



The problem of collective impact: why helping doesn't do the trick

Andrea S. Asker^{1,2} 

Accepted: 21 May 2023 / Published online: 17 June 2023
© The Author(s) 2023

Abstract

Collective impact cases are situations where people collectively bring about a morally significant outcome by each acting in a certain way, and yet each individual action seems to make no, or almost no difference to the outcome. Intuitively, the beneficial or harmful outcomes give individuals moral reason to act (or refrain from acting) in collective impact situations. However, if the individual action does not make a difference to the outcome, it is not clear what those moral reasons are. The problem of collective impact is the challenge of identifying such moral reasons. Julia Nefsky has presented an account of how an individual action can help without making a difference – call it the Helping Account – that claims to provide a general solution to the problem of collective impact while avoiding problems faced by previously suggested solutions. I present an internal critique of Nefsky's work. First, I argue that, based on the problems that Nefsky has raised against previously suggested solutions, three success conditions for a general solution to the problem of collective impact can be formulated: The Weightiness condition, the Generalizability condition, and the Connectedness condition. Second, I argue that the Helping Account fails to satisfy the three success conditions, thereby failing, by Nefsky's own standards, to provide a general solution to the problem.

Keywords Moral reasons · Collective impact · Difference-making · Helping · Julia Nefsky

✉ Andrea S. Asker
andrea.asker@philosophy.su.se

¹ Stockholm University, Stockholm, Sweden

² Institute for Futures Studies, Stockholm, Sweden

1 Introduction

Collective impact cases are situations where people collectively bring about a morally significant outcome, beneficial or harmful, by each acting in a certain way, and yet each individual action seems to make no, or almost no difference to the outcome.¹ In other words, the outcome seems to be the same (or at least not relevantly different) regardless of any one individual action. Common examples of collective impact situations are greenhouse gas emissions contributing to climate change, factory-farmed meat consumption, and voting in large elections. Intuitively, the beneficial or harmful outcomes give individuals moral reason to act (or refrain from acting) in collective impact situations.² However, it can be argued that this intuition is inconsistent with the notion that no individual action makes a difference to the outcome in question (McPherson, 2021). What moral reasons to act, stemming from the morally significant outcome, does an individual have if her action will not make a difference? The problem, which, following Nefsky (2017), I will refer to as the problem of collective impact, is to identify such moral reasons.³ I will also follow Nefsky (2017) in focusing on individual action in momentary choice situations, rather than individual behavior over time. Hence, the challenge here is to identify a moral reason for a single individual act.

Julia Nefsky (2015, 2021) suggests that previously proposed solutions to this problem can be roughly divided into two categories: instrumental approaches and non-instrumental approaches.⁴ Accounts in the instrumental category appeal to the instrumentality of each individual action with respect to the outcome. Generally, these accounts deny, in one sense or another, that an individual action will not make a difference to the outcome. Perhaps most prominently, accounts that appeal to expected utility hold that each individual action has some chance of making a significant difference to the outcome, and thus the expected utility of each action is such that it generates a moral reason for the individual to act (e.g. Broome (2019), Kagan (2011)). Nefsky's (2017) original proposal also falls in the instrumental camp. It argues that, even if an individual action does not make a difference, it can still play a relevant instrumental role with respect to the outcome by *helping* to bring it about – call this the Helping Account.

Accounts in the non-instrumental category either accept that an individual action will not make a difference to the outcome, or do not concern themselves with whether

¹ Depending on whether the morally significant outcome is beneficial or harmful, the concern will be either with bringing about the beneficial outcome, or preventing or mitigating the harmful one. I will refer mainly to the bringing about of outcomes, but this should be interpreted broadly to include whichever option fits the context.

² An individual action can be to do something or to refrain from doing something. Henceforth, I will use action to mean either doing or refraining – whichever is appropriate in the given context.

³ The problem is referred to in the literature by various names such as “the problem of collective harm” (e.g. Nefsky (2012)) and “the inefficacy problem” (e.g. Budolfson (2019)). It also has various stipulations, but it generally concerns the inefficacy of individual actions with respect to outcomes in collective impact situations, and the implications of this inefficacy for the individual's obligations and moral reasons in the collective impact context.

⁴ This distinction is not to be equated with the distinction between consequentialist and non-consequentialist approaches.

it makes a difference. Instead, they appeal to other considerations in order to identify moral reasons for individual action (Nefsky, 2015). Fairness accounts, for example, hold that there is a collective obligation to bring about beneficial outcomes or prevent harmful outcomes in collective impact situations, and that it is unfair for some not to do their part in fulfilling this obligation while others do (see e.g. Cullity (2000) for a representative version of this view). Participation accounts argue that an individual's action in a collective impact situation makes them a participant (or member) in a group that brings about or contributes to a beneficial or harmful outcome, and such participation is reason-giving. Parfit (1986), Kutz (2000), and Wieland and van Oeveren (2020) each present accounts that fit roughly into the participation category.

Nefsky (2015, 2017, 2019) has argued that leading previous accounts face problems that render them unsuccessful in providing a general solution to the problem of collective impact, but that her own proposed solution avoids these problems and succeeds in providing a general solution.⁵ In this paper, I present an internal critique of Nefsky's work. First, I suggest that, based on the problems that Nefsky raises against previous accounts, three success conditions for a general solution to the problem of collective impact can be formulated: The Weightiness condition, the Generalizability condition, and the Connectedness condition. Second, I argue that the Helping Account fails to satisfy the three success conditions, and therefore fails to provide a general solution to the problem of collective impact by Nefsky's own standards.

While I do not argue that these conditions are in fact conditions for a successful solution to the problem of collective impact, I do think that they are *prima facie* plausible, and that they can at least be viewed as candidate success conditions for a general solution to the problem. However, for those who are reluctant to accept these conditions, it is enough to view them simply as tools for representing Nefsky's criticism against previously suggested solutions, and highlighting central aims of her own account, for the purpose of the internal critique.

In Sect. 2, I give a brief overview of the Helping Account. In Sect. 3, I formulate the Weightiness and Generalizability conditions based on Nefsky's criticism against expected utility accounts. In Sect. 4, I show that the Helping Account does not satisfy the combination of Weightiness and Generalizability. In Sect. 5, I defend the arguments in section three against an objection that appeals to the strength of helping-based moral reasons. In Sect. 6, I formulate the Connectedness condition based on a problem that Nefsky raises against some non-instrumental accounts. In Sect. 7, I show that the Helping Account does not satisfy Connectedness. Section 8 concludes.

