

Fairness, Participation and the Real Problem of Collective Harm¹

1. Introduction

In a wide range of cases, people *collectively* cause a morally significant outcome but no *individual* act seems to make a difference. For example, while climate change is caused in part by millions of people acting in certain ways, any one such act does not seem to make a difference. Take a single driver off the road and the problem of climate change will surely be just as bad. Similarly, while large-scale consumer patterns can have a major impact on the lives of people across the globe, it's hard to believe that a single purchase will itself make a difference. Will things be any better for anyone if, for instance, the cup of coffee I buy this morning is fair trade rather than conventional? Or take the case of voting in a national election, where – let us suppose – it matters a great deal for the welfare of the population which candidate wins. While our votes collectively determine who wins, a single vote won't make a difference: the winner will be the same give or take any individual vote.

The problem in cases of this sort is that it seems each person can argue, “it won't make a difference whether or not I do *X*, so I have no reason to do it.” The challenge is to explain where this argument goes wrong. If we cannot, this would be very troubling; morality would be powerless in wide range of cases in which it should have force. I call this “the Problem of Collective Harm”, and the cases in which it arises “Collective Harm Cases”.²

¹ I am very grateful for helpful comments from Markus Kohl, Niko Kolodny, Mark Lebar, Alex Rennet, Abraham Roth, Carolina Sartorio, Sergio Tenenbaum, R. Jay Wallace, and two anonymous referees for Oxford University Press. I also owe much gratitude to the participants of the Fifth Annual Arizona Normative Ethics Workshop, and to audiences at Carleton University and the 2013 Toronto-Montreal Joint Ethics Centre Conference at the University of Montreal.

² This follows my terminology in Julia Nefsky, “Consequentialism and the Problem of Collective Harm: A Reply to Kagan.” *Philosophy & Public Affairs* 39 (2011): 364-395 and in my “How You Can Help, Without Making a Difference”, draft. I think these names, on balance, do a nice job of giving a sense of the relevant phenomenon, but they are in one way misleading. While in many Collective Harm Cases it is *harm* that is at stake, this need not be so. In a voting case, for instance, it could matter a lot morally which candidate wins, but not because there will be more harm if one candidate wins than if the other does. This would still count as a “Collective Harm Case” on my

There are two ways of trying to solve the problem.³ The first is to try to *deny the description* of the cases: to argue that, contrary to appearances, one's individual act actually does, or at least might, make a difference with respect to the morally significant outcome in question and that there is, thus, no problem after all. The second is to accept that there are cases in which one's individual act cannot itself make a difference, but to argue that there is still reason for action – a reason that doesn't have to do with the difference you make in outcome. Those who take this second type of approach, in other words, respond by *rejecting the implication* from 'it won't make a difference' to 'there is no reason to do it.'

In this paper I am interested in the possibility of taking the second kind of approach – a “Rejecting the Implication” approach.⁴ I centre the discussion around three such proposals. The first two – which I call “Weak Participation” and “Strong Participation” – both claim that we can solve the problem by attending to the fact that individuals in Collective Harm Cases are acting as part of a group that makes a difference. The third proposal – the “Fairness Approach” – attempts to solve the problem by appealing to considerations of fairness.

After raising some initial questions about each proposal individually, I argue that there is a common, basic problem with all three. While they do identify considerations *other than* the difference you make, it's not clear that these considerations can provide the reasons for action we are looking for *as long as* it is true that your act won't make a difference. More specifically, I argue that unless your act could play a significant instrumental role in bringing about the outcome in question, the explanations these views are offering as to why you have reason to do it cannot get off the ground. And yet *not* playing any such role seems to be just what we are

definition. We can think of “Collective Harm Cases” as naming a wider class of cases after a subset, namely the subset in which what is at stake is harm.

³ This excludes biting the bullet and accepting that there is no reason for action in Collective Harm Cases.

⁴ The “Denying the Description” approach will briefly come back into the discussion at the end of this paper. I discuss the plausibility of such an approach elsewhere (in “Consequentialism and the Problem of Collective Harm”).

accepting when we accept that your act won't make a difference. I bring out this problem through a discussion of the Fairness Approach in Section 6. In Section 7 we see how the same problem arises for the other two approaches.

The problem, I then argue, extends beyond these three proposals. In particular, if we try to apply certain general non-consequentialist moral theories to Collective Harm Cases, we encounter the same basic problem. This clarifies why it is that the Problem of Collective Harm poses a genuine challenge not just for consequentialists but also for non-consequentialists.

None of this will be to say that notions of fairness or participation will not be important in understanding the morality of Collective Harm Cases, or that non-consequentialist theories cannot ultimately handle these cases effectively. Rather, the aim of the discussion is to obtain a sharper understanding of what the core challenge in the Problem of Collective Harm really is. This, in turn, helps us identify what the central task needs to be if we are to properly address it.

2. Two Examples

I have given a few quick examples of real-world Collective Harm Cases, but I will largely focus on the following two imaginary cases. Both are from Derek Parfit.

Drops of Water: Ten thousand men lie 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 of us add our pints, this will be enough to alleviate their suffering. But while our acts of adding our pints would *collectively* alleviate their suffering, any individual such act does not seem to make a difference. A single pint added to the cart only allows each man to drink an extra *ten-thousandth* of a pint of water. This is no more than a single drop. And, however much water they are receiving, one drop more or less is simply too miniscule an amount to make any difference to how they feel.⁵

Harmless Torturers: A victim is hooked up to an electric shock machine that has a thousand settings, controlled by a dial. There are a thousand "torturers", each of whom turns the dial up a single notch. At the first setting there is no electric current. At the second setting

⁵ This is adapted from Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), p. 76 and follows my presentation in "How You Can Help, Without Making a Difference."

there is a tiny electric current, but it is too tiny to be perceived by the victim. In general, the settings increase the voltage by increments so tiny that the difference between any two adjacent settings is too miniscule to make a difference to how the victim feels. But, while any two adjacent settings are indistinguishable, after many increases the victim feels pain, and by the thousandth setting the victim is in excruciating pain. While the torturers' dial-turns together result in a lot of pain for the victim, it seems that no individual makes a difference for the worse. Take any given torturer's dial-turn away, and the pain would be just as bad.⁶

The Problem of Collective Harm arises in both cases. In Harmless Torturers, each torturer can argue, "my dial-turn did not make a difference; the victim would have been in just as much pain had I not done it. So I did not do anything wrong." In Drops of Water, each of us can say, "it won't make a difference whether or not I add my pint, so there's no reason to add it." True, the dehydrated men will not be relieved from their pain unless enough of us add our pints. But this does not change the fact that things will be the same for them give or take any single pint.

