

Olle Torpman¹

Responsibility-Based Reasons to Act in Collective Impact Cases²

What moral reasons to act could an individual have if her action would not make any difference? In this paper, I argue that there are responsibility-based reasons for individuals to act, and that these can help explain why an individual sometimes should act in so-called collective impact cases even if she cannot make a difference with respect to the outcome in those cases. I distinguish between retrospective and prospective kinds of responsibility, and argue that (i) an individual has prospective responsibility-based reasons to act in a specific way in collective impact cases, if she will thereby avoid contributing causally to harm, or contribute causally to good when that is desirable; and (ii) an individual has retrospective responsibility-based reason to act in a specific way in collective impact cases, if she would otherwise be blameworthy for making a (causal or constitutive) contribution to harmful outcomes in such cases.

¹ Institute for Futures Studies, correspondence: olle.torpman@iffs.se.

² Funding from Riksbankens Jubileumsfond (grant number: P22-0662) is gratefully acknowledged.

1. Introduction

It is common to think that an individual has no reason to act specifically in a situation if she cannot make any relevant difference to the outcome in that situation. For example, if it does not matter to climate change whether or not you stop emitting, then you have no reason – at least no climate change-related reason – to stop emitting. This poses a problem particularly in so-called *collective impact cases* where the impact (good or bad) stems from individuals' collective actions but where no individual member of the collective seems to be capable of making a relevant difference to that impact. Intuitively, however, it seems that individuals should at least sometimes act specifically even in such cases. But what moral reasons to act specifically could an individual have in collective impact cases if her action will not make any difference to the occurrence of the outcome?

There are basically two possible strategies of arguing for the existence of reasons for an individual to act specifically in such cases. One strategy would be to object to the idea that an individual's action in collective impact cases does not make any relevant difference, and show *how* it actually can make such a difference. In this regard, some have argued for the existence of expected utility-based reasons, according to which an individual has reason to act in virtue of the chance (however small) that her action will pass a threshold that leads to a (much) better outcome (see, e.g., Kagan 2011). Voting cases seem to be such cases, where there is initially a small chance for each vote that it will make a difference as to who will win.

The other strategy would be to argue that there are other reasons besides those that are *difference-based*, as we may call them, and that such other-based reasons are present in collective impact cases. In this regard, some have argued for the existence of *fairness-based* reasons, arguing that the collective has a duty to act in collective impact cases, and that the only fair thing to do for an individual member of this collective is to participate in the work that is needed (e.g., Cullity 2000; Baatz 2014). Others have argued in favor of *virtue-based* reasons for actions, the idea being that we have reason to perform actions that stems from virtuous motives or character traits, whether or not these actions make any difference to the outcome (see, e.g., Jamieson 2007; Hourdequin 2010). Another proposal refers to *helping-based* reasons, the idea being that an individual act can help to bring about an outcome in the sense that it makes a non-superfluous causal contribution to that outcome, even if it cannot make a difference to the outcome, and that this can in itself have reason-giving force (see Nefsky 2017). Yet others have argued in favor of *participation-based* reasons, where an individual has reason to participate in group activities that can make a difference, even if the individual herself cannot make this difference (e.g., Wieland & Oeveren 2020).

My aim in this paper is to offer another proposal along the lines of the second strategy, which, I will argue, is less problematic than the existing proposals of that kind. More precisely, I will argue for the existence of *responsibility-based* reasons as a distinguished type of reasons for individuals to act specifically in collective impact cases. In section 2, I briefly clarify the different notions of responsibility that are relevant to the present paper, distinguishing between prospective moral responsibility and retrospective moral responsibility. In section 3, I discuss prospective moral responsibility-based reasons for action. In section 4, I discuss retrospective moral responsibility-based reasons. Section 5–8 answers potential objections to the responsibility-based reasons account, most of which have been raised against other accounts in the debate. Section 9 concludes.

2. Different types of responsibilities

There are many different types of responsibility discussed in the philosophical literature (see, e.g., Williams 2010; Poel, Royakkers, & Zwart 2015). For instance, an agent can be *causally* responsible for something in the sense that she caused it, or *attributively* responsible in the sense that it is attributable to her agency, or *morally* responsible in the sense that she is either praise- or blameworthy for it or under a duty to do something about it. This paper is concerned with moral responsibility. More precisely, I will investigate the reasons for action that moral responsibility can ground in collective impact cases.

