

What Is the Point of Helping?

James Fanciullo

Forthcoming in *Philosophical Studies*

Abstract: In some cases, a group of people can bring about a morally bad outcome despite each person's individual act making no difference with respect to bringing that outcome about. Since each person's act makes no difference, it seems the effects of the act cannot provide a reason not to perform it. This is problematic, because if each person acts in accordance with their reasons, each will presumably perform the act—and thus, the bad outcome will be brought about. Recently, Julia Nefsky has argued that this problem is solved by rejecting the assumption that if an act makes no difference with respect to an outcome, then the act cannot do anything non-superfluous toward bringing that outcome about. Nefsky suggests that, even if an act makes no difference, the act may nevertheless help: it may make a non-superfluous causal contribution. If this is right, it means that the potential effects of an act may give us a reason to perform the act, even if the act wouldn't make a difference. In this paper, I offer some reasons to be wary of Nefsky's approach. I first argue that her account generates problematic results in a certain range of cases, and thus that we may have no reason to help in any case. I then argue that, even if we do sometimes have a reason to act when it seems we cannot make a difference, this reason cannot be the one that Nefsky identifies.

This paper concerns a troubling sort of story—a sort of story in which the moral mathematics seem not to add up.¹ Consider

Drops of Water: 10,000 wounded people lie out in the desert, suffering from intense thirst. We are an equally large number of altruists, each of whom has a pint of water. We could pour these pints into a water cart. This would be driven into the desert, and our water would be

¹ See Parfit 1984 (ch. 3).

shared equally between all these many wounded people. By adding our own pint, each of us would enable each wounded person to drink slightly more water—perhaps only an extra drop. The effect of each of these extra drops on each person would be imperceptible.²

If each of us adds our own pint, each person in the desert will receive a full pint of water, and the suffering of each wounded person will thus be relieved. It might seem obvious, then, that this is what each of us ought to do. The problem, however, is that it seems no individual contribution to the water cart will make any difference to the suffering of these wounded people.³ Adding our own pint will only enable each wounded person to drink an extra drop of water. This contribution isn't great enough to be noticed, let alone make any difference to whether any person's suffering is relieved. Since adding our own water simply isn't the kind of act that could have any influence on the wounded people's suffering, it seems that no one of us has any reason to add our own pint. So, while we might hope that each of the other altruists adds their pint, we have no reason to think that anyone with a pint ought to do so.

Obviously enough, if we are at all concerned with the suffering of these wounded people, this conclusion is likely to disappoint. After all, if each of us has no reason to contribute our own water, and each of us acts in accordance with our reasons, it seems no one will ultimately contribute. And if no one contributes their water, the suffering of the wounded people will not be relieved. It seems clear that morality should have something more to say here.⁴ But because no individual person's act makes any difference with respect to the wounded people's suffering, it's unclear what

² This is slightly adapted from Parfit 1984 (p. 76).

³ Of course, some people deny that each contribution would make no difference. See, for example, Kagan 2011 and Norcross 2004. These people, in particular, argue that each individual contribution *might* make a difference. This approach is untenable, however, for reasons Julia Nefsky (2011) provides.

⁴ Some people are willing to “bite the bullet,” however, and claim that collective action cases are merely unfortunate situations—much like coordination problems—where no individual has done anything wrong. See Tannsjo 1989. See also Sinnott-Armstrong 2005, who nevertheless concludes that governments have a responsibility to make a difference in these cases. For reasons why these approaches, among others, are not tenable, see Nefsky 2015.

exactly this something more might be. There is no obvious error in reasoning or flaw of character to which we can attribute our collectively failing to prevent the harmful outcome. We are then left either to give up the intuitive thought that collectively satisfying the demands of morality will not lead us to such harmful outcomes, or else to find some previously unrecognized reason, issuing from morality, to contribute our water to the cart.⁵