3. Weak Participation

Weak Participation and Strong Participation both claim that we can solve the Problem of Collective Harm by attending to the fact that, even if you do not yourself make a difference, you may be part of a group that does. *We together* make a great deal of difference in Collective Harm Cases, and this can explain why each of us has acted wrongly or rightly. Weak Participation essentially starts and stops with this claim.

Parfit tries this simple appeal to what people together do in *Reasons and Persons*, giving the following principle, which he calls "C7":

(C7) Even if an act harms no one, this act may be wrong because it is one of a *set* of acts that *together* harm other people. Similarly, even if some act benefits no one, it can be what someone ought to do, because it is one of a set of acts that together benefit other people.⁷

⁶ This is adapted from *Reasons and Persons*, p. 80 and follows my presentation in "Consequentialism and the Problem of Collective Harm".

⁷ *Reasons and Persons*, p. 70. Parfit intends this to be a consequentialist proposal. For a discussion of whether it should be accepted by consequentialists see Frank Jackson, "Which Effects," *Reading Parfit*, Ed. Jonathan Dancy (Oxford: Blackwell, 1997), 42-53, and Ben Eggleston, "Should Consequentialists Make Parfit's Second Mistake: A Reply to Jackson," *Australasian Journal of Philosophy* 78 (2000): 1-15. In this paper, though, I am interested in the proposal independently of whether it would be attractive to a consequentialist.

In Harmless Torturers, for instance, while turning up the dial one notch makes no difference to the harm suffered by the victim, it is one of a set of acts that together cause him a lot of pain. The torturers *together* make a difference for the worse, and so each of them has acted wrongly.

Does this response work? Not as it stands. After all, what determines who is in the group?⁸ Why can't each individual say: my act makes no difference, and so it's *not* part of the set of acts that together makes a difference for the worse? It is unclear, in other words, why I should count as part of the group that harms if what I do *makes no difference* with respect to that harm.

The point is not (or not just) that we need some criteria that will tell us which acts are in the group and which are not. The point is that if we want to derive normative implications directly from my being part of a group, we need an understanding of what this sort of participation consists in that explains why it has these normative implications.

Still, Weak Participation might seem to be pointing down the right path. Weak Participation says: even if you do not by yourself make a difference, if *you together with others* do, this can explain why you have acted wrongly or rightly. Strong Participation can be thought of as an attempt to substantiate this idea.⁹

4. Strong Participation

Strong Participation, unlike Weak Participation, employs a robust conception of acting with others – of participating in a collective action. Using this conception of what it is to participate, it attempts to explain why mere participation, regardless of whether you can make a difference, has normative implications.

⁸ Parfit attempts to address this question. He sets out various principles that attempt to further elaborate on C7 and, if they succeed, would specify who is in the group. But, as he admits in a later paper, they do not succeed: the further principles exclude the very acts he is trying to capture, acts that cannot themselves make a difference. See his "Comments," *Ethics* 96 (1986), pp. 846-848.

⁹ When substantiated in this way, though, it is clearly not consequentialist. (i.e. See note 7 above.)

Christopher Kutz develops such a proposal. On Kutz's account, you participate in a collective action when you act with a "participatory intention", an intention to do your part in a collective project.¹⁰ Suppose we are building a sandcastle. Your act of filling the bucket with water is part of our collective action because it is performed with the intention of contributing to our collective end, a sandcastle. You fill the bucket so that we can mix water with sand, which will allow us to sculpt the castle. If you had instead filled it with the intention of pouring it on my head, this would not have been part of our collective action. "Jointly acting groups," Kutz writes, "consist of individuals who intend to contribute to a collective end."¹¹

Next Kutz argues that (1) you are accountable for actions that are ascribable to you, and for the consequences (whether intended or not) of those actions, and that (2) when you participate in a collective action, that action is ascribable to you. When you take part in a collective action, you are an "inclusive" author of that action.¹² You are one of a number of people who can say, "We did it." It is, in this sense, *your* action, and so you can be held accountable for it and its consequences.¹³

Kutz claims that this holds regardless of whether you individually make a difference. He gives examples like the following. Suppose that you and I are having a picnic. While I am getting the drinks from the car, you put the blanket down on a bed of flowers and ruin it. Even though it is your individual act and not mine that causes this damage, you put the blanket down

¹⁰ Christopher Kutz, *Complicity* (Cambridge: Cambridge University Press, 2000), Chapter 3. There is a further condition, namely that your conception of the joint project sufficiently overlaps with that of the other members of the group. (*Complicity*, pp. 94-95.)

¹¹ *Complicity*, p. 89. There is disagreement about how to understand the phenomena of collective action. Kutz's account – as he explains – is minimalistic. Other accounts bring in stronger conditions, conditions requiring, for instance, mutual knowledge or mutual responsiveness. And they often invoke not just intentions to do *your part* but intentions that *we do the group action*. In part because Kutz's account is already a minimalist one, it is not as though we could resolve the core problem that I will raise (in section 7) by substituting in another available account of collective action.

¹² Where this is contrasted with being the "exclusive" author of an action. (*Complicity*, pp. 105-106.)

¹³ *Complicity*, pp. 137-139.

as part of our collective action of picnicking together. So, while none of my individual acts make a difference for the worse, the damage is a result of a collective action that is ascribable to me.

Thus, I can be to some extent accountable for it.¹⁴ A similar thing can be said in Collective Harm Cases: even if one's individual act makes no difference with respect to the outcome of concern, one can be accountable if one is a participant in a collective action that results in this outcome.

One issue that arises in trying to solve the Problem of Collective Harm in this way is that there are, as Kutz points out, many Collective Harm Cases that simply do not involve collective action. These are cases of “unstructured collective harms”: cases in which harm results from many people acting in certain ways, without there being any joint project that they are all intentionally participating in.¹⁵ Much environmental damage is unstructured harm; drivers, for instance, are not intending to do their part in some collective project whenever they drive. Similarly, people who buy conventional rather than fair trade goods are not intending to do their part in a joint project.

There might be, however, a simple way of extending the picture to cover unstructured cases. Typically when there is a risk of unstructured collective harm, there is also an opportunity to participate in a collective action of preventing that harm. In *Drops of Water*, if none of us add our pints and the men continue to suffer, we would not be intentionally participating in a collective project. So, this would be a case of unstructured collective harm. But if we *do* add our pints, it seems we would be engaged in a joint project of alleviating their suffering. Similarly, while drivers are not engaged in a collective project when they drive, they could participate in a collective action aimed at mitigating environmental damage if they each reduce the amount they drive with the aim of doing their part in such a project. If Strong Participation explains why there

¹⁴ *Complicity*, p. 154.

¹⁵ *Complicity*, p. 166.

is reason to refrain from participating in harmful collective actions even if you make no difference, it should also be able to explain why there is reason to participate in beneficial collective actions even if you make no difference.