As the above description unveils, there are (at least) two types of moral responsibility. On the one hand, an agent is *retrospectively* morally responsible if and only if she is worthy of praise or blame for her choices of actions, or the outcomes of her actions. Roughly speaking, retrospective moral responsibility is backwards-looking, and regards *things one has done (or omitted doing)*. On the other hand, an agent is *prospectively* morally responsible if and only if she has a certain duty to act with respect to a certain situation – e.g., to care for someone, to solve a problem, or to pay certain costs. Prospective moral responsibility is thus forward-looking, and means responsibility *to do something* or, in other words, *to see to it that something is the case* (Poel, Royakkers, & Zwart 2015).

Any account of responsibility identifies what we may call *responsible-making features* – that is, the (set of) features in virtue of which an agent is responsible. Different responsible-making features might be relevant depending on whether we have retrospective moral responsibility or prospective moral responsibility in mind. One such feature that is relevant to retrospective responsibility concerns foreseeability, in the sense that an agent can be worthy of blame or praise for an action or outcome only if she understands the situation and can foresee the connection between her

action and the outcome. Another such feature concerns voluntariness, implying that an agent can be praised or blamed for an action only to the extent it is within her own control. A third retrospective responsible-making feature concern avoidability, in the sense that an agent can be blamed for a choice of action only if she could have chosen otherwise. One of the prospective responsible-making features concerns ability, in the sense that an agent can be responsible to do something only if she has the ability to do so. I will get back to these differences below.

While neither causal nor attributive responsibility in themselves implies reasons for action, moral responsibility might do. Given the two types of moral responsibility clarified above, there are potentially two different types of moral responsibility-based reasons: (i) retrospective responsibility-based reasons, and (ii) prospective responsibility-based reasons. Below, I will discuss both types. Given the tight connection between prospective responsibility and duties, it is clearer that prospective responsibility might yield reasons for action. Hence, I will start with that.

However, the connection between prospective responsibility and duties might put into question the relevance of the notion of prospective responsibility, and hence the relevance of prospective responsibility-based reasons for action. Why not just say that I have a duty or obligation to care for my daughter, and a duty or obligation not to do harm, etcetera, and skip the talk about prospective responsibility? Saying that an agent A has a prospective responsibility to φ might thus be a different way of saying that A has a duty or obligation to φ . If so, talk about prospective responsibility would be redundant and uninformative. Hence, it would make no sense either to talk about prospective responsibility-based reasons for individuals to act specifically in collective impact cases.

The way to address this worry, I think, is to point out that there is a sense of prospective responsibility that differs from duty and obligation. There are several ways in which this can be done. For instance, in the entry “Collective Responsibility” on the *Stanford Encyclopedia of Philosophy*, Marion Smiley says that “[i]n cases where we use the language of moral obligation, we signal that the agent has to perform a particular act. In cases where we use the language of responsibility, we allow the agent to use its own judgment in deciding how to bring about the desired state of affairs” (Smiley 2023). If this is correct, prospective responsibilities concern generic types of actions, whereas duties mainly concern sub-types and particular tokens of actions. This means that I might have a responsibility to do something in general, without having a duty to do anything in particular.

A similar view is found in Robert Goodin, who moreover argues that what matters to prospective responsibility is that the agent “*see[s] to it that X*” (see Goodin 1995: 83). He says that “[s]eeing to it that X” requires, minimally; that [the agent] satisfy himself that there is some process (mechanism or activity) at work whereby X will be

brought about; that [he] check from time to time to make sure that that process is still at work, and is performing as expected...” (ibid.).

This idea is shared by Ibo van de Poel, Lambèr Royakkers, and Sjoerd D. Zwart, who moreover argue that the “...sense in which responsibilities are different from duties [...] is that responsibilities do not require the agent to achieve the outcome ϕ by her own actions” (Poel, Royakkers, & Zwart 2015: 28–29). The idea is, in other words, that prospective responsibilities can, whereas duties cannot, be fulfilled by external factors. They can, for instance, be delegated to other agents or realized by natural causes. Again, what matters is that the agent *sees to it that* ϕ . This moreover implies, they argue, that “[a]lthough this responsibility is aimed at realizing ϕ , the occurrence of ϕ is not the main criterion whether [an agent] actually fulfilled her forward-looking responsibility” (Poel, Royakkers, & Zwart 2015: 29). Interestingly, this moreover means that an agent can fulfil her prospective responsibility to see to it that ϕ even in cases where ϕ does not occur.