Recently, Julia Nefsky has taken the second of these two options, and offered an account on which each of us does, in fact, have a morally significant reason to add our water to the cart.⁶ This reason has gone largely unrecognized, according to Nefsky, because most have assumed “that if an act won’t make a difference with respect to an outcome, then it cannot do anything non-superfluous toward bringing that outcome about. In other words, *helping* to bring about an outcome requires *making a difference*” (Nefsky 2017, p. 2744). If we reject this assumption, we see that there is in fact a reason to contribute in cases like Drops of Water. In particular, there is a reason stemming from the fact that our act could *help* relieve the wounded people’s suffering—that is, our act could make a non-superfluous causal contribution to bringing this morally good outcome about. If this is right, we may be able to avoid the outcome in which the suffering of those in the desert is not relieved. More generally, we may be able to avoid bad outcomes in any case where acts or omissions of a certain type collectively bring about a bad outcome but individually make no difference—or, any *collective impact case*.⁷ Rejecting the assumption that helping requires making a difference thus seems to reveal our reason, for example, to vote in national elections, and to reduce our individual carbon footprints.

⁵ Taking the first option would also force us to reject the plausible thought that collectively satisfying the demands of morality will not lead us to morally suboptimal outcomes. This thought is a rough version of what has elsewhere been called “the principle of moral harmony.” Proponents of the principle include Parfit (1984, p. 54), Pinkert (2015), Portmore (2016), and Regan (1980).

⁶ Nefsky 2017.

⁷ I borrow this terminology from Nefsky (2017, p. 2744). Strictly speaking, though, collective impact cases also include those where acts or omissions of a certain type collectively bring about a *good* outcome but individually make no difference.

My aim in this paper is to explore the implications of Nefsky's account of helping. In the following section, I'll offer a more complete explanation of her account. I'll then argue, in section 2, that her account generates problematic results in a certain range of cases. Finally, in section 3, I'll argue that if we do have a reason to contribute in collective impact cases, this reason cannot be the one that Nefsky identifies.

1. Nefsky on helping

On Nefsky's view, an act needn't make a difference with respect to an outcome in order to make a non-superfluous contribution toward bringing the outcome about. The above line of reasoning is then flawed in its claim that adding one's water isn't the type of act that can influence whether the people's suffering is relieved, since adding one's water might make a non-superfluous causal contribution to the good outcome. That is to say, adding one's water might *help*. Still, it is not entirely clear what, exactly, this claim amounts to.

Consider first the notion of causal contribution. It seems obvious enough that, in cases like Drops of Water, one's act can make a causal contribution to the relevant outcome. Indeed, as Nefsky notes, if one's act couldn't make such a contribution, there would be no problem to begin with: one would not face any choice as to whether one should contribute. The fact that one's act is of a kind that, when collectively performed, causes the good outcome is a basic feature of the case. Thus it seems adding one's water to the cart just is to make the relevant kind of contribution—one's act may be part of what causes the outcome. One's own contribution is not significant enough to be noticed, of course, but still it may be part of what relieves the wounded people's suffering. So it's clear that one's act makes some contribution—just not one that is noticeable. The task for Nefsky, then, is to show that one's contribution here is also non-superfluous.

Consider next then the notion of non-superfluosity. According to Nefsky, your act is non-superfluous with respect to an outcome just in case, at the time at which you act, it is an “open possibility” both that the outcome may come about, and that the outcome may fail to come about due, at least in part, to the fact that not enough such acts have been performed.⁸ Obviously, if the outcome is guaranteed to come about, or guaranteed not to come about, any act performed in the service of bringing about or avoiding the outcome will be merely superfluous. That is why it must be an open possibility both that the outcome may come about, and that it may fail to come about. More specifically, Nefsky thinks, it must be an open possibility that the outcome might fail to come about due, at least in part, to the fact that not enough acts of the relevant type have been performed. To illustrate, Nefsky offers an example in which, after seeing that a parking meter has been paid until six o’clock, a subject adds money despite knowing that parking there after six o’clock is free. When the police come by, they see that the meter has a time on it that is past six o’clock, and so they do not give the subject a ticket. Clearly, this person’s act plays a causal role in bringing about the desired outcome: the police do not give her a ticket because they see that the meter has a time on it that’s past six o’clock. But, equally clearly, this person’s contribution is merely superfluous, since there is no chance at the time they act that the bad outcome will be brought about due to a lack of money added to the meter. It remains true, of course, that when she adds the money, there is a chance that the bad outcome will be brought about—perhaps she has parked over the line, or too far away from the curb—but what makes her act here merely superfluous is that there is no chance of the bad outcome being brought about due, specifically, to a lack of money in the meter. So, this person’s act of adding money could not be an instance of helping, since it was not an open possibility at the time at which they acted that the outcome would not be brought about due, at least in part, to the fact that not enough such acts had been performed.