The next question to ask, though, is: how exactly does Strong Participation explain why there is reason to do either? So far, what it has explained is why participants in a collective action can be accountable, even if they make no difference. But, at least on Kutz's notion of accountability, asserting that someone can be accountable for a harm is not equivalent to saying that this person did something that she ought not to have done. Accountability, for Kutz, concerns "warranted response", where this is not just a matter of responses owed "because of the rights and wrongs" one has done.¹⁶ Think of the picnic example. Perhaps it is true that I am partially accountable for the damage caused to the flowers: perhaps I ought to express apology or help compensate the gardener because of the way in which I am connected to the damage through our collective action. But one thing that is clear is that I did not do anything I ought not to have done. Everything I did was perfectly fine. As Kutz writes:

I am not at fault with regard to the gardener, but it seems I do nonetheless owe some form of response ... Because I have made our picnic mine by my intentional involvement, I have also made its consequences mine. My accountability is therefore intermediate, between what one owes for faultless harms (for *I* was not at fault) and what one owes for faulty harms (for *you*, and *hence we*, were). Your fault sticks to me but not with its full force.¹⁷

Are things similar in Collective Harm Cases? In Harmless Torturers, my individual act of turning the dial makes no difference to the suffering the victim undergoes. Perhaps I am to some extent accountable because of the way in which I am connected to the harm through my participation. But if what I did couldn't have made any difference with respect to the harm, how does Strong Participation give us grounds for thinking that I ought not to have done it?

¹⁶ *Complicity*, p. 18. That is, accountability for Kutz is specifically not a retributivist or desert-based notion.

¹⁷ *Complicity*, p. 154.

I take it the reply is: in the picnic case I have no reason to expect the bad outcome; this is why we cannot say that I ought to have acted otherwise. In Harmless Torturers, on the other hand, the collective project *is* a harmful one and so I can expect that the collective action will result in harm. If you expect that a collective action will result in harm, then you ought to refrain from participating in it, and – crucially – you ought to refrain from participating regardless of whether you could make any difference for the better by refraining. This is because if you participate, you will be an “inclusive author” of the harm. The harmful collective action will be, in a sense, yours – something that *we*, and thus inclusively *you*, did.

Intuitions will likely diverge here. Some might find this picture appealing. Others might think: while I can expect that the collective action will result in harm, I also know (in advance) that my “inclusive authorship” won’t make any difference with respect to that harm. Why should that sort of inclusive authorship make my individual act wrong? Why not think that the most it can do is make me accountable in a faultless, associative sense?¹⁸

5. The Fairness Approach

A common idea is that the reason one ought to act in the relevant way in Collective Harm Cases is that not to do so would be unfair. Garrett Cullity advances such a proposal, and it is his version that I will make use of here.¹⁹

¹⁸ Of course, I might not want to be even faultlessly accountable; I might not want to owe or be owed any response. In that case, the risk of being accountable could itself give me reason to refrain from turning the dial. But, divorced from any claim that I would be doing something I ought (or, at least, have independent reason) not to do, this appeal to the risk of accountability seems to simply give a reason concerning how turning the dial could affect *me*. This seems to be the wrong sort of reason to respond to the Problem of Collective Harm that arises in Harmless Torturers.

¹⁹ The appeal to fairness is common for those concerned with Collective Harm Cases in which not acting in the relevant way is considered *free-riding*. These are cases in which the relevant agents are also the beneficiaries. For example, why should I pay my subway token if I can sneak on? \$3 more or less surely won’t make a difference to subway operations. A common answer is: I am only able to gain the benefits of the subway *because* others pay their tokens; so, to not pay would be unfair. Not all Collective Harm Cases are potential free-rider cases. In Drops of Water, if I keep my pint, I am not gaining a benefit that is only possible because others have given up their pints. (Note also that not all free-rider cases are Collective Harm Cases.) But while the two kinds of cases only overlap, and are not identical, it seems plausible that not acting in the relevant way is wrong for the same reason in both

The Fairness Approach first claims that in Collective Harm Cases there is a collective obligation to prevent (or to avoid bringing about) the relevant bad outcome. Unless enough people act in a certain way, bad consequences will result, and so there is an obligation on the group as a whole to ensure that this does not happen. But if we have a collective obligation – the argument goes – it is unfair for some but not others to do the work of satisfying it. If I stand by while others do their part in fulfilling the obligation, this is not fair to the others in the group and is therefore wrong. Thus, regardless of whether or not one could make a difference by doing so, there is reason to act: “it is a matter of pulling one’s weight in what we all ought to be doing.”²⁰

Suppose that each of the potential torturers would receive a benefit for turning up the dial – say, fifty dollars. Each knows that he won’t make a difference to the victim’s pain by doing so. So, why should he refrain? According to the Fairness Approach, without making any assumptions about the torturers’ individual obligations, we can say that there is an obligation on the *group as a whole* not to cause the victim to suffer. Most people in the group must refrain from turning the dial if this collective obligation is to be met, and if most people refrain while you go ahead and claim the fifty dollars, this is unfair. So, the right thing to do is to refrain.

Now, the point is not that from a collective obligation we can always derive individual obligations. If no one is trying to satisfy the collective obligation, I am not being unfair by refraining from acting myself. The idea is, rather, that when there is a collective obligation *and* people are trying to fulfill it, it can be unfair for me not to do my part too. Cullity explains:

I have not been arguing that individual imperatives can be derived from every collective imperative... Before the collective action is started, none of us is arrogating any special privileges in refusing to get it started. We are collectively acting wrongly, but no individual is taking advantage of others’ propensity to contribute to doing what we ought to be doing.²¹

kinds of cases. This is the idea that Cullity pursues in “Pooled Beneficence”, *Imperceptible Harms and Benefits*, Ed. M. Almeida (Dordrecht: Kluwer Academic Publishers, 2000), 1-23.

²⁰ “Pooled Beneficence,” p. 17.

²¹ “Pooled Beneficence,” p.18.

This, however, seems to be a problem for the Fairness Approach. This means that as long as all the torturers continue turning up the dials, or as long as no one has given a pint to the cart, the Fairness Approach has nothing to say about why one ought to act otherwise. But the Problem of Collective Harm is that it's not clear why *any* individual has reason to give up her pint, or to refrain from turning the dial, and so on. Each individual such act seems to make no difference, and thus it is hard to see why anyone has reason to do it. If the Fairness Approach only kicks in once people have already started acting in the relevant way, it can't count as a full solution to the problem.

