather that there are reasons having to do with what you would be *expressing* about yourself to not participate in harmful groups and to participate in beneficial groups (Wieland 2022). This is a version of a different response to the problem of collective impact: the appeal to expressive reasons. And it faces the same issue that I have raised in earlier work with such an approach (Nefsky 2018). Put in terms of Wieland's account, the problem is that the fact that one would be a participant in Simple Participation's sense is simply not sufficient for counting as expressing the relevant attitudes. Take Locks again. You could know that you are contributing to the underlying dimension (the water level) by going into the canal, and so you could know that you are participating in what Dr. Evil and Mustafa are doing in Simple Participation's sense of the

term. But it is simply not true that you therefore are expressing a lack of concern for Powers, or displaying a willingness to help in Dr. Evil's project, in virtue of doing so. Your action expresses no such thing, at least insofar as you also know that your doing so could not possibly help Dr. Evil's project in any way at all. Similarly, in Vending Machine, if C chooses not to add the quarter, we cannot say that C is thereby expressing a lack of concern for the hungry hikers. If C is aware that her quarter won't help at all, it expresses no such bad attitude to choose not to add it. (If she does express a lack of concern, this will be in virtue of other things she says or does.) When we move back to collective impact contexts, driving a gas-guzzling SUV might indeed often express a lack of concern for climate change. But if so, this is *not* simply in virtue of the fact that it counts as 'participation' in Simple Participation's sense of the term. As the case of Locks and Vending Machine demonstrate, something else or further must be doing the work. I argue elsewhere that when we look at what could explain why acts like driving a gas-guzzling SUV express the relevant attitudes (if they do), we find that expressive reasons cannot themselves address what is at issue in the problem of collective impact (see Nefsky 2018: 279–83).¹⁵

5. The Disconnect Problem

I have argued that Simple Participation faces superfluity problem 2.0: it is not actually identifying a morally significant reason for action, and if it looks like it does, this is just because we are implicitly assuming that individual acts are not instrumentally merely superfluous—an assumption this approach cannot help itself to. But suppose I am wrong and there is a moral reason of the sort that Simple Participation identifies. The problem then, I'll now argue, is that it faces the disconnect problem.

According to Simple Participation, you have moral reason to add your pint of water coming from the fact that doing so would make you a 'participant' in the group that helps the suffering men. Doing so would make you a participant because, even though (we are assuming for the sake of argument) the addition of your pint is instrumentally useless, it does add to the underlying dimension (how much water is in the cart). Notice though that, insofar as this is correct, this moral reason can be equally well-responded to by adding just a drop of water from your pint (and keeping the rest), or if it were possible, adding just a

15. From a correspondence, my impression is that Wieland thinks I have shown that their account of participation is not right, but that *with* the right account, we would be able to appeal to expressive reasons to participate in beneficial groups, and to avoid participating in harmful groups, as an answer to the problem of collective impact. But I see no reason to think there is a fully non-instrumental account of participation that would be able to do this trick.

molecule. Adding just a drop, or just a molecule, adds to the underlying dimension. So, it does make you a participant in the relevant sense. And if we are accepting that adding a pint is of no instrumental significance, adding a pint seems *no better* a way of counting oneself as a participant than just adding a drop or just a molecule. Both make you a participant, and neither does anything—we are assuming—to progress the cause of alleviating the men’s suffering. So, neither seems to be a superior way of responding to the participation-based moral reason.¹⁶

Why is this a problem? Because *just* adding a drop of water, or *just* adding a molecule, is not (in Drops of Water, as described) doing the kind of thing that if enough people do it, the harmful outcome will be prevented. The continued suffering will *not* be prevented, or even reduced, if people just give at most a single drop. As we said, to solve the problem, we need an explanation as to why each of us has moral reason specifically to do the sort of thing that *if* enough of us do it, will prevent the harm of concern. A reason to add a pint that could be responded to just as well by, instead, adding just a drop or just a molecule is no such reason.¹⁷