⁵ Note that the Helping Account does not purport to identify the only relevant moral reason for individual action in collective impact cases. The account is consistent with there being other important moral reasons. However, the account does claim to "address the inefficacy challenge at the level of reasons" (Nefsky, 2019, p. 11), which I take to mean that identifying the helping-based reasons is supposed to be sufficient for solving the problem of collective impact as it is stipulated here.

2 The helping account

Unlike the other accounts in the instrumental category, Nefsky's (2017) Helping Account does not focus on difference-making. It argues that the instrumentality of an action does not hinge on the action's ability to make a difference. Rather, an individual action can *help*, meaning it can make a non-superfluous, causal contribution with respect to the outcome (and thereby be instrumentally relevant), without making a difference.

More specifically, the Helping Account holds that an individual act of a certain type, an individual act of X-ing, can help to bring about an outcome, Y, if and only if, at the time when the individual performs the action, the following conditions hold:

(1) It is possible that Y will fail to occur at least in part due to an insufficient number of acts of X-ing, and.

(2) It is possible that Y will occur at least in part due to a sufficient number of acts of X-ing.⁶

When both conditions are met, the individual action is non-superfluous and so could help to bring about the outcome, and this is what yields a moral reason for the individual to act.⁷ In other words, as long as it is "up in the air" whether Y will occur (or to what degree it will occur), and an (in-)sufficient number of acts of X-ing could be part of what determines its occurrence, there is a helping-based moral reason for the individual to act. However, if the outcome is already guaranteed (to occur or to not occur), or it is not possible that acts of X-ing will be part of what brings the outcome about, then the individual action is superfluous and there is no helping-based moral reason to do it (Nefsky, 2017, pp. 2753–2754).

Note that in many collective impact cases, a morally significant outcome will occur no matter what we do, but, depending in part on collective action by individuals, the outcome might be more or less severe (or more or less beneficial). For example, climate change will continue to occur regardless of collective action by individuals, but it can be mitigated to various degrees depending in part on what individuals collectively do. In the following, the notion of "bringing about an outcome" should be interpreted as referring to bringing about a change that is substantial enough to be morally significant.

To better understand what helping is, it is useful to note how the Helping Account distinguishes between three different ways in which an action can stand in instrumental relation to an outcome: being part of the cause of the outcome, making a difference to the outcome, and helping to bring about the outcome. Nefsky plausibly suggests that an action can be part of the cause of an outcome without making a difference to it. For instance, in the case of voting, individual votes are part of what causes a

⁶ Nefsky does not list (2) as one of the conditions, but she does implicitly stipulate it (Nefsky, 2017, p. 2754 fn. 23).

⁷ Note that when the conditions for helping are fulfilled, the individual action *is* non-superfluous and *could* thereby help, but whether the action will in fact help depends on what actually happens. An action does not help if the outcome in question does not occur or if the individual action is not part of the cause of the outcome (Nefsky, 2017, p. 2754). However, it is the non-superfluity of the individual action and the potential to help that is reason-giving on the Helping Account.

candidate to win, even if the election does not turn on any one vote (Nefsky, 2017, p. 2750).

Moreover, an action can be part of the cause of an outcome without helping to bring it about, that is, the action can be a *superfluous* part of the cause as opposed to a *non-superfluous* part of the cause (Nefsky, 2017, p. 2751). To see this, take the following example: A large group of people are on a sinking ship. They each have a bucket and are trying to scoop out enough water before it is too late to prevent the ship from sinking. However, there are pumps on board that will start automatically and drain enough water just before it is too late (suppose the pumps are infallible). If the people manage to scoop out enough water in time, an individual act of scooping water can be part of what causes the prevention of the sinking, but since the pumps would have started anyway, the outcome was already guaranteed, so the individual action was causally superfluous. An action can thus be part of the cause of an outcome without making a difference or helping. Finally, Nefsky thinks that if an action makes a difference to the outcome, it also helps to bring it about, but the reverse need not apply. Hence, an action can help without making a difference (Nefsky, 2017, p. 2746).

3 The weightiness and generalizability conditions

With the Helping Account, Nefsky aims to provide a general solution to the problem of collective impact, one that will identify a moral reason for individual action in a range of collective impact cases that we are concerned with (2017, 2019, Sect. 5.2). Nefsky thinks that the moral reason identified should be weighty enough not to be insignificant and easily outweighed by considerations such as minor advantages in pleasure or convenience to the individual (2017, pp. 2754, 2764). The relevant range of cases does not include all situations that could be characterized as collective impact cases. For instance, it does not include situations where a given benefit or harm is inevitable, e.g. if catastrophic climate change were inevitable and it were no longer possible to mitigate it by reducing greenhouse gas emissions, this would not be considered a collective impact situation of the relevant kind (Nefsky, 2021). In general, the collective impact cases of concern are those where collective action by individuals⁸ can bring about (or at least play a part in bringing about) a morally significant outcome, but no individual action seems to make a difference to that outcome, and where real-world examples include voting, greenhouse gas emissions, meat consumption, and typical charity donations to poverty relief.⁹

Nefsky (2011, 2017, 2019) has argued that expected utility accounts are not suitable for providing a general solution to the problem of collective impact because they fail to identify a weighty enough moral reason in some collective impact cases of

⁸ The term “collective action by individuals” simply refers to the set of relevant individual actions. It can correspond to a structured or an unstructured collective. In the real-world collective impact cases of concern, it is typically a matter of an unstructured collective.

⁹ I will not attempt to give an exact definition of the relevant range of collective impact cases, but it at least includes real-world situations with the relevant structure that we are generally concerned with, and where we judge that there are moral reasons for individual action.

concern. The expected utility accounts claim that it is a mistake to think that an individual action does not make a difference to the outcome, because each action at least makes an *expected* difference, and thus a difference in expected value. Recall that the general argument on the expected utility approach is that each individual action has some chance of making a significant difference to the outcome, and therefore the expected utility of each action is such that it yields a moral reason for the individual to act.¹⁰ For instance, Kagan (2011) illustrates how this reasoning can be applied in the case of meat consumption and animal suffering, while Hiller (2011) and Broome (2019) each propose an expected utility account with respect to climate change.