In response, perhaps we can say that while the collective obligation does not itself entail individual obligations, its presence does give individuals, right from the start, at least some morally relevant reason for action. While I may not have an obligation to add my pint if others have not yet done so, perhaps the fact that we are collectively obligated to alleviate the suffering explains why I have at least some reason to add my pint. Once others, or enough others, have added their pints, I then become obligated to do so because not doing so would be unfair.

I do not think this reply succeeds. But this is a consequence of a more basic and important problem – one that arises for all three of the views under discussion. Let's turn to that problem.

6. Fairness and Difference-Making

Jonathan Glover dismisses the Fairness Approach, claiming that it invokes a notion of fairness that we do not or, at least, should not care about. He considers a case in which a car needs to be pushed up a hill; only six people are necessary for the job, but there are eight altogether. Even though six could manage the task on their own, it would be unfair for two to sit by while others do all the work. This sounds much like the point that the Fairness Approach tries

to make in Collective Harm Cases. In the voting case, for instance, the Fairness Approach says: even though we do not all need to vote, it would be unfair for some but not others to do so. But, Glover replies:

We are free to accept the argument from injustice in the car pushing case without accepting it in the voting case. If I do not push the car, the others will have to push a bit harder. Many of us are against the kind of injustice that involves giving benefits to some at the cost of additional hardships to others. But no-one has to vote harder because I do not vote. It seems a dog-in-the-manger version of justice that objects to one person benefiting because others are left unchanged.²²

Glover is right that there is an important disanalogy between the car-pushing case and the voting case.²³ But it misunderstands the Fairness Approach to say that it invokes a “dog-in-the-manger version of justice” or that it “objects to one person benefiting because others are left unchanged.” Take Drops of Water. Unlike a “dog in the manger” situation, each of us would benefit from keeping our pints; we would each have it as a refreshing drink. The point the Fairness Approach makes is that it would be unfair for some to keep their pints, receiving the benefits that *each of us* would have enjoyed, while others give theirs up in order to achieve what *we as a group* are obligated to achieve. The type of unfairness involved in keeping my pint is not that doing so benefits only myself and not others – that it benefits one person while leaving others unchanged. It is, rather, that if I keep my water, while others give theirs up, I am not – as Cullity writes – “pulling my weight in what we all ought to be doing.”²⁴ I am “relying on others to satisfy ... [the collective] imperative, while excepting myself from doing so.”²⁵ It seems reasonable to propose that this sort of unfairness is morally objectionable. And someone who

²² Jonathan Glover, “It Makes No Difference Whether or Not I Do It,” *Proceedings of the Aristotelian Society* 49 (1975), p. 182.

²³ Here is the disanalogy. In the car pushing case, my act is not necessary for achieving the desired outcome – there are ways of bringing about the outcome without it. But in Collective Harm Cases, it’s not just that my act is not necessary; it’s that it simply *can’t* make any difference. We can take my act away, and not replace it with anything (e.g. not replace it with more effort on the part of the others), and the outcome of concern will be the same in all relevant respects. Because of this disanalogy, while appealing to fairness might work to explain why I ought to push in the car-pushing case, it won’t work as a solution to the Problem of Collective Harm, as we will see shortly.

²⁴ “Pooled Beneficence”, p. 17.

²⁵ “Pooled Beneficence”, p. 17.

Julia Nefsky

finds this view compelling will not be moved by the reply that we should only care about unfairness *if* it is matter of one person gaining a benefit by making things worse for others. Such a reply seems to simply contradict the proposal, rather than identifying a problem with it.

Glover's objection is, thus, not effective as it stands. But there is an important issue in the vicinity. The Fairness Approach says: if we collectively ought to bring about some outcome, it is unfair to let others do the work of satisfying this obligation without doing so yourself. You ought to pull your weight in what we collectively ought to do, regardless of whether in doing so you will make any difference. Glover's reply is: we shouldn't care about the unfairness in not pulling one's weight if it won't make any difference. But the problem is more basic. If acting in the relevant way won't make any difference then it does not seem that it *pulls any weight at all*. It doesn't seem to count as *doing the work* of satisfying the collective obligation.

Consider this revised version of the car-pushing example. We need to get the car up the hill, but one of us – call him “Small” – has a rare physical condition: while he is otherwise healthy, he has the muscles of a two-year-old. He is, let's say, an adult trapped in a two-year-old's body. Because of this, it makes no difference whether or not Small pushes. If he pushes with all his might, this will exert a tiny amount of force on the car, but this tiny amount of force is too tiny to make any difference at all with respect to the task at hand. The car won't progress up the hill any faster if he pushes than if he doesn't, nor will it be any easier for the others involved. And this is so regardless of how many others push. Whether or not enough other people push to get the car up the hill, adding or removing Small from the equation simply won't change things in any relevant respect.

In this situation, does it make sense to say that Small ought to push because this is his fair share of our collective burden? His act of pushing wouldn't help us achieve the goal. It would be

merely superfluous. It is true that there is a burden on the collective to get the car up the hill. But if his pushing won't help discharge this burden, how could it be his fair share of the burden? It doesn't seem to be any real share at all. Instead it seems like a mere waste of his efforts. Perhaps Small should do *something* out of fairness. If there is some useful task he could perform, he might have a duty of fairness to do that. But it doesn't make much sense to say that what he ought to do is something merely superfluous and unhelpful.

In Collective Harm Cases each individual act seems to be like Small's act of car-pushing. It appears that it can't make any difference with respect to the outcome of concern, and so seems to be merely superfluous. But what we are seeing is that the Fairness Approach cannot get off the ground if this is so. The Fairness Approach appeals to (a) the presence of a collective obligation to bring about some outcome, and (b) the idea that it is unfair to let others take on the task of satisfying this obligation while excepting yourself from doing so. But as long as my act won't help satisfy the collective obligation – as long as it would be merely superfluous with respect to what we are collectively obligated to do – (a) and (b) cannot explain why I should do it. Even if there is a sense in which we are collectively obligated to relieve the men's suffering in Drops of Water, if adding my pint to the cart *won't help* to relieve their suffering, this collective obligation wouldn't give me reason to add my pint. And bringing in a requirement to do one's fair share does not fill the gap. How could doing something merely superfluous with respect to the alleviation of their suffering count as doing my fair share of our obligation to alleviate their suffering? It seems to be no real share at all.²⁶

²⁶ Cullity considers and tries to respond to a similar objection. Part of his response is to argue that one's individual act counts as making a real contribution toward meeting the collective obligation because "the contribution is perceptible" even though "it has no perceptible effect on what is achieved by the group" ("Pooled Beneficence", p. 16). In Drops of Water, for instance, my contribution is perceptible: a pint is a perceptible amount of water. This does not work. The perceptibility of a pint of water is irrelevant to the question of whether it actually contributes in any morally significant sense to the task of alleviating their suffering. The problem here is similar to a problem I discuss in "Consequentialism and the Problem of Collective Harm"; there is a tendency in writing about this topic to

I am not claiming that the idea that I ought to do my fair share has no force in this case. As with Small's case, if there is some useful task I could perform I might have that sort of duty to do it. The point is that as long as it's true that *pouring my pint into the cart* wouldn't do anything useful toward alleviating their suffering – that it will not help matters in any real way – it does not seem that *that* would be the thing to do in a response to a requirement to do my part in the task of alleviating their suffering.