⁸ Nefsky 2017 (p. 2752).

Reconsider Drops of Water. It is true of each of us that adding our own pint to the cart cannot make a difference with respect to the good outcome. But it is not true of each of us that adding our own pint cannot help, on Nefsky's view. That's because, first, each of our acts is of a kind that can make a causal contribution to the relief of the wounded people's suffering; and second, there is a chance, at the time at which we act, that the suffering will fail to be relieved due at least in part to a lack of such contributions. Accordingly, as Nefsky sees it, while the act of adding our own pint cannot make any difference with respect to the outcome, this act might help bring the outcome about.

On Nefsky's view, then, one's act can help bring about an outcome just in case it can make a non-superfluous causal contribution to bringing the outcome about. More precisely, the view says the following:

Suppose your act of X-ing could be part of what causes outcome Y.

In this case, your act of X-ing is non-superfluous and so could help to bring about Y if and only if, at the time at which you X,

(*) It is possible that Y will fail to come about due, at least in part, to a lack of X-ing.⁹

Nefsky notes that the view contains the following three conditions.

- (1) It is possible that Y will occur.
- (2) It is possible that Y will fail to occur.

⁹ Nefsky 2017 (p. 2753).

- (3) It is possible that Y will fail to occur at least partly as a result of there not having been enough acts of X-ing.¹⁰

So the suggestion is that, if one's act could help bring about a good outcome, this constitutes a reason for performing the act. The problem we face in cases like Drops of Water is then dissolved, since we do in fact have a reason, issuing from morality, to perform the relevant acts. In particular, this reason is that one's act could help bring about a good outcome, which is a reason one will have whenever the above conditions obtain.

Well and good, then each of us has a reason, stemming from the fact that our act could help, to add our water to the cart. Suppose so. Still, we must ask where else these reasons to help might arise. If this account does not get us the right results in other sorts of cases, its success in Drops of Water may be far less significant. So let's turn to consider a few potential such cases.

2. Helping and happenings

The first thing I want to do here is suggest a reason for thinking that Nefsky's account generates the wrong result in a certain sort of case. I'll then argue that if Nefsky can avoid this problematic result, she can do so only by appealing to a claim whose implications are implausible. So, either way, we have good reason to doubt her view.

The relevant problem for Nefsky's view is that it seems to imply that we have a reason to add our water, not just in Drops of Water, but in another sort of case as well. And so if our intuitions about whether we should contribute in this other sort of case are not the same as they are in Drops of Water, her account will appear far less plausible. To see the sort of case I have in mind, consider

¹⁰ Nefsky 2017 (p. 2753).

One Mechanism. Your situation is just as it is in Drops of Water, except now there are no other altruists. Instead, there is a mechanism that will either release 9,999 pints of water into the cart, or it will not. You are just as confident that the mechanism will add the water as you are that, in Drops of Water, enough other altruists will add their water to ensure the relief of the people in the desert's suffering.

Do we have the same intuition—that you have a reason to contribute—in this case? Most people, I suspect, will not. After all, it seems utterly pointless to add your water here, given that whether the good outcome will come about depends only on whether the mechanism will release the water. Your pint would never be noticed if it were added to 9,999 others, and your pint alone would do nothing to relieve the suffering of the people in the desert. Either way, then, adding your pint seems pointless. And so, you do not seem to have any reason to add it.

On Nefsky's account, however, there doesn't appear to be any relevant difference between One Mechanism and Drops of Water. So long as you think there is a good enough chance that the mechanism will release the water, it seems the view tells you to add your pint. This is because, at the time at which you act, it is possible that the suffering of the people in the desert will not be relieved due, at least in part, to a lack of water being added to the cart. So, in Nefsky's terms:

- (1*) It is possible that the suffering of those in the desert will be relieved.
- (2*) It is possible that the suffering of those in the desert will fail to be relieved. And
- (3*) It is possible that the suffering of those in the desert will fail to be relieved at least partly as a result of there not having been enough water added to the cart.