Nefsky (2019, Sect. 4) contends that providing a general expected utility-based solution requires arguing that there is always a chance of one's individual action making a difference to the outcome in collective impact cases,¹¹ and that this in turn requires arguing that all collective impact cases have a form of threshold structure where, for some number of individual actions of a certain type, there is a threshold where one additional such action will trigger a change in the outcome (Kagan, 2011; Nefsky, 2019, Sect. 4.1). Even if the thresholds are few and far between, each individual action has some small chance of being the triggering action and thereby making a difference. It has been debated whether all collective impact situations have this threshold structure, or whether there can be non-threshold cases that instead have a gradual structure with no sharp boundaries (Barnett, 2018). Some argue that collective impact cases must have a threshold structure on pain of paradox or inconsistency (Barnett, 2018; Carlson et al., 2021; Kagan, 2011). Nefsky (2011, 2019) disagrees and argues that attempts to establish that collective impact cases cannot have a non-threshold structure fail.¹²

Since she thinks it is a controversial and ambitious project to deny the existence of non-threshold cases, and since she is not convinced by arguments against the existence of non-threshold cases, Nefsky thinks that this undermines the generalizability of expected utility-based solutions (2012, Sect. 3.13, 2019, Sect. 4). That is, if the expected utility approach will not identify a moral reason for individual action in a collective impact case unless it has a threshold structure, and there are relevant cases that do not have a threshold structure, then the expected utility approach will not identify a moral reason in the full range of relevant cases.

Nefsky (2019, Sect. 4) goes on to suggest that even if we assume that there is always a non-zero chance of making a difference, another problem remains – namely one concerning Weightiness. In large collective impact cases, such as factory-farmed

¹⁰ The expected utility of an action is taken to be the sum of the weighted utilities of each possible outcome, where the weights are the respective probabilities of each outcome occurring, conditional on the action. Hence if the disutility (utility) of a particular outcome (e.g. the individual action making a significant difference) is very high (a large negative (positive) number), then the expected disutility (utility) of an action can be great despite the probability of that outcome being very low (Briggs, 2019).

¹¹ It is not clear from this whether Nefsky claims that the argument must be that there is always an objective chance of making a difference or that there is always a subjective chance, but we can plausibly assume that in this context the argument refers to rational subjective probabilities.

¹² The debate about threshold vs. non-threshold structures goes beyond the problem of collective impact and involves discussions on topics such as vague boundaries, Sorites paradoxes, and the reliability of experiential reports, which are beyond the scope of this paper. Therefore, I will not delve further into this debate here.

meat consumption and voting, there is (we now assume) a small chance that an individual action, e.g. a meat purchase or a vote, will make a great difference in terms of animal suffering or the outcome of an election (see e.g. Kagan (2011) and Parfit (1986)). In other cases, there may be a significant chance of making a very small difference. Consider a stylized example:

Drops of water A group of 10.000 people are stuck in the desert, suffering from severe thirst. Another group of 10.000 people are at the edge of the desert, each holding a pint of water. Each person at the edge of the desert is facing a choice of whether to donate their pint by pouring it into a cart that is to be driven into the desert where the water will be distributed equally among the suffering people. For each pint that is donated, each person in the desert will get an extra drop of water ($1/10.000$ of a pint), an amount so small that it will not make any perceptible difference to their suffering.¹³ Hence, an individual donation does not seem to make a difference to the outcome, that is, to the (degree of) relief of suffering; the people in the desert will not suffer any more or less give or take one pint (adapted from Nefsky (2017, p. 2751) and Parfit (1986, pp. 76–77)).

If Drops of Water has a threshold structure, it is reasonable to think that it would have many thresholds where an additional drop of water would trigger a small increase in the relief of suffering for some thirsty person.¹⁴ The greater the number of thresholds, the greater the chance of hitting one of them and thereby making a small difference through one's individual action (Nefsky, 2019, Sect. 4.2). Both the small-chance-great-difference scenario and the significant-chance-small-difference scenario seem promising for generating an expected utility-based moral reason for individual action. Yet, Nefsky (2011, 2019), as well as Budolfson (2019), argue that, in both types of cases, the expected utility calculation might not come out in favor of the individual act in question. For instance, depending on empirical facts about the marketplace and supply chains, the expected value of a meat purchase might not be negative, given that the chance of triggering additional production is very small and given that the consumer enjoys eating meat (Budolfson, 2019). Similarly, a chance of making a tiny difference in suffering for the thirsty people in the desert might be outweighed by something like the individual's interest in keeping the pint of water for herself.

If this is correct, the expected utility approach will not yield the moral reasons we are looking for in some relevant collective impact cases. It will either identify no moral reason at all, or an insignificant, easily outweighed reason, depending on how we look at it. On one view, expected utility accounts will conclude that there is no expected utility-based moral reason for individual action in these cases. For instance, if the overall expected utility of a meat purchase is positive, then there is no expected

¹³ Parfit (1986, pp. 76–77) points out that, since the benefit from receiving water is merely the relief of suffering from thirst in this scenario, and not some improvement to the people's health, these benefits must be perceptible in order to be benefits at all.

¹⁴ As opposed to there being few thresholds where an additional drop would trigger a large increase in the relief of suffering for some thirsty person.

utility-based moral reason to refrain from the purchase. But we could also look at it in the following way: the meat purchase comes with a small expected harm, so small that it is outweighed in the utility calculation by the expected benefit to the individual of eating meat. However, we might still say that this small expected harm yields a moral reason to refrain, but that this reason is insignificant and outweighed by the small (but greater) expected benefit to the individual.

Based on this criticism against the expected utility approach, as well as Nefsky's aims for her own proposal, the first two success conditions for a general solution to the problem of collective impact can be formulated:

Weightiness The successful account should identify a moral reason that is weighty enough not to be insignificant and easily outweighed by considerations such as minor advantages in pleasure or convenience to the individual.

Generalizability The successful account should identify a weighty enough moral reason in all the collective impact cases of concern.

A proposed solution to the problem of collective impact satisfies the combination of Weightiness and Generalizability if and only if it identifies a weighty enough moral reason for individual action in the full range of relevant collective impact cases. An account that identifies a moral reason in some but not all relevant cases fails with respect to Generalizability.¹⁵ An account that identifies a moral reason in all relevant cases, but that identifies a *weighty enough* reason in none or only in some of these cases, fails with respect to Weightiness. Since a general solution is supposed to solve the problem across the board, it is not enough for an account to identify the right kind of reason only in some subset of relevant cases.

Given Nefsky's criticism, the expected utility approach fails to satisfy the combination of Weightiness and Generalizability.

4 Weightiness, generalizability, and the helping account

The Helping Account holds that as long as the conditions for helping are satisfied, the individual action is non-superfluous and thus could help to bring about the morally significant outcome. Part of this view is that regardless of whether it in fact ends up helping, the non-superfluous action makes progress towards the outcome (Nefsky, 2017, pp. 2756–2757). Based on this, there is a helping-based moral reason to perform the individual action whenever the conditions for helping are satisfied. However, this reason may vary in strength, and in some cases, it can be weakened to the point of being insignificant and easily outweighed. In this section, I argue that Nefsky's Helping Account faces a problem similar to that which was raised against

¹⁵ There is no problem if an account omits to identify a moral reason for individual action because there is good reason to think that there is none given the circumstances in the particular situation. The issue is rather with accounts that fail to identify a moral reason in some types of relevant collective impact cases even though it is plausible to think that there is a moral reason in those cases.

the expected utility approach, and that the Helping Account also fails to satisfy the combination of Weightiness and Generalizability.