For another way of seeing the disconnect problem, consider a point similar to the one I made about holding the granola bars in Vending Machine. Suppose

16. This is something Wieland and van Oeveren embrace in their article. They say that their account is ‘degreeless’: ‘If you add to D, you participate all the same’ (19). In a later article, Wieland (2022) revises this, saying that participation-based reasons do come in degrees. But his point there is that they vary in degree based on how likely it is that the group will cause harm or will cause benefit. There is no argument to suggest that the strength of the reason varies with the size of the contribution you would make. One might think, though, that obviously it does. Since ‘contribution to the underlying dimension’ makes one a participant, doing so more must make you more of a participant, and there is good reason to participate more. But while participation sometimes works like this, it does not always. I might count as participating more in the class by speaking up more, or as participating more in the party by dancing more. But while my occupying a seat at the symphony does make me a participant in the event, stretching out so that I occupy two seats does not make me more of a participant in it. Making an instrumentally *irrelevant* contribution along the underlying dimension (e.g., the adding of the quarter in the vending machine example) seems like it would be participation of the second sort rather than the first sort—a simple ‘in’ rather than ‘out’ way of participating. If, on the other hand, the contribution to the underlying dimension is *not* merely superfluous to what the group is doing, it would make sense to regard a larger such contribution as more participatory than a smaller such contribution. So, a way out of this issue with Simple Participation would be to refine the conditions for participation and argue that, on the revised conditions, the individual contribution is not instrumentally irrelevant. But of course, this would then no longer be a non-instrumental solution to the problem; it would be a solution of exactly the sort I am saying we need.

17. This is not to say that a view that says there is reason to add just a molecule cannot count as a solution to the problem. If the reason to add a pint is substantially stronger than the reason to just add a molecule, this could be adequate. The issue with Simple Participation is that it does not give a basis for regarding the reason to add a pint as any stronger than the reason to add just a drop. In other words, the consideration identified can be just as well responded to either way.

I don't want to give up my pint of water, and you say to me, 'Why don't you see if you can instead just rock the cart back and forth a little bit before it is taken to the desert? That way you will still have participated in the group effort'. It's true that Simple Participation does not count this as a way of participating. But why not? If making a contribution to the underlying dimension is a morally recommended way of participating *even though* the contribution is entirely useless, then it seems like rocking the cart back and forth should be too. Both involve inserting oneself into the causal pathway but in, we are assuming, a useless, unhelpful way. In other words, just as with adding the quarter and holding the granola bars, there does not seem to be a principled reason being offered by Simple Participation to view these two ways of engaging in what the group is doing differently. So, I should be able to respond to the moral reason their account is getting at (if there really is such a reason) by rocking the cart instead of adding my pint. Such a reason cannot count as addressing what is at issue in our problem.

A proponent of Simple Participation could attempt to show that adding a pint does have a different participatory significance than merely rocking the cart and to show that there is good reason to give one's full pint of water and not just a drop. But this would require additional argument, and I cannot see a plausible, non-stipulative way that this could go that does not involve (effectively) arguing against the instrumental uselessness of adding one's pint to the cart. That is, I cannot see how a proponent of Simple Participation could make the case for these distinctions without (effectively) moving to an at-least partly instrumental approach.

To sum up, I've argued that while this participation view does avoid the original superfluity problem, it faces superfluity problem 2.0: it only seems to spell out a morally significant reason for action when we are (even if only implicitly) thinking of contributions to the underlying dimension as mattering instrumentally. When we look at examples in which we can better control for that sort of intuition (like Vending Machine and Locks), it becomes clear that the sort of consideration identified, taken by itself, is not a moral reason-giving one.¹⁸ But I've also argued that even if I'm wrong about that, and there is moral reason to participate in the bare sense identified by Simple Participation, the view still falls into the disconnect problem. It is not a reason that addresses what is at issue. In the last section of the essay, though, I want to make a different case for the necessity of an instrumental solution to the problem of collective impact.