It seems, then, that One Mechanism is in all relevant respects identical to Drops of Water, on Nefsky's view. The odds of success are the same, and the conditions for helping seem to obtain equally in each case. Hence any reason to add your water in Drops of Water seems also to be a reason to add your water in One Mechanism, which means that if you should add your water in the former case, you should also add it in the latter.

This conclusion is surely implausible. But does Nefsky's view actually imply it? After all, it may be objected: it's stipulated that in order to help, it must be the case that the good outcome might fail to occur at least partly as a result of there not having been enough *acts* of a certain type. A mechanism, however, isn't the kind of thing that can act. So, in One Mechanism, it isn't the case that your act can help, because it's not the case that the good outcome might fail to occur at least partly as a result of there not having been enough *acts* of adding water. Put differently, the problem with the above line of reasoning is that, applied to this case, the third condition of Nefsky's view should read:

(3**) It is possible that the suffering of those in the desert will fail to be relieved at least partly as a result of there not having been enough acts of adding water to the cart.

And since, in One Mechanism, this claim is false—the mechanism will unilaterally decide whether enough water will be added—it follows that your act cannot help, on Nefsky's view. And so, it seems, her account faces no worry here.

This response has some force. Notice, however, that it might lead proponents of Nefsky's view to the wrong conclusion in the following case:

9,999 Mechanisms: Your situation is just as it is in One Mechanism, except now there is more than just one mechanism. Instead, there are 9,999 mechanisms, each of which will either release a pint of water into the cart, or it will not. You are just as confident that enough mechanisms will add their water to the cart to ensure the relief of the people in the desert's suffering as you are that, in Drops of Water, enough other altruists will add their water.

Note that adopting the current response would force us to claim that, since a mechanism isn't the kind of thing that can act, the conditions for helping fail to obtain in this case as well. Even if the odds of each mechanism adding its own pint in this case are equal to the odds of each person adding their own pint in Drops of Water, on this response, you have a good reason to add your water in the latter case, but no reason at all to add it in the former. That is to say, you have a good reason to try to help relieve the suffering of the people in the desert in Drops of Water, but no reason at all to try to help relieve their suffering in 9,999 Mechanisms. Thus the mere fact that, in the latter case, the acts of other people will not be the method by which water is added to the cart is enough to make your causal contribution—which would be the same in either case—superfluous, and thereby eliminate any reason you had to help. I'm unsure whether proponents of Nefsky's view will find this result untenable, plausible, or somewhere in between. But, as I see it, it seems implausible that this difference between 9,999 Mechanisms and Drops of Water is significant enough to do the work being asked of it here. While my intuition that you can help in 9,999 Mechanisms is, admittedly, somewhat weaker than my intuition that you can help in Drops of Water, my intuition about 9,999 Mechanisms is not nearly as weak as Nefsky's account implies that it should be. Given you are not (explicitly) cooperating with the other altruists in Drops of Water, it does not seem all that important that the pints might be contributed by these actors, as opposed to some mechanisms. Either way,

your act makes an identical causal contribution, and the odds of success are the same. So, we may wonder, what's the difference?

Of course, a proponent of Nefsky's view might insist that the distinction between acts and mechanical happenings is weighty enough to separate cases where some act can help from cases where it cannot. But, if we assign to each mechanism a probability of adding water that corresponds to the probability of adding water that we assign to each altruist, I do not see how this distinction could lead us to conclude that, even though in each case your act would play the same causal role, in one case the act is pointless, and in the other you have good reason to perform that same act. That is, if we acknowledge that the odds that Mechanism 1 will contribute are the same as the odds that Actor 1 will contribute, and the odds that Mechanism 2 will contribute are the same as the odds that Actor 2 will contribute ... and the odds that Mechanism 9,999 will contribute are the same as the odds that Actor 9,999 will contribute, it is unclear to me how we could plausibly conclude that whether we have any reason to perform what would in any case be the same act—and whether we are to count as superfluous what would in any case be the same causal contribution—turns merely on which of these two types of potential contributor the others turned out to be. As I see it, this possible response is more implausible than the claim that we have no reason to act in either case. If this is right, it means that the fact that our act could help-without-making-a-difference gives us no reason to act, and that Nefsky's view should therefore be rejected.