Nefsky identifies three variables that she thinks will affect the strength of helping-based reasons for individual action: (1) the significance/severity of the outcome in question, (2) the size of the causal contribution that the individual action could make with respect to the outcome, and (3) the extent to which the outcome is “up in the air”, that is, the closer the chances of its occurrence are to 50–50 as opposed to being nearly settled to occur or to not occur (Nefsky, 2017, p. 2763). I will focus on variables (2) and (3), and show how they can affect the strength of helping-based reasons in some of the collective impact cases of concern.

Nefsky explains the third variable in terms of the extent to which it is up in the air whether enough acts of X-ing will be performed in order to achieve the outcome. But factors other than the (expected) number of acts of X-ing can also affect the extent to which the outcome is up in the air. The extent to which the outcome is up in the air in turn affects the likelihood that the individual action will actually help. Consider two scenarios in Drops of Water where an individual is considering whether to donate her pint: (1) it is unlikely that any of the others will donate their pints, or (2) there is already a highly reliable rescue mission underway to save the people in the desert. In 1), it is nearly guaranteed that the suffering will not be relieved, due to an insufficient number of water donations. In 2), it is nearly guaranteed that the suffering will be relieved, due to something other than water donations. In both scenarios, the individual action is non-superfluous and could help, but the probability that it will actually help is low, which means that the moral reason for the individual to donate her pint is weaker than it would have been if the chances of the outcome occurring were closer to 50–50 (Nefsky, 2017, pp. 2762–2763).¹⁶

We see from this that one aspect that affects the extent to which the outcome is up in the air, and thus the likelihood of an individual action actually helping, is the fact that there may be several different ways of achieving a given outcome besides the collective action that consists of individual acts of X-ing (the second condition for helping can be thought of as a requirement that there be at least one route to achieving the outcome, Y, that involves acts of X-ing (Nefsky, 2017, p. 2754, fn. 23)). For example, if the desired outcome is the shutdown of factory farms where animals (say chickens) are suffering, then one way to achieve this outcome is to pass legisla-

¹⁶ Nefsky thinks that it is plausible that variable (3) affects the strength of the helping-based moral reason, but that this requires further discussion. I think we should accept that variable (3) affects the strength of the helping-based moral reason as suggested above; it is plausible that the moral reason for performing an action is weaker when it is unlikely that the action will actually help to achieve the outcome. If so, it would at least be the case that the moral reason for individual action is weaker when the outcome is nearly settled to *not* occur than when the outcome is more “up in the air.” Whether the same is true in the case where the outcome is nearly guaranteed to occur is less clear. If the outcome is nearly guaranteed because it is highly likely that a sufficient number of individual actions have or will be performed, then the individual action is non-superfluous *and* likely to in fact help as long as it is performed while the outcome is still not guaranteed. If anything, it seems that the moral reason for individual action would be stronger in this scenario. On the other hand, if the possibility is small that the outcome will (fail to) occur at least in part due to an (in-)sufficient number of individual actions because the outcome is nearly guaranteed by some factor other than a sufficient number of individual actions, we might think that the helping-based moral reason would be weaker. In that case, the closer the outcome is to being guaranteed by another factor, the weaker the moral reason for individual action.

tion that prohibits large-scale production of chicken. Another way is for a significant portion of consumers to refrain from individual purchases, making production unprofitable.

Relatedly, something that affects the size of the causal contribution that an individual action could make (the value of variable (2)) is that there may be many different factors besides acts of X-ing that might be needed for, or that may play a part in the realization of the outcome (hence why the conditions for helping hold that it must be possible that the outcome will (fail to) occur *at least in part* due to an (in-)sufficient number of acts of X-ing). Take the chicken example again. On the route to the desired outcome that involves acts of refraining from buying chicken, a sufficient number of individual acts will be part of what brings about the desired outcome, but other factors will also play a role. For instance, the shutdown of production also requires that meat producers fail to capitalize on new markets and that they do not adopt cheaper ways of producing to compensate for the decrease in sales (see Budolfson's (2019) discussion about buffers on the supply side). Hence, if a given outcome is not brought about, it will be partly due to an insufficient number of acts of X-ing, partly due to other factors, and partly due to failure to bring about the outcome in some other way (Nefsky, 2017, p. 2754 fn. 23).

It seems, then, that the greater the likelihood that something other than collective action by individuals will determine the outcome, and the greater the significance of other factors, the smaller the expected causal contribution of an individual action – and the weaker the helping-based reason for performing that action. This is especially salient in large and very complex real-world cases where significant structural changes are likely to have a much greater impact on the relevant outcomes compared to collective action by individuals. Take climate change and animal suffering in meat industries; climate change can be mitigated through structural means such as global regulations, international cooperation, and technological innovation. Animal suffering can be remedied, for example, through investment in alternative food sources and formal sanctions against inhumane farming practices. These complex cases thereby differ from the stylized examples like Drops of Water where individual contributions are central. In the latter type of case, individual contributions collectively constitute the primary causal determinant of the outcome (all else equal).¹⁷ Hence, even if the Helping Account can identify a weighty enough moral reason in those cases, the above reasoning suggests that the account will not apply equally well in the complex cases.

It can be objected that, in the real-world cases that we have discussed, it is not meaningful to distinguish between routes to achieving the outcome that involve the

¹⁷ As an anonymous referee pointed out, some structural changes are importantly intertwined with individual action. For instance, some policies aim specifically to incentivize the relevant individual actions, so the effect of such policy is not separable from that of individual action. Implementing incentivizing policy increases the likelihood that there will be enough individual actions to affect the outcome, so the presence of such policies increases rather than decreases the expected causal contribution of an individual action. However, there are structural measures that are not aimed at increasing the likelihood of the relevant type of individual action, and it is safe to assume that more powerful structural measures, such as large-scale industries transitioning to clean energy, will be more effective. Hence, the general point still stands that the expected causal contribution of an individual action will likely decrease as the expected effect of other factors increases, even if this is not the case for all other factors.

relevant individual actions and those that do not. Presumably, the outcomes in question are most likely to be achieved through a combination of factors where collective action by individuals is one. Hence, in these cases it might be best to think of the third variable primarily in terms of the extent to which it is up in the air whether enough acts of X-ing will be performed in order to achieve the outcome – or rather in terms of whether enough acts of X-ing will be performed in order for the collective action consisting of these acts to be any significant *part* of what brings about the outcome, given that other factors play a part as well – and take this to be the primary factor affecting the likelihood that the individual action will actually help.