One might reply as follows: “putting the idea of ‘doing one’s fair share of the collective obligation’ aside, there is a simple explanation as to why fairness gives one reason to add one’s pint. If we really think that *fairness* matters morally in itself, then it does not need to be that adding a pint would be useful with respect to satisfying this collective obligation. It is unfair for some of us to keep our pints if others have given theirs up *just because* this would mean having something that those others no longer have.”

Now, this kind of fairness – unlike the one that occurs in the Fairness Approach – is the “dog-in-the-manger” type of fairness that Glover finds objectionable. It asks us to “level down”: to give something up, lowering our own position, *simply* for the sake of making things more equal. Still, while many object to the idea that there is any reason to level down, some do think that there is moral reason to do so and that it only seems objectionable if we do not attend to the fact that this reason can often be outweighed.

Regardless, this is really beside the point. Even if you do have reason to give up your pint coming simply from the fact that this would make things more equal amongst the original pint-holders – that is, even if there *is* a reason to “level down” – this is not a reason that can count as a solution to the Problem of Collective Harm. We can see this by noting that it does not explain

confuse *perceptible* with *morally significant*. Still, while Cullity's point about perceptibility does not work, some of his other comments hint in a potentially promising direction.

why I have reason to add my pint to the cart, rather than – for example – to pour it onto the ground. As far as bare equality in the resources of the original pint-holders is concerned, it would be just as effective if I pour my pint onto the ground. To respond to the Problem of Collective Harm we need an explanation of why I have reason *specifically* to add my pint to the cart. That is, we need an explanation of why each of us ought, or at least has reason, specifically to do the sort of thing that if enough of us do it, will alleviate the suffering. A reason that is just as much a reason to, instead, pour one’s pint onto the ground is no such reason.

Notice, moreover, that if I have reason to add my pint coming simply from the fact that this would make the resources more equal amongst the group of original pint-holders, this is a reason that has nothing particularly to do with the suffering men in the desert. It only concerns how my resources compare to the others in this group. Any such reason cannot count as a solution to the Problem of Collective Harm. We need a reason for action that connects appropriately to the fact that the men in the desert are suffering. Otherwise, we are essentially just changing the topic.

The Fairness Approach tries to provide this connection by invoking a specific kind of fairness – that of pulling one’s weight in what we collectively ought to do. But, as we have seen, it’s not clear that this can work as long as we are accepting the claim that one’s individual act won’t make any difference.²⁷ Let’s call this “the Superfluity Problem”.

²⁷ Some might object to the comparison between Small’s case and Drops of Water. In Small’s case, everyone *but* Small is able to do something instrumentally significant; Small is an exception. In Drops of Water, on the other hand, everyone is symmetrically situated: each can add at most a pint, and so *no one* can make a significant difference, though we together can get the job done. If we made Small’s case like that it would have to be that everyone was like Small. There could be thousands of people with Small’s condition who need to get a car up a hill; they can do so if most of them push, but any one of them won’t make a difference. When we move to this “symmetrical” version it no longer seems unintuitive that Small is doing his fair share by pushing.

My response is largely to agree. I agree that there is an important difference between the original Small case, on the one hand, and Drops of Water and this “symmetrical” Small case on the other. I agree that, while there might be no (morally relevant) reason for Small to push in the original case, Small *does* have reason to push in the symmetrical version, and each of us has reason to add our pints in Drops of Water. This is because I think there is a

7. A Common Problem

The Superfluity Problem arises just as much with the other two views. We have essentially already seen this for Weak Participation. Weak Participation says that even if you don't harm anyone yourself, if *you together with others* do this explains why you have acted wrongly. But in Collective Harm Cases, it's not just that you don't cause the harm in question all by yourself; it's that you *make no difference* with respect to that harm. Your act doesn't seem to play any significant role in bringing it about. So, why should you count as part of the group to begin with? Why is your act part of the set of acts that harms, if it does not contribute in any non-superfluous way to bringing about the harm? To come at the point another way: as long as your act doesn't play any significant role in what the group does, we cannot simply *point* to what the group does to explain why you have acted wrongly (or rightly). And yet this is just what Weak Participation tries to do.

Strong Participation does not do any better. Whether or not we find appealing the idea that there is something wrong with participating in a harmful collective action even when one's doing so can't make a difference, there is a more basic issue. If your act can't make any difference, we may not be able to count it as an act of participation to begin with. In Harmless Torturers, Strong Participation says: 'even if each torturer makes no difference to the harm, the torturers would be engaged in a collective action of harming the victim. By turning the dial, you are participating in this collective action. This makes the collective action in a sense *yours*, and

solution to the Problem of Collective Harm. (Note that the symmetrical Small case *is* a Collective Harm Case.) Moreover, I do not deny that it ultimately will make sense to think of adding a pint as one's fair share of the collective obligation. The point about symmetry is certainly not enough to establish this, but it does – I think – help refocus our attention in the right direction: namely, one of trying to identify features that will allow us to explain why in Collective Harm Cases it is not actually true that one's individual act is instrumentally superfluous. What our discussion in this section is meant to reveal is that we cannot deal with the "it makes no difference" argument by invoking the idea of a duty to "pull one's weight", or to do one's "fair share." Instead, *before* we can make use of these ideas in understanding the morality of Collective Harm Cases, we first need to be able to debunk the claim that my acting in the relevant way would be instrumentally superfluous and unhelpful.

so you are accountable for it.’ But, on closer inspection, it’s not clear this story can get off the ground. For any given torturer, how can we say that she is participating in this collective action? To participate in a collective action is to intentionally *do your part* in a joint project – in this case, presumably, a project of torturing the victim.²⁸ But for any given “torturer”, her act of turning up the dial a single notch makes no difference to the pain of the victim. It doesn’t seem to play any non-superfluous role in causing the harm. So, how can we say that in turning it she is intending to *do her part* in this joint project? Her intention in turning up the dial could just be to acquire fifty dollars at no cost to anyone. Indeed, if she decided to turn the dial only after realizing that doing so won’t make any difference, that probably is her intention.