But, of course, you may think that this is not right—that these mechanism cases do nothing to show that Nefsky's view has any unacceptable implications. Then consider one more case:

Raised Stakes: Your situation is just as it is in 9,999 Mechanisms, except now every person on earth (except you) is in the desert, and these people will all die if they do not receive enough

water to relieve their intense suffering. For each additional person, there is an extra pint, and these pints are distributed equally among you and the mechanisms.

Again, the current response forces us to claim that, since mechanisms are incapable of acting, the conditions for helping fail to obtain in this case. Thus even if the fate of the world depends on whether enough water is added to the cart, you have no reason to add your water, on Nefsky's view. And this is explained just by the fact that the other potential contributions would not result, specifically, from acts.

As may be expected, I do not find it all that plausible that you both have a reason to help in Drops of Water, and have no such reason in Raised Stakes. If I concede that you have a reason to help, or that your act can make some appreciable causal contribution, in one, it seems to me more plausible that you have a reason also to help, or that your act can also make that contribution, in the other. And if I do not concede this, it seems to me more plausible that you have no reason to help, and that your act cannot make that contribution, in either. That is because I am primarily concerned—perhaps too narrowly, perhaps not—with *your* act, and what *its* effects would be. The odds of the other potential contributors adding their water no doubt seems relevant to this: if no one else contributes, your act is certain to do nothing. Their identity, on the other hand, seems beside the point. This irrelevance seems obvious in cases where the numbers are smaller: if, for example, two sufferers needed slightly more than half a pint of water each, and just you and one other potential contributor had pints, your reason to contribute would not seem to depend on whether this other potential contributor was a mechanism or an actor. All that would seem to matter is how likely they were to contribute. In the same way, once I concede that your act is likely enough to have certain effects, or play a certain causal role, in Drops of Water, I am immediately inclined to think that it is likely enough to have corresponding effects in Raised Stakes. And, naturally, when I

do not concede this, I am immediately inclined to think that it is not likely enough to have any relevant effects in either case. My intuitions are not decisive, of course, but they at least suggest that Nefsky's view may be getting us the wrong result in at least one of these cases. If this suggestion proves correct, it seems we will have good reason to reject her approach.

In any case, what I ultimately find implausible is not just the implication that there is this difference in our reasons, but the explanation that Nefsky's account provides for it. The fact that the other potential contributions would result from acts only in Drops of Water does not seem, to me, to properly explain this difference. Providing an argument for this thought is the task of the following section.

3. Helping by accident

You may at this point remain unconvinced by my claims about the mechanism cases. You may think that we surely have no reason to contribute in these cases, and that we surely have such a reason in Drops of Water. What I now want to argue is that, even if that is right—even if we sometimes have a reason to contribute, and the identity of the other potential contributors matters to whether we have it—our reason to contribute still cannot be the one that Nefsky identifies.

Suppose then that we have no reason to contribute our water to the cart in 9,999 Mechanisms, despite our having such a reason in Drops of Water. What would explain this difference? According to Nefsky, it is the fact that the other potential contributors in 9,999 Mechanisms cannot act, and it is therefore not possible that the good outcome might fail to come about due, at least in part, to a lack of acts of adding water to the cart, that explains the difference. Since the other potential contributors in Drops of Water are capable of acting in this way, your act there might help. And since the other potential contributors in 9,999 Mechanisms are not capable of acting in this (or any) way, your act there cannot help.

I do not think that this explanation can be correct, however. As I see it, whether the other potential contributions would be the result of acts does not seem like the kind of difference that could explain our having a reason to contribute in one case and not the other. It seems to me more likely that this explanation would have more to do with the fact that the other contributors lack certain attitudes—that it is in virtue of some feature of our shared attitudes, or shared agency, that we each have a reason to contribute—than it would have to do with the fact that their contributions would not be the result of acts.