It is also possible that prospective responsibilities can *ground* duties, meaning that they would be more fundamental than duties. Overall, on the basis of these observations, I will assume that there is a notion of prospective moral responsibility that is not redundant, and, hence, that it makes sense to investigate the possibility of prospectively moral responsibility-based reasons for actions.

3. Prospective responsibility-based reasons: Do no harm

Prospective moral responsibility connects agents with possible future actions. As mentioned above, it provides reasons for seeing to it that a certain state of affairs obtains. As I see it, prospective moral responsibility can yield reasons to act in basically two ways: (i) to abstain from wrongdoing in the first place (i.e., unconditionally); and (ii) to correct for wrongdoings that have already taken place (i.e., conditionally). Since most collective impact cases discussed in the literature do not involve prior wrongdoing, I will here focus on unconditional prospective responsibility. Note that I here use ‘wrongdoing’ in a broad sense to be compatible with different moral theories.

I assume there are two types of unconditional wrongdoings in this regard: to contribute causally to the presence of harms where avoidable, or to contribute causally to the absence of benefits where desirable. I here use “harms” and “benefits” also in a broad sense to be compatible with different moral theories. A necessary condition for wrongdoing is, thus, to contribute causally to harms where avoidable, or to not contribute causally to benefits where desirable. This implies that a sufficient condition for abstaining from doing wrong is to contribute causally to neither the presence of avoidable harms, nor the absence of desirable benefits. In fact, I can see no other way

in which an agent can abstain from wrongdoing. If I am correct about that, then this is also necessary for abstaining from wrongdoing.

Consequently, individuals have prospective moral responsibility-based reasons to not contribute causally to harm where avoidable, and to contribute causally to benefits where desirable, in collective impact cases. This means that the relevance of prospective moral responsibility-based reasons for individuals to act specifically in collective impact cases thus hinges on the meaning of “causal contribution”. There are different ways in which “causal contribution” could be analyzed. For the sake of argument, I will here assume the so-called NESS (Necessary Element of a Sufficient Set) account, which takes a cause to be a necessary element of a set that is sufficient in the circumstances for their effects (see, e.g., Beebe & Kaiserman 2020). This view builds on the views of J. L. Mackie (1965) and has in different versions been proposed by, for instance, Brahm and van Hees (2009) and Kaiserman (2016). The definition provided by the NESS account can be formulated as follows:

An agent A (in circumstances C) contributes causally to an outcome O if, and only if, A performs an action such that (i) the action is a member of a set of actions that is sufficient (in C) for O, and (ii) no subset of that set of actions is sufficient (in C) for O.

This merits clarification. First, it implies that the set of actions is minimally sufficient for O, meaning that there is no proper subset relative to the set at issue that would also realize O. This does not require that there is no proper subset relative to the full set of involved actions. For example, if you and four other agents act in a way that leads to O, but only four agents’ actions are necessary for the realization of O, then there are several subsets of actions – e.g., the original set of actions minus your action – that would also have realized the outcome. This, however, does not mean that your action does not contribute causally to the outcome. Given that you actually perform your action in this case, your action is itself a member of a set (indeed four sets) which is minimally sufficient for the realization of the outcome. This means that you do make a causal contribution in this case. This moreover shows that the definition applies to cases of overdetermination.

Second, it is important to note that the occurrence of “performs an action” is a simplification. In fact, what is relevant is what the agent chooses to do – whether it is to act or to omit. If intentional, an omission could also make a causal contribution, since an agent’s choice to omit can constitute a member of a set that is sufficient for the realization of an outcome. Suppose that it is enough that three out of four people intentionally omit to push a button in order to realize O. If all four intentionally omit to push the button, then each of their individual intentional omissions belongs to a

set that is minimally sufficient for the realization of O. This moreover implies that the only way in which an agent can make sure not to contribute to the outcome in such a case is to push the button. When doing so, the agent's choice no longer belongs to any set of actions that is minimally sufficient for that outcome. Still, the definition has the welcome implication that it does not count all omissions as causal contributions to outcomes. Indeed, omissions sometimes do not constitute members of any set of actions that are minimally sufficient for realizing the outcome at issue. If you choose to omit to take part in a joint activity that saves 10 lives, and if your omission has no effects on that activity, then your choice does not belong to any set of acts that is minimally sufficient for realizing the outcome. Hence, your omission does not count as a causal contribution to saving those lives.