There may be some grounds for believing that it is unlikely that a sufficient number of acts of X-ing will be performed in the complex cases like that of climate change. An extremely large number of individual actions would likely be required to achieve a substantial mitigation of climate change that is comparable to what might be achieved through structural change. Then again, it may be that *over time* the number of individual actions will be sufficient to achieve significant mitigation, but compared to structural changes, this is likely a matter of a much longer time frame. It may be more accurate to say, then, that it is unlikely that there will be enough individual actions *early enough* in these cases.

Given the demonstrated impact of variables (2) and (3), we see that the weight of helping-based reasons can be significantly reduced in the overall judgment. Take the example of an individual meat purchase to summarize: if the collective action of refraining from buying meat is one of many ways to reduce animal suffering in meat industries, and there are other, structural factors that may be decisive, and it is unlikely that a sufficient number of acts of refraining will be performed in time, then the moral reason for individual action is based on a small chance that the action will make a small causal contribution to a potential decrease in animal suffering.

Just like the corresponding expected utility-based reason, such a reason may be considered weak and easily outweighed by other considerations. Therefore, like the expected utility approach, the Helping Account fails to satisfy the combination of Weightiness and Generalizability. It should be noted, though, that the Helping Account does better in some respects compared to the expected utility approach. In particular, if it is true that the expected utility approach will not identify a moral reason for individual action in non-threshold cases, then the Helping Account achieves greater generalizability since it will identify a moral reason in threshold cases and non-threshold cases alike.¹⁸ Still, given the problems raised in this section, this is not enough for the Helping Account to satisfy the combination of Weightiness and Generalizability.

5 The strength of reasons

A potential objection to the arguments in Sect. 4 is that, although helping-based reasons for individual action can be weakened as suggested, it need not be the case that they are weakened to the point of being insignificant and easily outweighed. It can be

¹⁸ I thank an anonymous referee for pushing this point.

argued that helping-based reasons are strong as a default, and although in some cases these reasons will be weaker than in others, they will not be so weak as to be easily outweighed. If strong helping-based reasons are the default, then the weakening of these reasons need not prevent the Helping Account from satisfying Weightiness and Generalizability.

I do not think that an argument to this effect will set helping-based reasons apart from expected utility-based reasons in terms of weightiness, nor that it can block the inference from the problems raised in Sect. 4 to the result that the Helping Account fails to satisfy Weightiness and Generalizability. If we reject the claim that helping-based reasons can be weakened to the point of being easily outweighed, it seems that we must reject this claim for expected utility-based reasons as well. It seems to me inconsistent to hold that a small expected harm yields no moral reason or only a very weak moral reason while a small expected contribution in the form of helping yields a strong moral reason as a default. I see no reason why there would be such a difference in reason-giving force between expected difference-making and expected helping.

Nefsky might stress that the helping-based reason is not grounded primarily in the expected causal contribution (expected helping), but in the fact that an individual action is non-superfluous and makes progress towards the outcome. Sometimes the helping-based reason will be stronger (weaker) because the expected causal contribution is greater (smaller), but the default reason-giving force of the action's non-superfluity and progress-making is unaffected, or so it can be argued. But why could we not say something similar on behalf of the expected utility approach? We could argue that the fact that an individual action has a chance of making a difference yields a strong moral reason to act regardless of how small that chance is. On this view, just the fact that an individual is doing something that has a chance of making a difference to the morally significant outcome gives her a reason to act, never mind if that chance is miniscule and the expected harm very small. I do not see why this would be any less plausible than attributing a strong reason-giving force to an individual action's non-superfluity and treating the expected causal contribution as secondary.

It can be objected that the above argument ignores a relevant difference between the expected utility approach and the Helping Account in terms of what is reason-giving. What is reason-giving on the Helping Account is not just the fact that an individual action could help, but that it makes progress towards the outcome; one additional drop for each thirsty person is one drop closer to the relief of suffering (Nefsky, 2017, pp. 2756–2757). While the fact that there is a chance of an individual action making a difference is merely a matter of what the individual action could do, making progress is something that the individual action actually does. But once again we can think of ways of construing the expected utility approach in a similar way. For instance, we could say that, given that each individual action has a chance of making a difference, each action increases the likelihood that a difference will be made, which is also a form of progress-making.

At this point, it can be responded that both the Helping Account and the expected utility approach identify moral reasons that are strong as a default; neither helping-based reasons nor expected utility-based reasons will be weakened to the point of being easily outweighed. Both approaches can appeal to the fact that there is a poten-

tial instrumental relationship (and some kind of actual, progress-making relationship) between an individual action and the morally significant outcome, and hold that this is enough to yield a strong moral reason for individual action. This reason may be weaker (stronger) depending on variables (2) and (3), but these factors will not undermine the default strength of these reasons.

I find this response unconvincing for both approaches. If these accounts are really concerned with what the individual action can do in terms of affecting the outcome, then it seems that the reason-giving force lies to a significant extent in the individual action's expected instrumental contribution – be it in terms of difference-making or helping – and not in the mere fact that the individual action could make such a contribution. The latter does not seem to yield strong moral reasons for action on its own. If one thinks that it does, this assumes that there is something reason-giving about this fact itself, beyond the reason-giving force of the action's expected instrumentality. If so, more needs to be said about what this something is. Until then, I think that the size of the potential instrumental contribution and the likelihood of actually being instrumental are central to the presence and strength of the moral reason for individual action. If so, both helping-based reasons and expected utility-based reasons will sometimes be weakened in the ways and to the extent argued in the previous section.

However, there is one ameliorative aspect that I have not explicitly touched upon so far, namely the distinction between objective and subjective moral reasons. Nefsky (2017, p. 2755) argues that if the conditions for helping are not satisfied, the individual has no *objective* helping-based reason to act, but as long as the individual does not know whether the conditions are satisfied, she still has a *subjective* reason to act since, as far as she knows, her action could help – even if it in fact cannot. This reasoning can presumably be extended to the strength of these moral reasons as well. That is, given that an individual does not know whether her reason for acting is weakened by facts about the situation at hand, her reason may still be subjectively strong. Moreover, this applies to the expected utility approach as well; as long as the individual believes that her action has a non-negligible chance of making a difference, she may have a (strong) subjective moral reason to act.