This suggestion is, I must admit, at this point largely speculative. But I would like to briefly offer an argument in support of it, just to get the idea on the table. So consider

Slips on Water. Your situation is just as it is in Drops of Water, except now the altruists are not altruistic. They do not care in the least whether the suffering people receive enough water, and they have no intention of contributing their pints to the cart. Happily, though, these people are standing on a very slippery floor, and there is a good chance that when each goes to announce that they will not be contributing, they will accidentally slip, fall, lose hold of their water, and propel it into the cart. You are just as confident that enough of these people will accidentally add their water to the cart in this way to ensure the relief of the people in the desert's suffering as you are that, in Drops of Water, enough other altruists will add their water.

Should you contribute? Intuitions here may diverge. But let me report mine. As someone who does not see much of a difference between Drops of Water and 9,999 Mechanisms, I also do not see much of a difference between Drops of Water and Slips on Water. Still, the differences I do see are similar: the claims that you should contribute in 9,999 Mechanisms and Slips on Water seem

less plausible than the claim that you should contribute in Drops of Water to roughly the same degree. That is, as I see it, the claim that you should contribute in Slips on Water seems less plausible than the claim that you should contribute in Drops of Water to roughly the same degree that the claim that you should contribute in 9,999 Mechanisms seems less plausible than the claim that you should contribute in Drops of Water. We can agree on this much even if we disagree about how plausible these claims about 9,999 Mechanisms and Slips on Water are. If this pecking order of plausibility seems correct, it suggests that our reasons (or lack thereof) to contribute in 9,999 Mechanisms and Slips on Water should be fairly similar.

Yet, on Nefsky's account, they are not. As we have seen, her account implies that you have no reason to contribute your water in 9,999 Mechanisms. But it also seems to imply that you do have such a reason in Slips on Water. After all, when these people accidentally contribute their water, they are performing *acts* of adding water to the cart (albeit unintentional ones). It is then true that, at the time at which you act, it is possible that the suffering of the people in the desert will not be relieved due, at least in part, to there not having been enough acts of adding water to the cart. So, again in Nefsky's terms:

- (1*) It is possible that the suffering of those in the desert will be relieved.
- (2*) It is possible that the suffering of those in the desert will fail to be relieved. And
- (3**) It is possible that the suffering of those in the desert will fail to be relieved at least partly as a result of there not having been enough acts of adding water to the cart.

The condition that was not met in 9,999 Mechanisms then certainly seems to be met here. Since Slips on Water is in all other relevant respects like Drops of Water, it seems Nefsky's view implies that your reason to contribute in one will be no different from your reason to contribute in the

other. And so, it implies that you have a reason to add your water to the cart in Slips on Water.

Nefsky's view thus reaches contrasting results in cases that intuitively seem quite similar.

This is, of course, no knockdown argument against Nefsky's view. It is simply an appeal to intuition. But be that as it may, it seems to me to get at something important. In particular: it seems to suggest that, if the identity of the other potential contributors matters to whether we have a reason to contribute in cases like Drops of Water, it matters only insofar as it identifies whether these other potential contributors possess certain intentions. Thus, on this picture, if we have no reason to contribute in 9,999 Mechanisms, this is not because the other potential contributors are incapable of action; rather, it's because they are incapable of intention. Similarly, if we have no reason to contribute in Slips on Water, this is not because the other potential contributors are not performing a certain action—on the contrary, they are—but is instead because they fail to form the intention to perform that action. These judgments are evidenced by the fact that the only relevant difference between Drops of Water and Slips on Water is that, in the latter case, the subjects do not intend to add their water to the cart. So, just so long as it seems implausible that we have a reason to contribute our water in Slips on Water, the most plausible explanation as to why we have no reason to contribute in either this case or the mechanism cases appears to be that our reason to contribute in collective impact cases depends on the intentions of the other potential contributors.

Importantly, notice that the conclusion of this argument does not depend on any claims or intuitions about the mechanism cases. My aim in drawing a comparison between Slips on Water and 9,999 Mechanisms above was to show that while something seems similarly “off” about these two cases, Nefsky's account reaches distinct verdicts about them. The more important point, however, is that her account implies that our reason to contribute in Slips on Water is no different from our reason to contribute in Drops of Water. This means that, so long as it seems even *slightly* less plausible that we should contribute in Slips on Water than it seems plausible that we should

contribute in Drops of Water, proponents of Nefsky's account face a significant worry. Short of giving up their view entirely, it seems they must either offer some reason for thinking that our intuition here may in fact be untrustworthy, or else amend their view to account for the intentions of others in determining our reason to help.