18. And that it cannot by itself support the applicability of other normative reasons, like expressive reasons.

6. Open vs. Closed Cases, and the Importance of Distinguishing Distinct Problems

Consider the following two pairs of situations:

Pair 1:

Actual Climate Situation: It is true that serious harmful outcomes will occur no matter what. But still, a huge amount depends on what happens going forward. If global emissions are massively reduced, things will be much less catastrophic than various other potential higher-emitting paths forward.

Imaginary Climate Situation: Things are already set to be maximally catastrophic, and there is no way any longer to slow climate change or reduce the severity of climate harms by way of reducing global emissions. We have passed some sort of tipping point with our emissions, such that it all will come crashing down on us, no matter what happens going forward. (I'm not saying this makes any scientific sense. This is purely imaginary.)

Pair 2:

Drops of Water: as given above.

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 adds their pint 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.¹⁹

As far as Simple Participation goes, the moral reason one has to not drive a gas-guzzling SUV, or to not fly, or to not eat beef (connected to climate change) is *the same* (in both type and strength) whether we are in the Actual Climate Situation or in the Imaginary (all is inevitable) Climate Situation. On this view, whether things are already totally doomed, or whether they depend a huge amount on what people do going forward, makes no difference to my decision

19. See Nefsky (2019: 9–10) for a version of this point in connection with a different sort of view.

about what to do. Similarly, on Simple Participation, the reason one has to add one's pint of water to the cart is the same (in both type and strength) whether we are in Drops of Water or Certain Water—whether the relief of the suffering depends on what people do going forward or not at all. This cannot be right. There should be additional, or substantially stronger, reasons for action in the Actual Climate Situation and in Drops of Water than in these alternative 'the outcome is already inevitable' scenarios.²⁰

Let's call Drops of Water and Actual Climate Situation 'open' cases, and Certain Water and Imaginary Climate Situation 'closed' cases. So, open cases are those in which the outcome is up in the air and in which it depends on what people do going forward, while closed cases are those in which the outcome is, for some reason, already settled in the relevant respects. If you agree that there should be additional or stronger reasons in the open cases, this is a different reason to be unsatisfied with Simple Participation and with other purely non-instrumental proposals. It is doubtful that we can capture why there is additional, or stronger, reason for the individual to act in the open cases while embracing that there is no difference in instrumental significance between your contribution in the open cases and your contribution in the closed cases. Just like Simple Participation, other truly non-instrumental proposals (ones that are non-instrumental all the way down), do not seem to be capable of capturing the difference either. Take an appeal to expressive reason, which says that, while your refraining from driving a car is instrumentally irrelevant, it does express your concern about climate change. This view, in addition to facing a version of the dilemma discussed above, faces the present problem: it cannot differentiate our reasons for action in the open and closed cases.

One might reply, though, that while purely non-instrumental approaches have this problem, instrumental approaches just face the opposite problem: they cannot explain why there is still moral reason to act in the relevant ways in the closed cases. I don't myself have a clear intuition that one should add one's pint in Certain Water, or that one should not drive a gas-guzzling SUV in the Imaginary Climate Situation. But I know that some people do, and they might be right. And I do share that sort of intuition in various other cases.²¹ But even if those

20. In Wieland's revised view (Wieland 2022) in which the reasons do vary in strength, it is the reasons in Certain Water and Imaginary Climate Situation that are stronger. On this view, the participation-based reasons get stronger the more certain it is that the group will together bring about the harm (or benefit) in question. So, if the outcome is inevitable, the reasons are the strongest. This might be fine for an answer to a different question (which I am about to identify), but this is very problematic for a proposed solution to the problem of collective impact.

21. For example, I share this intuition in Parfit's *Harmless Torturers*. In *Harmless Torturers* a thousand people each increase the voltage being applied to a thousand victims by a miniscule, imperceptible amount. While each change in voltage seems to make no difference for the worse to the victims, all of these actions together result in tremendous pain (Parfit 1984: 80.) In this example

conclusions are correct, this is *not* an issue for the point that I have been making in this essay, nor is it a problem for instrumental solutions to the problem of collective impact.