In his critique of the expected utility approach (particularly Kagan's (2011) meat consumption account), Budolfson (2019, pp. 1720–1721) argues roughly that people typically have access to evidence about real-world situations (e.g. about the workings of the meat industry) that suggests that the difference an individual action could make and the probability that it will make that difference are not such that the expected value of one's individual action would yield a (strong) moral reason to act in the relevant way (e.g. to refrain from a meat purchase). If so, the individual would have neither an objective nor a subjective (strong) moral reason to act.

Hedden (2020), on the other hand, disagrees with Budolfson, both about the empirical claims regarding the effects of individual purchases on meat production and about how informed individuals tend to be in collective impact situations. He thinks that, in real-world cases, people are often ignorant about where any potential thresholds are, how many other people will act in the relevant way, and how close one might be to hitting a threshold. He also points out that, as one becomes more informed, the expected effects of one's action will most likely become smaller. But in order for the expected utility approach not to yield a moral reason for the relevant

individual action, individuals would need information at a level of detail that in reality they do not have. Hence, Hedden suggests, the expected utility approach will yield the verdicts we are looking for in the real-world cases that we are concerned with, contrary to what Budolfson argues (Hedden, 2020, pp. 536–539).

Who is right in these matters is largely an empirical question that I will not attempt to sort out here. But, even if it is true that individuals are in fact often ignorant of the details of a given collective impact situation, it does not directly follow that they have a strong subjective moral reason to act in those situations. It is reasonable to think that having (strong) subjective moral reasons would require that the individual has good reason to believe that the circumstances are such that her action would be instrumental in the relevant sense (in terms of difference-making or helping). Otherwise, the individual could have such reasons just by choosing to remain as ignorant as possible or by forming baseless beliefs about what her action could accomplish. This would be an unacceptable result.

Insofar as ignorance entails having no good reason to believe that one's action can be instrumental in the relevant sense, it may very well preclude (strong) subjective moral reasons rather than foster them. The individual needs evidence in favor of her action's instrumentality in order to ground subjective reasons. It is hardly enough that she is ignorant of evidence against it. But in the process of acquiring evidence about the situation in which she is acting, if Hedden is right, she might actually end up undermining instead of supporting any subjective moral reasons for acting. Sometimes the individual will have good reason to believe in the instrumentality of her action given the evidence, in which case she will have (strong) subjective moral reason to act. But, given the discussion in this and the previous section, it seems clear that this will not always be the case. Sometimes the individual might even have sufficient evidence about the situation to instead infer that she does not have a weighty enough instrumentality-based moral reason to act. Hence, a simple appeal to subjective reasons will not solve the problem with satisfying Weightiness and Generalizability.

6 The connectedness condition

In the previous sections, we looked at Nefsky's criticism of the expected utility approach, which is the dominant approach in the instrumental category. We then saw how the first two success conditions for a general solution to the problem of collective impact can be formulated based on that criticism. In this section, we will consider Nefsky's critique against accounts in the non-instrumental category, and formulate the third success condition. Nefsky's main critique against some of the central accounts in the non-instrumental category is that they face the following problem:¹⁹

¹⁹ Specifically, Nefsky (2015, 2019) considers Parfit's (1986) membership account and Kutz's (2000) complicity account (both of which Nefsky (2015) refers to as participation accounts), as well as Cullity's (2000) fairness account.

The disconnect problem The identified moral reason for individual action does not connect appropriately to the morally significant outcome in question (Nefsky, 2015, 2021).

More specifically, Nefsky claims that:

“[T]o address what is at issue in the Problem of Collective [Impact] we need to be able to explain how it is 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” (Nefsky, 2015, p. 268).

The problem is that, in identifying a moral reason for individual action that is independent of the instrumentality of that action with respect to the outcome, the non-instrumental accounts fail to identify a moral reason for the relevant individual action that is not also a reason to do other instrumentally superfluous things. For instance, consider participation in Drops of Water. If one can count as a participant even though one's action is instrumentally superfluous with respect to the outcome, for example in virtue of the action's expressive significance, then it is not clear why the individual has a moral reason specifically to donate her pint rather than to perform some other instrumentally superfluous but expressive action, such as holding a sign with a supportive message (Nefsky, 2015, p. 263). In other words, if participation is what is reason-giving, and an individual can count as a participant by doing X^* (e.g. holding a sign) just as well as by doing X (e.g. donating a pint of water), then the account fails to identify a reason specifically to do X (Wieland & van Oeveren, 2020, p. 167), where X is such that if enough people do it, it will bring about the outcome, whereas X^* is not. Hence, such an account fails to identify a moral reason that tells us specifically to do the type of acts that can collectively bring about the outcome of concern, thereby failing to address what is at issue in the problem of collective impact according to Nefsky.

It seems that, if an account faces the Disconnect Problem, and has no way to overcome it, the account will fail to provide a general solution to the problem of collective impact on Nefsky's view. Therefore, I suggest that we can formulate the third success condition for a solution to the problem of collective impact based on the Disconnect Problem:

Connectedness The successful account should identify a moral reason that connects appropriately to the morally significant outcome of concern. This means identifying a moral reason specifically to do the type of acts that could collectively cause (prevent) the outcome.²⁰

An immediate concern is that it is not clear what the *type of acts that could collectively cause (prevent) the morally significant outcome* is. In order for an account

²⁰ As we have noted, in many collective impact cases it is a matter of the collective action by individuals being *part of* what brings about or prevents the outcome, given that other factors also play a role.

to satisfy Connectedness, it must identify a moral reason specifically for this type of action, but it is unclear how it can do this (and how we can determine whether it does) if it is unknown what the relevant type is. In the next section, I argue that the challenge of delimiting the relevant action type reveals an underlying, more general problem that poses a challenge to Nefsky's Helping Account and prevents it from satisfying Connectedness.

7 Connectedness and the helping account

On the Helping Account, one has a moral reason to perform an act of a certain type, X, as long as it is possible that the relevant outcome will (fail to) occur at least in part due to an (in-)sufficient number of that particular type of action. I interpret from this that the relevant action type is one to which actions belong in virtue of some feature that they have in common other than belonging to the set that could together bring about the outcome. In turn, this type should be such that a sufficient number of actions of that type can collectively bring about the outcome. The problem is that it is not clear how to describe the type X, nor how to determine whether actions of type X are such that they could collectively bring about the outcome. I argue that this reveals the following more general problem that underlies the Disconnect Problem and the Connectedness condition:

The action individuation problem In order to solve the problem of collective impact, we want to identify moral reasons for actions of a certain type, but the lack of a principled way of delimiting that type stands in the way of satisfactorily identifying those reasons.