Now, admittedly, this may not seem all that difficult a problem to solve. After all, it may be thought, all that proponents of Nefsky's view must do is amend the third condition of their view. More specifically, they must simply amend this condition to claim that helping requires, not that the other potential contributions would be the result of acts, but that the other potential contributions would be the result of intentional acts. This amended view would then imply that you have no reason to help in Slips on Water, since the view's third condition would not be satisfied: it would not be true that the good outcome might fail to come about due at least in part to a lack of intentional acts of adding water to the cart. That's because yours would be the only intentional such act. Accordingly, just as Nefsky's view implies that you have no reason to help in 9,999 Mechanisms, so this revised version of her view would also imply that you have no reason to help in Slips on Water. And so it appears my objection to Nefsky's view is easily avoided.

In fact, however, I think that my worry runs deeper than this response suggests. Notice that, on this amended view, the explanation as to why we have no reason to contribute in Slips on Water is that the other contributions would not result from intentional acts. But this explanation seems, to me, rather unsatisfying. That's because it is only derivatively concerned with the other potential contributors' intentions to contribute. On this way of seeing things, whether the other potential contributors have the relevant intentions matters only insofar as it tells us something about their potential acts—in particular, whether the acts would be of the required, intentional kind. There is no deeper fact about their intentions in virtue of which we have a reason to contribute in Drops of Water and have no such reason in Slips on Water. Instead, it is a fact about their actions that

explains this difference: namely, the fact that their actions are intentional only in Drops of Water. It is unclear to me, however, why it should matter whether their acts would be intentional, unless there is some deeper fact about their intentions that explains this difference in our reasons. To put the point here slightly differently: because of its inherited focus on the other potential contributors' acts, the revised version of Nefsky's view does not provide an adequate explanation as to why we have a reason to contribute in collective impact cases. It does not tell us what it is about the other potential contributors' intentions that makes our reason to contribute depend upon those intentions, or what it is about them that makes our own contribution non-superfluous. It is simply required, for your act to help, that the others have these intentions. That, I think, is an unsatisfying explanation. To adequately defend their view, then, proponents of Nefsky's view must not only amend it to account for the intentions of others in determining our reason to help, but also offer some explanation as to why the intentions of others should matter to whether we have any such reason. Though we've seen that the first task can be easily accomplished, success in the second may prove considerably more difficult. So long as these proponents remain focused on the other potential contributors' acts, rather than on the intentions to perform those acts, the significance of the latter seems likely to remain opaque.

In what way do our reasons to contribute in collective impact cases depend on the intentions of others? That is a topic for another paper.¹¹ Besides, it is not something that I must take a stand

¹¹ But here is an abridged version of my likely, again speculative response. Our reason to contribute in Slips on Water seems weaker because the other potential contributors are unwilling to add their water to the cart (and in the mechanism cases, it is because they are incapable of so willing). Since a willingness to act is, I think, no more than an intention to act under certain conditions, or a conditional intention to act, it seems that our reason to contribute partly depends on whether the others have a conditional intention to add their water to the cart. Conditional on what? Well, given the aim of cooperating is to produce the good outcome, it seems the intention must be conditional on whether the collective set of acts would produce this outcome. Thus the intention must be to add one's water if and only if enough others so intend—enough, that is, to ensure that the suffering of those in the desert will be relieved. In this way, one's own intention would be conditional on the intentions of the others: if none of them have formed an analogous intention—as in Slips on Water—one would have no reason to act on one's own intention; and if enough of them have formed an analogous intention—as, perhaps, in Drops of Water—one would have such a reason. (Here I have in mind something similar to what Michael Bratman (2014) calls a shared intention involving obligation-based interdependence.) Of course, none of this proves that we have any reason to form or act on conditional intentions of this sort in collective impact

on: recall that I remain uncommitted to the claim that we have any reason to contribute in collective impact cases. I have offered this argument just to convince those who have rejected my arguments in the previous section of the fact that, if we have a reason to contribute in collective impact cases, this reason is not properly identified by Nefsky's account. Contrary to what she suggests, the relevant difference between Drops of Water and the mechanism cases—if there is one—is not that the other potential contributions would be the result of acts in only the first case. It is instead that the other potential contributors would have certain intentions in only that case. While her view can be revised to accord with this fact, it remains unclear whether it can also be revised to adequately explain it.