Equipped with this notion of causal contribution, we can draw some conclusions regarding the prospective moral responsibility-based reasons for individuals to act specifically in collective impact cases. In general, it gives an individual reasons to (i) abstain from taking part in any collective activity that produces harm where avoidable, and to (ii) take part in collective activities which produce benefits where desirable, since by doing so she sees to it that her action (i) is no member of any set that is minimally sufficient for the presence of that harm, and (ii) is a member of a set of actions that produce that benefit. In the case of climate change in particular, the NESS account of causal contribution implies that an individual can fulfill her prospective unconditional responsibility by not emitting. Only thus can she see to it that her action is no member of any set that is minimally sufficient for the production of harmful climate change.

Of course, objections may be raised. But since many objections apply equally well to the account of retrospective responsibility-based reasons, I will first have a look at that account.

4. Retrospective responsibility-based reasons: Avoid blame

There is a widespread view in the literature on moral responsibility that an agent can be retrospectively morally responsible for an outcome if and only if they voluntarily, foreseeably, and avoidably contribute somehow to that outcome (see, e.g., Williams 2010; Braham & van Hees 2012; Poel, Royakkers, & Zwart 2015; Goodin 2018). This means that contribution, voluntariness, foreseeability, and avoidability are conditions for retrospective moral responsibility. Consequently, if an agent A fulfills these conditions with respect to a certain outcome O, then A is retrospectively morally responsible for O. In addition, if O is (sufficiently) morally bad or undesirable, then A is

blameworthy for O. If O is instead (sufficiently) morally good or desirable, then A is praiseworthy for O.

Under the plausible assumption that agents have reason, other things being equal, to avoid blame, we may assume that they have reason, other things being equal, to end any relevant responsibility-relation between themselves and the outcomes for which they are blameworthy. Given the conditions for retrospective moral responsibility, this means that the agent has a reason to avoid contributing voluntarily and knowingly to such outcomes.

Although it is possible in principle to fulfill this requirement through involuntary or ignorant action, it is hard to see how this would be possible in practice. If an agent puts himself in a situation where he is forced to do something (in order to fail with respect to voluntariness), then he would most certainly be blameworthy for having put himself in that situation. Likewise, if an agent puts himself in a state of ignorance (in order to fail with respect to foreseeability), then he would most certainly be blameworthy for having put himself in that state. This indicates that the only practical way of avoiding blame for a certain outcome is to not contribute to that outcome.

This line of reasoning suggests that there are retrospective moral responsibility-based reasons – or blame-avoidance reasons – for action. Moreover, it suggests that individual agents thus have such reasons not to contribute to morally bad or undesirable outcomes – be them individually or collectively produced.

Note that I have so far left it open what type of ‘contribution’ is at stake in the case of retrospective moral responsibility. The reason is that some have argued that causal contribution is not necessary for retrospective responsibility. Braham and van Hees, for instance, argue that “...holding a person morally responsible in the sense of blameworthiness appears to require something weaker than actual causal contribution to some state of affairs. A person may be blameworthy if, inter alia, the action they performed is at least a potential causal factor” (Braham & Hees 2009: 342). If this is correct, there is an interesting difference between the conditions for prospective and retrospective moral responsibility, respectively.

Robert Goodin (2018) offers another suggestion along these lines. He says that even if you do not causally contribute to a certain outcome stemming from a collective activity, you could still be constitutively responsible for it in virtue of taking part in, and hence being part of, the collective activity as a whole which produces this outcome. Being constitutively responsible in this respect, he argues, simply means being a part of a whole. This means that human beings and their actions can be constitutively responsible for group activities of which they are part.

Moreover, Goodin argues that “[y]ou bear constitutive responsibility in the sense that you are liable to credit or blame for voluntarily and knowingly being a part of that whole” (2018: 41). Presumably, if the outcomes of such activities are (sufficiently)

bad, and if the agent contributes voluntarily and knowingly, then that agent is blameworthy for taking – and hence being – part of that activity. If the outcome is (sufficiently) good, the agent is instead praiseworthy. Given that agents have blame-avoidance reasons for action, and given that constitutive responsibility in collective harm cases implies blameworthiness, an agent has retrospective responsibility-based reasons not to take part in such activities. This holds whether or not it could be argued that they contribute causally to that outcome.