To see the difficulty in delimiting the relevant action type, consider Drops of Water again. What is the action type X in this case and which actions belong to it? It is stipulated that each individual is considering whether to donate their pint of water to the suffering people in the desert by pouring it into the cart. We could simply say, then, that the action type X is that of pouring a pint of water into the cart, and the only actions that belong to that type are those that do just that. According to the case description, actions of this type could collectively bring about the relief of suffering. So far, so good. But this type-description is too narrow. If an individual poured 9/10s of their pint into the cart, would they not be performing an action of the same type as those donating the full pint? If not, it seems each fraction of a pint poured into the cart would constitute a distinct action type, which would correspond to an extremely large number of action types – one for every additional unit of the smallest amount of water that could be separated from the amount in the individual's water container. It seems more plausible to, as Nefsky does, describe the type more broadly as acts of "pouring in water" or "making water donations" (2017, p. 2753). On this description, an act of X-ing in Drops of Water would be any act of pouring water into the cart. It turns out, however, that this type-description is too broad on Nefsky's own view.

I suggest that employing this kind of broad type-description amounts to saying that the relevant action type is that of adding to the so-called underlying dimen-

sion. The underlying dimension is something that individual actions add to that is such that when enough is added to it, it causes changes in the morally significant dimension, i.e. changes to the outcome of concern (Nefsky, 2021; Wieland & van Oeveren 2020). For example, in Drops of Water, the underlying dimension is the amount of water in the cart, whereas the morally significant dimension is the relief of suffering. In the case of climate change, the underlying dimension is the amount of greenhouse gas in the atmosphere, and the morally significant dimension is climate change-related harms (Nefsky, 2021). This is the type-description used on Wieland and van Oeveren's (2020) participation account, which holds that one is part of the group that might bring about the morally significant outcome when one adds to the underlying dimension (Wieland & van Oeveren, 2020, pp. 179–180). That is, an individual action counts as participatory as long as it adds to the underlying dimension – call this the UD Account.

However, Nefsky (2021) argues that the UD Account faces the Disconnect Problem. First, since all it takes to participate is to add to the underlying dimension, an individual can participate even by contributing a very small amount – for example by pouring a single drop of water into the cart in Drops of Water. Second, the UD Account is degreeless (Wieland & van Oeveren, 2020, p. 181), and holds that one participates all the same regardless of how much one contributes to the underlying dimension. Hence, the participation-based moral reason that the UD Account identifies is just as much a reason to donate a single drop as it is a reason to donate a full pint. But, Nefsky (2021) argues, donating a single drop is not the type of act that if enough of them are performed, it will bring about the relevant outcome – a significant relief of suffering, not all additions to the underlying dimension fulfill this criterion. So, if the participation-based moral reason for donating a pint is also a reason to donate just a drop, the account fails to identify a moral reason *specifically* to do the type of acts that could collectively bring about the outcome of concern.

It is worth noting that it is not true that adding just a drop of water to the cart is not the kind of thing that if enough people do it, it will bring about significant relief of suffering for the people in the desert. It seems that there would just have to be a very large number of drops added (this would be the case if, for instance, each potential water donor donated their pint one drop at a time), and we can say something similar about other small contributions, such as refraining from very small emissions of greenhouse gas (e.g. by shortening one's Sunday drive by a minute or so). Nevertheless, thinking realistically about collective impact cases, the claim is relevant. If the outcome of concern is significant mitigation of climate change, it seems that even if everyone on earth did something like shortening a drive by one minute, it would not bring about this outcome. In Drops of Water, there is a limited number of people in a position to act, and if we add the assumption that each person may only pour water into the cart once, it is clear that even if they all added a drop it would not bring about any significant relief of suffering.

Granted that, in at least some cases, very small contributions to the underlying dimension do not constitute the type of act that could collectively bring about the outcome, there is a sense in which the type description “acts that add to the underlying dimension” is too broad. As noted above, Nefsky's Helping Account seems to be (at least implicitly) operating with the same kind of broad type-description. For

instance, in the case of climate change, Nefsky (2017, p. 2744) mentions individual actions such as refraining from going for a drive or taking a flight as examples of acts of X-ing. Without further specification, the suggested action type seems to be that of reducing or avoiding personal emissions, broadly understood. Hence, the Helping Account appears to be facing the same problem: if the action type X is something like “donating water” or “refraining from emissions,” then this includes very small contributions such as donating a single drop or refraining from an additional minute of driving. This means that the helping-based reason for making a larger contribution is also a reason for making a very small contribution.

One way to try to avoid this is to accept the view that each amount added to the underlying dimension constitutes an action type in itself, which would separate larger contributions from smaller ones. This could allow the Helping Account to identify moral reasons specifically for the contributions that are large enough to collectively bring about the outcome. But there is still the challenge of determining precisely which contributions are large enough. Realistically, it seems that there is no principled way of doing this unless we could point to some precise threshold in the relationship between the underlying dimension and the morally significant dimension, e.g. between the amount of water in the cart and the relief of suffering from thirst. But, as we have seen, relying on the notion of thresholds comes with its own problems. Besides, the Helping Account is supposed to apply in the same way regardless of whether there are thresholds or not (Nefsky, 2017, p. 2757). Hence, it seems that the Helping Account cannot escape the Action Individuation Problem so easily, and that it fails to satisfy Connectedness.

It can be objected that, unlike the participation-based reason on the UD Account, the helping-based reason to donate a pint of water would not be just as much a reason to donate just a drop. Given the discussion in Sect. 3 about the size of an action’s causal contribution and the strength of reasons, it seems that the Helping Account may hold that the helping-based reason is stronger the greater the contribution to the underlying dimension.²¹ On this view, the helping-based reason would be much stronger for donating the full pint than for donating just a drop. I think that adopting this view goes a long way towards overcoming the Disconnect Problem, at least in cases where the individual actions in question add to an underlying dimension. After all, it seems implausible to think that an account would fail to address what is at issue in the problem of collective impact just because the moral reason it identifies supports relevant contributions of any magnitude.

However, while arguing that the moral reason varies in strength depending on the size of the contribution ameliorates the problems with employing a broad action type-description, it is not the same as identifying a moral reason *specifically* for the type of acts that could collectively bring about the outcome. As we have seen, the Helping Account offers no principled way of doing this. As long as identifying a moral reason specifically for this type of action is taken to be part of what it takes to avoid the

²¹ At least given that the size of the contribution to the underlying dimension correlates appropriately with the size of the action’s potential causal contribution to the outcome. But a proponent of the Helping Account could also argue that the reason is stronger because a greater contribution makes more progress towards the outcome.

Disconnect Problem, specifying a difference in the strength of the moral reason will not be enough to overcome this problem and satisfy Connectedness. To do this, the account must solve the underlying Action Individuation Problem by finding a way to delimit the type of acts that could collectively bring about the outcome.