4. Conclusion

Let me sum up. Julia Nefsky has argued that, in certain cases where any individual act would make no difference with respect to a good outcome, but many such acts are needed to bring the outcome about, each person will have at least some reason to perform their individual act—specifically, that their act may help bring the good outcome about. But for all that, it seems, Nefsky's account has at least two problematic implications. As we have seen, her account says that our reason to help in collective impact cases partly depends on the identity of the other potential contributors. If these potential contributors are not capable of action—or, more precisely, if their potential contributions would not result from acts—we can have no reason to help, on Nefsky's view. This, however, implies that even if the fate of the world depended on whether enough contributions were made, and the odds of each contribution being made were the same as they were in Drops of Water (a case

cases. But, at the very least, I think that it provides the best explanation of what our intuitions are tracking in these cases. It suggests that our reason to contribute seems stronger in Drops of Water because the condition that must be met for us to act may have been met, and that our reason to contribute seems weaker in both Slips on Water and the mechanism cases because this condition has surely not been met.

where we would have a reason to contribute, on this view), we still could have no reason to contribute so long as—and simply because—the other contributions would not result from acts. Furthermore, even if this implication can be adequately defended, there is a more pressing worry to address. On Nefsky's view, our reason to help in collective impact cases partly depends on whether the other potential contributions would be the result of acts. Yet, as we have seen, what seems to matter to our reason here is not whether the other potential contributions would result from acts, but whether the other potential contributors would have certain intentions. Nefsky's account thus problematically implies that we would sometimes have a reason to contribute even if none of the other potential contributors had any intention of contributing themselves. And, more fundamentally, the account does not seem apt to adequately explain why the other potential contributors' intentions should matter to whether we have any reason to help. This means that our reason to contribute in collective impact cases—if we have one—is evidently not the one that Nefsky identifies. And this means that we have at least two good reasons for rejecting her approach.¹²

¹² Many thanks to an anonymous reviewer for truly helpful comments. And special thanks to Cheshire Calhoun and Doug Portmore for extremely helpful comments, discussions, encouragement, and advice.

References

- Bratman, M. E. (2014). *Shared Agency*. Oxford: Oxford UP.
- Kagan, S. (2011). "Do I Make a Difference?" *Philosophy and Public Affairs* 39: 105-141.
- Nefsky, J. (2011). "Consequentialism and the Problem of Collective Harm: A Reply to Kagan." *Philosophy and Public Affairs* 39: 364-395.
- . (2015). "Fairness, Participation, and the Real Problem of Collective Harm" In M. Timmons (ed.), *Oxford Studies in Normative Ethics*, Vol. 5. Oxford: Oxford UP.
- . (2017). "How You Can Help, Without Making a Difference." *Philosophical Studies* 174: 2743-2767.
- Norcross, A. (2004). "Puppies, Pigs, and People: Eating Meat and Marginal Cases." *Philosophical Perspectives* 18: 229-245.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford UP.
- Pinkert, F. (2015). "What If I Cannot Make a Difference (and Know It)." *Ethics* 125: 971-998.
- Portmore, D. W. (2016). "Maximalism and Moral Harmony." *Philosophy and Phenomenological Research*. doi: 10.1111/phpr.12304.
- Regan, D. (1980). *Utilitarianism and Co-operation*. New York: Oxford UP.
- Sinnott-Armstrong, W. (2005). "It's Not My Fault: Global Warming and Individual Moral Obligations." In W. Sinnott-Armstrong and R. Howarth (eds.), *Perspectives on Climate Change*. Elsevier: 221-253.
- Tanssjo, T. (1989). "The Morality of Collective Actions." *Philosophical Quarterly* 39: 221-228.