Kutz explains that the intentions that make an act participatory need not be “explicit in deliberation”; they can be “functionally implicit” in one’s behavior.²⁹ But this does not help here. We cannot say that while an intention to do her part in a collective project of torturing the victim may not be explicit in her deliberation, it is functionally implicit in her behavior. Her behavior, after all, is behavior that *won’t make any difference* to the torture the victim undergoes. The extent to which the victim is tortured will be exactly the same give or take her turn of the dial. So, we can’t simply appeal to that action to reveal that she is implicitly trying to do her part in a joint project of torturing the victim. Nor can we say that there are other aspects of her behavior that reveal her implicit intention to do her part in the collective project; that would be stipulating extra features that need not be present in the case.

We have said that to handle cases like Drops of Water or environmental damage – cases of “unstructured collective harm” – Strong Participation can talk about reason to participate in a collective action aimed at preventing or mitigating the harm. But the same problem arises here.

²⁸ There is no other potential joint aim specified in the example.

²⁹ *Complicity*, p. 82.

Again, participating in a collective action is to intentionally *do your part* in a collective project. In Drops of Water it would be a project of alleviating the men's suffering; in the case of environmental damage it would be a project of mitigating environmental damage. But how can doing something that would be merely superfluous with respect to our bringing about these outcomes count as doing my part in our bringing them about? Kutz writes, "my part is defined as the task I ought to perform if we are to be successful in realizing our shared goal."³⁰ A task that can't make any difference with respect to our shared goal seems to be no such task. Indeed, if we could have said that adding my pint is the task that I ought to perform if we are to relieve the men's suffering, there would not have been a Problem of Collective Harm there to begin with.

Advocates of Strong Participation might object that I am presupposing an instrumental conception of "doing my part". Why can't Strong Participation invoke a broader conception – a conception under which my act need not do anything useful toward the achievement of the shared goal in order to count as doing my part? Indeed, while in some places Kutz writes things like "participatory intentions can, thus, be seen as merely a species of ordinary, instrumental intentions, differentiated by the group-oriented content of the goal they specify"³¹, in other (nearby) places, Kutz explicitly endorses a broader conception. For example:

Contributory relations ["my part" relations between individual act and collective end] might take an **instrumental** form if what the agent does helps cause the collective outcome (my pushing helps to move the car), or if the agent's part is a constitutive element of the group act (stepping this way is part of dancing a tango). The relation might be **expressive** if by doing one's part, one thereby exemplifies one's membership in a group or participation in an activity, as when by voting I express my membership in a political community. And the relation might be **normative** if one performs one's part because of norms internal to some group or institution that demand certain behavior (I wear a dark suit as an IBM employee.) ... What makes my behavior participatory is nothing more (and nothing less) than my conception of what I do as *related* to the group act.³²

³⁰ *Complicity*, p. 81.

³¹ *Complicity*, p. 84.

³² *Complicity*, p. 82, my emphasis.

If Strong Participation invokes this broad understanding of doing one's part, can't it avoid the Superfluity Problem by saying that it is *non-instrumental* contributory relations that are at play in Collective Harm Cases? For instance, while adding a pint will not do anything instrumentally significant, it can still count as doing my part in our collective project of alleviating the suffering, because it expresses my support for the project or my solidarity with those who have donated.³³

The problem is, though, that on this proposal, Strong Participation cannot explain why I have reason *specifically* to add my pint to the cart, rather than – for instance – to wear a T-shirt that says, “I support the rehydration project!” Especially insofar as we are accepting that neither amounts to doing anything instrumentally significant, wearing the T-shirt could just as well express my support, or my solidarity, as adding my pint could. Strong Participation, in other words, would be unable to differentiate between an act of pouring in one's pint and any other instrumentally insignificant act that could similarly express support or solidarity: neither makes any significant causal contribution to the achievement of the collective goal, and either can make you a participant in the collective project because of its expressive features.

The problem here is similar to the one we saw with the attempt to appeal to a broader notion of fairness. Turning to non-instrumental ways of doing one's part in the joint project avoids the Superfluity Problem only by deflating Strong Participation's potential to address what is at issue in the Problem of Collective Harm. As we've said, to address what is at issue, we need an explanation of why each of us *specifically* has reason to do an act that is of the sort that if

³³ Why not say that adding a pint counts as “doing my part” in the *constitutive* instrumental sense? We could say that my adding a pint is a constitutive element of the collective act of *together adding our pints to the cart*. The problem with this is that, at least on Kutz's account, a collective act is defined in terms of the collective end. The collective end could be either (a) a state of affairs, *or* (b) an activity (e.g. dancing the tango), *or* (c) an institution of some kind. (Kutz, p. 82) In *Drops of Water*, the most plausible understanding of what our collective end would be is the *alleviation of the men's suffering* (a state of affairs), *not* the activity of together adding our pints to the cart. We would not be adding our pints for its own sake, with the alleviation of the suffering as a mere foreseeable consequence. (When we dance the tango, on the other hand, we may be doing this for its own sake, rather than to bring about some further state of affairs.)

enough of us do it, this will collectively cause (or, depending on the case, prevent) the morally significant outcome of concern. Finding a reason to add a pint to the cart that is just as much a reason to instead wear a T-shirt that expresses one's support or solidarity does not do this.

To summarize: Rejecting the Implication approaches try to solve the Problem of Collective Harm by accepting that your individual act can't make any difference, and explaining why nonetheless there is still reason to do it (or, depending on the case, reason to refrain). But the problem is that – at least for the three views we have seen – the explanations they offer do not seem to work as long as your act would be instrumentally superfluous. And yet that seems to be exactly what we are accepting when we accept that your act can't make any difference.

8. Can't We Just Apply a Non-Consequentialist Theory and Call it a Day?

It is tempting to think that it is only consequentialist theories that will have trouble with Collective Harm Cases.³⁴ After all, non-consequentialists do not think that all that matters morally is the difference you make in outcome. So, the suggestion goes, their theories will do just fine. Our discussion of the Superfluity Problem helps reveal why this view about non-consequentialist theories is mistaken.

Recall, we cannot solve the Problem of Collective Harm by giving a reason for action that has nothing particularly to do with the morally significant outcome in question. I may have all sorts of reasons to refrain from driving: to get more exercise, to save money, to reduce the risk of causing or being in an accident, to avoid stressful traffic jams. But, of course, none of these can provide a solution to the Problem of Collective Harm that arises with respect to car driving and environmental damage. We need a reason for action that connects in some appropriate way to the fact that widespread car use causes environmental damage.

³⁴ This is a common reaction I have received from philosophers in conversation.

The problem for non-consequentialists is that, even though they do not think that all that matters morally is the difference you make in outcome, it's not clear that there can be any reason to *X*, which connects appropriately to outcome *Y*, if *X*-ing cannot make any difference with respect to *Y*. If acting in the relevant way 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 you have reason to do it.