First, my point is that an instrumental approach is a necessary and central part of a satisfactory answer to the problem of collective impact. I'm not arguing that it is sufficient for a full picture of our individual moral reasons and obligations in collective impact contexts.

Second, even if we do want to be able to explain why there is reason for making the choices in question in the closed cases, this is not a necessary part of addressing the problem of collective impact itself. Indeed, it is important to see that there are two distinct problems here that should be separated.

The problem of collective impact is a problem about why we (as individuals) ought, or at least have good moral reason, to not act in ways that, while individually seem too small to make a difference, could collectively result in serious avoidable harm. In other words, it is a problem about how serious bad outcomes that are at risk of occurring in the future *could be prevented or avoided* if enough people make certain sorts of choices that seem individually insignificant. A distinct problem is, roughly: How can we explain why there is moral reason not to be a cause, or to associate oneself with the cause, of a harmful outcome *even when* that outcome is inevitable?

This second problem might arise in some collective impact contexts, but it is not at all unique to them. It comes up in various other sorts of areas. An example is Bernard Williams's case of George the chemist, who is trying to decide whether to take a job in morally problematic chemical weapons research. It is stipulated in the example that someone else will take the job if he does not; the idea is supposed to be that, therefore, the outcomes of the research will be just as bad, or worse (because the other job candidate lacks the moral scruples George has), if George does not take the job. Leaving aside what we think George should do all-things-considered, many people think there is at least strong moral reason for him to not take the job. But if the harmful outcomes of the research are inevitable, why would this be? This is an instance of the second problem, and not the problem of collective impact. It is a problem about why there could be reason to not be a cause of harm, or involve oneself in the production of harm, even *if* the harm is inevitable. If we take away the stipulation that the harm is inevitable, there is no question as to why George has moral reason not to take the job.²²

I find it intuitive that there is at least some moral reason not to turn the dial up a notch even if one is the very last one and the full harm is entirely inevitable at that point.

22. For a different example, take an 'overdetermined shooter' case: A is going to shoot B no matter what, and A's shot will be enough to kill B. Knowing this, and supposing that the pain involved will not be any worse if C shoots too, why is there moral reason for C to not also shoot B simultaneously? This is *not* the problem of collective impact. Here the harm is supposed to be

It is possible that a single solution will address both problems, but that does not need to be the case. Indeed, if you agree with me that there should be different, or stronger, reasons for action in the open cases than in the closed cases, then trying for a single solution to both problems is probably not the way to go. In any case, it is no objection to a proposed solution to the problem of collective impact that it does not also address this second different sort of problem. I think many of those who are drawn to pursuing a purely non-instrumental approach to the problem of collective impact, and particularly those who resist my point about the need for an instrumental solution, are making the mistake of running these two problems together and failing to see them as distinct.

Acknowledgments

For helpful comments and discussions, I am very grateful to Antti Kauppinen, Maria Lasonen, Jan Willem Wieland, two anonymous reviewers, an editor at this journal, and audiences at the Invited Symposium on Ethical Consumption and Complicity at the 2020 Central APA and at the Workshop on Individual Responsibility for Collectively Caused Outcomes hosted by the University of Helsinki, June 2021. My work on this paper was supported by an Insight Grant from the Social Sciences and Humanities Research Council of Canada; I am very grateful for this support.