The Action Individuation Problem also seems to pose a specific challenge to the Helping Account given that action types play a crucial role on the account. As we have seen, the individual action's ability to help depends on whether it is possible that the outcome will (fail to) occur at least in part due to an (in-)sufficient number of that particular type of action (Nefsky, 2017, p. 2753). If it is unclear how to describe the relevant action type, it might be difficult to evaluate whether the individual action could help, which in turn can create ambiguity in the application of the Helping Account. For instance, this might be particularly salient when it comes to evaluating individual actions that do not add to the underlying dimension but that still seem to be part of the collective impact situation since they contribute in other, perhaps more indirect ways. Examples include various kinds of political action such as demonstrations and online activism. Categorizing such actions into types and evaluating them against the conditions for helping might not be straightforward.

8 Concluding remarks

Julia Nefsky's Helping Account claims to provide a general solution to the problem of collective impact by identifying a helping-based moral reason for individual action. First, I have argued that, based on the problems that Nefsky has raised against previously suggested solutions, three success conditions for a general solution to the problem of collective impact can be formulated: The Weightiness condition, the Generalizability condition, and the Connectedness condition. Second, I have argued that the Helping Account fails to satisfy the three success conditions, thereby failing, by Nefsky's own standards, to provide a general solution to the problem.

However, the fact that the Helping Account fails to satisfy these conditions does not mean that it does not make useful progress towards solving the problem. Even though it does not identify a weighty enough moral reason in the full range of relevant collective impact cases, it may well do so in some subset of the relevant cases, i.e. in the less complex cases where collective action by individuals is crucial in determining the outcome. This can at least be seen as a partial solution. Moreover, when it comes to the challenges concerning action individuation, it may often be enough to acknowledge that the moral reason is stronger the greater the contribution.

As I said at the beginning, I do not argue that the three conditions formulated here are in fact conditions for a successful solution to the problem of collective impact, but the question of what is required for a successful general solution is of interest for future research. However, the demonstrated difficulty in satisfying the three conditions formulated here, coupled with the recognized merits of previously suggested solutions like the Helping Account, might instead lead us to call into question the need for and feasibility of a general solution to the problem. Perhaps a pluralistic approach, one that employs different accounts to identify moral reasons for individual action in different types of collective impact cases, is within closer reach. If a

pluralistic approach turns out to be better applicable in practice, it might also be the more desirable option.

Acknowledgements I would like to give special thanks to Krister Bykvist and H. Orri Stefánsson for continuous discussion of and many rounds of valuable comments on this work. I am also grateful to Romy Eskens, Jens Johansson, Olle Risberg, Jules Salomone-Sehr, Jan-Willem Wieland, and two anonymous reviewers for *Philosophical Studies* for their helpful comments. I thank organizers and participants at the Practical Philosophy PhD Seminar at Stockholm University, and the Higher Seminar in Practical Philosophy at Uppsala University, for valuable comments and discussion. Finally, I thank Såde Hormio, Julia Nefsky, and all other participants at the Small Acts, Big Harms Workshop on individual responsibility for collectively caused outcomes at the University of Helsinki, for productive and highly engaging discussions.

Funding The author did not receive support from any organization for the submitted work. Open access funding provided by Stockholm University.

Declarations

Competing interests The author has no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Barnett, Z. (2018). No free lunch: The significance of tiny contributions. *Analysis*, 78(1), 3–13. <https://doi.org/10.1093/analys/anx112>.
- Briggs, R. A. (2019). Normative Theories of Rational Choice: Expected Utility. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2019). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2019/entries/rationality-normative-utility/>
- Broome, J. (2019). Against Denialism. *The Monist*, 102(1), 110–129. <https://doi.org/10.1093/monist/ony024>.
- Budolfson, M. (2019). The inefficacy objection to consequentialism and the problem with the expected consequences response. *Philosophical Studies*, 176(7), 1711–1724. <https://doi.org/10.1007/s11098-018-1087-6>.
- Carlson, E., Edling, M. J., & Johansson, J. (2021). The significance of tiny contributions: Barnett and Beyond. *Utilitas*, 33(4), 488–496. <https://doi.org/10.1017/S0953820821000248>.
- Cullity, G. (2000). Pooled Beneficence. In M. J. Almeida (Ed.), *Imperceptible Harms and Benefits* (pp. 1–23). Springer Netherlands. https://doi.org/10.1007/978-94-011-4144-4_1
- Hedden, B. (2020). Consequentialism and collective action. *Ethics*, 130(4), 530–554. <https://doi.org/10.1086/708535>.
- Hiller, A. (2011). Climate change and individual responsibility. *The Monist*, 94(3), 349–368. <https://doi.org/10.5840/monist201194318>.
- Kagan, S. (2011). Do I make a difference? *Philosophy & Public Affairs*, 39(2), 105–141. <https://doi.org/10.1111/j.1088-4963.2011.01203.x>.
- Kutz, C. (2000). *Complicity*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511663758>.

- McPherson, T. (2021). The puzzle of Inefficacy. In M. Budolfson, T. McPherson, & D. Plunkett (Eds.), *Philosophy and Climate Change* (pp. 222–248). Oxford University Press.
- Nefsky, J. (2011). Consequentialism and the problem of collective harm: A reply to Kagan. *Philosophy & Public Affairs*, 39(4), 364–395. <https://doi.org/10.1111/j.1088-4963.2012.01209.x>.
- Nefsky, J. (2012). *The Morality of Collective Harm* [Doctoral Thesis, UC Berkeley]. <https://escholarship.org/uc/item/9s49q22t>
- Nefsky, J. (2015). Fairness, participation, and the real problem of collective harm. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics* (5 vol., pp. 245–271). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198744665.003.0012>.
- Nefsky, J. (2017). How you can help, without making a difference. *Philosophical Studies*, 174(11), 2743–2767. <https://doi.org/10.1007/s11098-016-0808-y>.
- Nefsky, J. (2019). Collective harm and the inefficacy problem. *Philosophy Compass*, 14(4), e1–e17. <https://doi.org/10.1111/phc3.12587>.
- Nefsky, J. (2021, June 2). *Participation, Progress, and Superfluity*. Small Acts. Big Harms, Helsinki, Finland.
- Parfit, D. (1986). Five mistakes in Moral Mathematics. *Reasons and persons*. Oxford University Press. <https://doi.org/10.1093/019824908X.001.0001>.
- Wieland, J. W., & van Oeveren, R. (2020). Participation and superfluity. *Journal of Moral Philosophy*, 17(2), 163–187. <https://doi.org/10.1163/17455243-01702002>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.