Suppose, for instance, that you subscribe to a theory of Virtue Ethics. Suppose you think that the right thing to do is what the virtuous person would do. Can this view explain why you ought to add your pint to the cart, even if doing so won't make a difference? It's not clear that it can. If it really won't make a difference to the well-being of the suffering men, then it's not clear that the virtuous person would add her pint. Surely, the virtuous person does not act wastefully, and insofar as it is true that adding a pint won't make any difference, doing so appears to be a waste. Unless we can say that adding a single pint could genuinely *help* to alleviate the men's suffering, it does not seem we can say that this is what the virtuous person would do.

Or, consider the virtues and vices that a Virtue Ethicist might appeal to in order to explain why one ought to add one's pint. They might want to say that keeping your pint when these men are suffering would be greedy or selfish, and that to donate it would be generous or compassionate. But if adding your pint would be superfluous with respect to the alleviation of the men's suffering, it doesn't seem any of these would apply. Making a 'donation' that will not do anything useful does not seem generous or compassionate; it seems foolish and wasteful. And

if you have considered whether it would be of any use to pour in your pint and have correctly determined that it would not be, then keeping this water does not seem greedy or selfish.³⁵

Now, suppose you are a Kantian. Suppose you think that Kant's Formula of Humanity gets things right. You agree that you should "act that you use humanity, whether in your own person or in the person of any other, always at the same time as an end, never merely as a means."³⁶ Can this explain why you ought to refrain from turning up the dial a notch in Harmless Torturers? If turning the dial won't make a difference for the worse to anyone, it's not clear why it would count as treating someone as a mere means. If you are not contributing in any non-superfluous way to bringing about the victim's suffering by turning the dial, then how can we say that in doing so you are treating him as a mere means? You don't seem to be treating him in any way at all by doing that particular act. Unless your act would play some significant role in causing the outcome, it does not seem we can appeal to the Formula of Humanity to explain why you ought not to do it.³⁷

'Okay', you might say, 'but surely the more promising place to look is Kant's Formula of Universal Law.' That formulation might seem to be perfectly suited to avoid the Superfluity Problem. Let's see if this is so.

The Formula of Universal Law says that an act is right if and only if you can will that the maxim of the action become a universal law of nature without contradiction.³⁸ To determine whether this condition is met, we must imagine a world in which, by your will, everyone acts on

³⁵ Walter Sinnott-Armstrong makes a similar point in "It's Not My Fault: Global Warming and Individual Moral Obligations," *Perspectives on Climate Change: Science, Economics, Politics, Ethics* (Advances in the Economics of Environmental Research), Volume 5 (2005), pp. 303-304.

³⁶ Immanuel Kant, *Groundwork of the Metaphysic of Morals*, Trans. and Ed. Mary Gregor (Cambridge University Press: Cambridge, 1997), 4: 429.

³⁷ The Formula of Humanity might imply that you have other duties with respect to that victim: for example, to do what you can – if anything – to alleviate the pain, or to help him get back on his feet afterward.

³⁸ *Groundwork of the Metaphysic of Morals*, 4: 421.

the maxim in question, and we must look for, first, a contradiction in conception³⁹ and, second – if there is no contradiction in conception – a contradiction of the will.⁴⁰ If there is neither, acting on the maxim is permissible.

There is a difficult question concerning what should go into the description of the maxim. But it seems that for many Collective Harm Cases, however we describe the maxim, we are not going to get a contradiction in conception.⁴¹ In Drops of Water, whether we think of the maxim as ‘refrain from adding my pint of water to the cart, so as to keep it for myself’, or as simply ‘keep my water supply for myself’, or as ‘don’t give something away when doing so won’t make a difference’, there is no contradiction in conception. For each of these, it is possible to act successfully on the maxim in a world in which everyone has the same maxim.

What about the ‘contradiction of the will’ test? Here we might get different results depending on how we describe the maxim. Suppose, first, that we describe the maxim with a fair amount of specificity. For example, let’s take the maxim to be ‘refrain from adding my pint of water to the cart, so as to keep it for myself’. We can universalize this maxim without a contradiction of the will. For there to be a contradiction of the will it must be that my willing this maxim as a universal law is in contradiction with something else that I must will as a rational agent. There is no such contradiction. *I* will not be negatively affected if no one puts water into this cart. The maxim described in this way passes the test.

But the problem might be that we are describing the maxim in an overly narrow way.

Perhaps if we give a more general description, we can obtain the contradiction of the will that we

³⁹ As I’m understanding it, there is a contradiction in conception if it is not possible to act successfully on the maxim in the world in which everyone has that maxim. (But note that the point I will make remains if we go with a stronger interpretation under which there is a contradiction in conception only if the world in which everyone has the maxim is logically impossible.)

⁴⁰ As I’m understanding it, there is a contradiction of the will if willing the universalization of the maxim contradicts something else that you must will as a rational agent.

⁴¹ I say ‘many’ rather than ‘all’, because in free-rider Collective Harm Cases (see note 19 above), we may get a contradiction in conception under some descriptions of the maxim.

are looking for. How can we describe the maxim of refraining from adding a pint more broadly so that the universalization test yields a contradiction of the will? Here is what seems to me to be the best (if not the only good) answer: we should describe the maxim as, ‘refrain from helping those in need.’ Willing the universalization of this maxim produces a contradiction of the will. This, after all, is the maxim of non-beneficence, one of Kant’s own examples of a maxim that fails the contradiction of the will test. But here – just when it seems we have found the contradiction we were looking for – we hit the Superfluity Problem. We can only describe my maxim of refraining from adding my pint as a maxim of ‘refraining from helping those in need’ if my adding the pint *would* help those in need. But the claim that my act won’t make any difference to those in need seems to be precisely a claim that it will *not* help them – that it would, instead, be instrumentally superfluous when it comes to alleviating their suffering.

In any case, we can at least conclude that it is far from clear how the Formula of Universal Law would yield the result that you ought to add your pint to the cart, unless we can first debunk the claim that doing so will not help.⁴²

9. The Problem of Collective Harm (and How to Solve It)

In Collective Harm Cases, one’s individual act (or omission) doesn’t seem to make any difference. It appears that things will be the same in all relevant respects, whether or not it is performed. And what this seems to mean is that acting in this way would be instrumentally superfluous – that it would not itself play any significant role with respect to bringing about the

⁴² Perhaps there is a (correct) way of describing the maxim that avoids this problem. Even if that is so, there is a different level at which the Superfluity Problem might show up. If we are going to appeal to the Formula of Universal Law, we need an understanding of *why* this test tells us what we ought to do. We need to know, in other words, why universalizability matters. It could very well be that the explanation we want to give about why universalizability matters will not work in Collective Harm Cases unless we can first show that one’s individual act is not instrumentally superfluous. For instance, a common idea is that acting in a way that is not universalizable is wrong because it exhibits a certain kind of unfairness: it amounts to relying on others to do something, while excepting yourself. For reasons similar to those given in our discussion of the Fairness Approach, this story might not work in Collective Harm Cases *unless* we can first debunk the impression that one’s act would be instrumentally superfluous. Many thanks to Sergio Tenenbaum for this point.

outcome in question. This is *why* it is hard to see how there could be any point in acting. While one knows that unless enough people act in this way bad consequences will result, one is faced with the apparent *uselessness* of acting in this way oneself.