References

- Baatz, C. 2014. 'Climate Change and Individual Duties to Reduce GHG Emissions', *Ethics, Policy & Environment*, 17: 1–19 <<https://doi.org/10.1080/21550085.2014.885406>>
- Benn, S. 1979. 'The Problematic Rationality of Political Participation', in *Philosophy, Politics, and Society*, 5th edn, ed. by P. Laslett and J. Fishkin (Oxford: Basil Blackwell)
- Brennan, G., and L. Lomasky. 1989. 'Large Numbers, Small Costs: The Uneasy Foundation of Democratic Rule', in *Politics and Process: New Essays in Democratic Thought*, ed. by G. Brennan, and L. Lomasky (Cambridge: Cambridge University Press), pp. 42–57 <<https://doi.org/10.1017/CBO9780511528156.003>>
- Brennan, G., and L. Lomasky. 1993. *Democracy and Decision: The Pure Theory of Electoral Preference* (New York: Cambridge University Press) <<https://doi.org/10.1017/CBO9781139173544>>
- Brennan, J. 2009. 'Polluting the Polls: When Citizens Should Not Vote', *Australasian Journal of Philosophy*, 87: 535–49 <<https://doi.org/10.1080/00048400802587309>>

inevitable, and if it were not inevitable, then C's shooting very clearly could make a difference for the worse. This sort of case is often grouped in with collective impact cases and regarded as posing the same problem. But these are actually two distinct problems to which one might take very different approaches and which might require different tools to address.

- Broome, J. 2019. 'Against Denialism', *The Monist*, 102: 110–29
- Cullity, G. 2000. 'Pooled Beneficence', in *Imperceptible Harms and Benefits*, ed. by M. Almeida (Dordrecht: Kluwer Academic Publishers), pp. 1–23
- Jamieson, D. 2007. 'When Utilitarians Should Be Virtue Theorists', *Utilitas*, 19: 160–83 <<https://doi.org/10.1017/S0953820807002452>>
- Kagan, S. 2011. 'Do I Make a Difference?' *Philosophy & Public Affairs*, 39: 105–41 <<https://doi.org/10.1111/j.1088-4963.2011.01203.x>>
- Kutz, C. 2000. *Complicity* (Cambridge: Cambridge University Press)
- Lane, M. 2018. 'Uncertainty, Action and Politics', in *Political Thought and the Environment*, ed. by K. Forrester and S. Smith (Cambridge: Cambridge University Press) <<https://doi.org/10.1017/9781108185509.009>>
- McPherson, T. 2016. 'How to Argue for (and against) Ethical Veganism', in *Food, Ethics and Society* (New York: Oxford University Press)
- Nefsky, J. 2011. 'Consequentialism and the Problem of Collective Harm: A Reply to Kagan', *Philosophy & Public Affairs*, 39: 364–95
- 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 *The Oxford Handbook of Food Ethics*, ed. by A. Barnhill, M. Budolfson, and T. Doggett (Oxford: Oxford University Press), pp. 267–86
- Nefsky, J. 2019. 'Collective Harm and the Inefficacy Problem', *Philosophy Compass*, 14: e12587
- Nefsky, J. 2021. 'Climate Change and Individual Obligations: A Dilemma for the Expected Utility Approach, and the Need for an Imperfect View', *Philosophy and Climate Change*, ed. by M. Budolfson, T. McPherson, and D. Plunkett (Oxford: Oxford University Press) <<https://doi.org/10.1093/oso/9780198796282.001.0001>>
- Norcross, A. 2004. 'Puppies, Pigs, and People: Eating Meat and Marginal Cases', *Philosophical Perspectives*, 18: 229–45
- Parfit, D. 1984. *Reasons and Persons* (Oxford: Clarendon Press)
- Singer, P. 1980. 'Utilitarianism and Vegetarianism', *Philosophy & Public Affairs*, 9: 325–37
- Sandler, R. 2010. 'Ethical Theory and the Problem of Inconsequentialism: Why Environmental Ethicists Should Be Virtue-Oriented Ethicists', *Journal of Agricultural and Environmental Ethics*, 23: 167–83 <<https://doi.org/10.1007/s10806-009-9203-4>>
- Strang, C. 1960. 'What If Everyone Did That?' *Durham University Journal*, 53: 5–10
- Wieland, J. 2022. 'Participation and Degrees', *Utilitas* 34(1), 39–56.
- Wieland, J., and R. van Oeveren. 2020. 'Participation and Superfluity', *Journal of Moral Philosophy*, 17: 163–87