The three Rejecting the Implication views that we discussed approach this problem by trying to show that, even if acting in the relevant way can't make a difference, there is still reason to do so. They do this by turning the focus *away* from the concern about the apparent instrumental superfluity of your individual act, and onto other sorts of considerations: being part of a group that makes a difference, participation in a collective action, fairness in the context of a collective obligation. But these considerations can't do the work they are supposed to do as long as your act would be merely superfluous. Unless your act could play a significant, non-trivial role with respect to bringing about the outcome in question, the explanations these views are offering as to why you have reason to do it cannot properly address what is at issue in the Problem of Collective Harm. The same problem seems to arise for various non-consequentialist theories, if we think we can turn to them to provide a solution to the Problem of Collective Harm.

In general, to address what is at issue in the Problem of Collective Harm we need to be able to explain how it is that the fact that the men in the desert are suffering from painful thirst means that each of us has reason to add our pints, and that the fact that we are collectively causing climate change means that each of us has reason to take a bicycle to work instead of a car, and so on. That is, we need to give a reason for action that has to do with – in a central way – the morally relevant outcome of concern, and that tells us specifically to do the sort of acts that could collectively cause (or, depending on the case, prevent) that outcome. It is doubtful we can

do this adequately while at the same time accepting that acting in this way would be merely superfluous with respect to bringing about that outcome.

What all of this reveals is that to solve the Problem of Collective Harm, we need to show that one's individual act *isn't* merely superfluous. We need to show that we are mistaken when it seems to us that an individual act of the relevant sort won't itself do anything instrumentally significant. Of course, it's not clear that we can do this. But unless we can, a satisfactory answer to the problem seems unlikely.

If we can show that an individual act of the relevant sort isn't merely superfluous, and rather could significantly help to bring about the outcome in question, this would address the Superfluity Problem that Weak Participation, Strong Participation and the Fairness Approach face. So, that core problem with those views would be resolved. Similarly, the general non-consequentialist theories we considered would then, it seems to me, be able to handle Collective Harm Cases. For instance, if one's turning up the dial on the shock machine (in order to get \$50) might play a significant role in harming the victim, it makes sense to say – as the Kantian would want to – that in doing so one is treating him as a mere means.

But the point isn't just that this would allow these views about our moral reasons in Collective Harm Cases to get off the ground. Rather than just being a matter of filling in the hole that we've identified with these views, showing that an individual act can do something significant toward bringing about the relevant outcome is the *central task* of solving the Problem of Collective Harm. I've argued that it is doubtful that we can solve the problem in a satisfactory way without showing this. But, importantly, if we can explain how it is that your individual act can do something instrumentally significant, this would itself solve the core issue of the Problem of Collective Harm. What leads us to fail to see how we could have reason for action in

Collective Harm Cases *is* the impression that doing so would be useless – that it wouldn't do anything significant toward bringing about the relevant good outcome, or preventing the relevant bad one. If we could show that this were not true, that core issue would be addressed.

Now, what this might seem to mean is that if we want to solve the problem, we need to take a *Denying the Description* approach. That is, we must try to argue that in Collective Harm Cases, contrary to appearances, individual acts do, or at least might, make a difference with respect to the morally relevant outcome of concern. This line of response has largely been taken by those seeking a consequentialist solution to the problem. But, since it would address the Superfluity Problem if it worked, one implication of our discussion is that non-consequentialists also have a stake in whether such a view can succeed.

I discuss the Denying the Description strategy elsewhere. Unfortunately, if my conclusions there are correct, such an approach does not work.⁴³ I don't doubt that in some cases we can and should deny the "it won't make a difference" claim: we are surely sometimes making a mistake when we think that one act won't make a difference. But, while that will surely be true in some cases, I argue that we cannot draw that conclusion in general. We are not, I argue, in a position to deny that there are Collective Harm Cases in which one's individual act simply can't make a difference.

Suppose, for the sake of argument, that I am right about this. It might seem, then, that we have hit a standstill. On the one hand, the upshot of the present paper is that we cannot solve the Problem of Collective Harm *unless* we can maintain that one's individual act can play a significant instrumental role with respect to the outcome of concern. But on the other hand, if my conclusions elsewhere are correct, we cannot deny that there are Collective Harm Cases in which one's individual act simply can't make a difference.

⁴³ "Consequentialism and the Problem of Collective Harm: A Reply to Kagan".

But rather than a standstill, these conclusions together carve out a new space in which a solution to the problem might be found. Instead of taking a Denying the Description approach, we could try to show that an act can play a significant instrumental role in bringing about an outcome, *even if* it cannot make a difference with respect to that outcome. If an act cannot make a difference, we take this to mean that it would be instrumentally superfluous. A way of solving the Problem of Collective Harm would be to show that this assumption is mistaken.

What would it mean for this assumption to be mistaken? Whether one's act makes a difference is a particular, counterfactual matter. It is a matter of whether had one *not* done it, things would have been (relevantly) different. The assumption would be mistaken if this were not the correct test for determining whether an act is instrumentally useful. Perhaps my adding a pint to the cart could play a significant, non-superfluous role in alleviating the men's suffering even though the extent to which their suffering is alleviated will not be different depending on whether or not I do so.⁴⁴

In summary: the upshot of our investigation of Rejecting the Implication approaches is that to respond to the "it makes no difference" argument, we need to be able to show that one's individual act in Collective Harm Cases is not instrumentally superfluous. This is the real challenge of the Problem of Collective Harm. Attempting to address the problem by turning our attention to other sorts of reasons for action will not work. One way of trying to meet the challenge is to take a Denying the Description approach: to try to show that an individual act actually can make a difference in Collective Harm Cases. But that is not the only way. There is a

⁴⁴ I am not claiming that this strategy will work. For the purposes of this paper, the point is just that this strategy exists as a second way – besides denying the "it makes no difference" description – of trying to address the superfluity challenge. I attempt to develop a solution of this sort in "How You Can Help, Without Making a Difference."

Julia Nefsky

second strategy that remains open: we could question whether an act needs to be a difference-maker in order to be instrumentally useful.