Summing up thus far: An individual has prospective responsibility-based reasons to act in a specific way in collective impact cases given that she will thereby avoid contributing causally to harms where avoidable, or contribute causally to benefits where desirable. Also, an individual has retrospective responsibility-based reason to act specifically in collective impact cases given that she will otherwise be blameworthy for making a voluntary, foreseeable and avoidable (causal or constitutive) contribution to harmful outcomes in such cases.

Let us now consider objections.

5. First objection: The problem of emergent properties

The above proposed account(s) of responsibility-based reasons for individuals to act specifically in collective impact cases could be objected to by pointing out a distinction between aggregative and emergent properties. This argument has been raised by Kingston & Sinnott-Armstrong (2018). Although they raise it against what they call “the partial causation argument” (i.e., against the applicability of the notion of causal contribution) in collective impact cases, it might also apply to the notion of constitutive responsibility, hence making it potentially effective against both types of responsibility-based reasons for action.

Aggregative properties are properties that belong to both parts and wholes, where the property of the whole equals the aggregate sum of that same type of property of the parts. Size and weight are examples of such properties. If every single piece of the puzzle is 3 cm², then the thousand-piece puzzle as a whole is 3000 cm². Or, to use Kingston & Sinnott-Armstrongs example: “[C]onsider a quantity of oil that has a mass of one kilogram and contains, say, 3 times 1025 molecules of oil. Then we can calculate the mass of one molecule of oil simply by dividing one kilogram by 3 times 1025” (Kingston & Sinnott-Armstrong 2018: 175).

Emergent properties, on the other hand, belong only to the whole, and are hence not properties of its parts. Such properties emerge out of parts that lack that property. Kingston & Sinnott-Armstrong exemplify:

The quart of oil is very slimy, but an individual molecule of oil by itself is not slimy at all. It is not that the molecule has a little sliminess, but much less than the whole quart. An individual molecule is not slimy in the least. We cannot feel any individual molecule at all, so it cannot feel like slime. The same point applies to other properties of the oil, including appearing yellowish and causing moving parts to last longer. (2018: 175)

On their view, climate change is emergent in this way. They say that, “[j]ust as individual molecules of oil do not cause parts of sensations of sliminess (or yellowish color), so individual molecules of greenhouse gas do not cause parts of dangerous climate impacts. Instead, as with the sliminess and color of oil, what increases the dangerous impacts of climate change is larger groups of molecules of greenhouse gases” (2018: 175). What is more important, is that they take the emergent property of climate change to imply that individual emitting actions, lacking that property, cannot be partial causes of climate change. Against their opponents in the debate about “joyguzzling” (joyriding in a gas guzzler) as an example of questionable emitting activity, they say that if “...global climate change as well as specific climate events and their harms are all emergent phenomena [...] they cannot cite partial causation to justify their claim that there is a moral requirement to refrain from joyguzzling” (2018: 176).

However, the mere distinction between aggregative and emergent properties does not rule out that emitting actions can be parts (i.e., members) of wholes (i.e., sets) that are themselves minimally sufficient for climate change – yet climate change is an emergent property which is lacking in individual emitting actions. If an individual emitting act is such a member, then it is a causal contributor, whether or not climate change is an emergent property. Consider voting for example. No single vote for candidate A possesses a ‘winner-making’ feature. But if more than 50% of the electorate vote for A, then these votes will together possess that feature. This means that ‘winner-making’ is an emergent property. Nevertheless, individual votes may contribute causally to A’s election win. Consequently, just because climate change would be emergent rather than aggregative, this does not imply that emitting activities could not contribute causally to climate change. Kingston & Sinnott-Armstrong are therefore wrong when they say that “the partial causation argument [...] assumes that climate change is aggregative, not emergent” (2018: 178).

Still, their objection might have force against the applicability of the notion of constitutive contribution in the context of retrospective responsibility-based reasons. It seems plausible, for instance, to say that one molecule of oil is not a constitutive part of the sliminess of a gallon of oil. Likewise, it seems plausible to say that a single act of emissions is not a constitutive part of the climate change harm. The underlying explanation would be that no individual act (such as an emitting action) which lack

an emergent property (such as climate change harm) can constitute a part of an emergent property (such as climate change harm).

At a closer look, however, this seems to be false. Suppose that I knowingly and voluntarily add a certain chemical, C1, into a bowl, another person knowingly and voluntarily adds another chemical, C2, into that same bowl, a third person knowingly and voluntarily adds yet another chemical, C3, into that bowl, and these three chemicals together give rise to a chemical composition with a corrosive emergent property that is lacking in each of the single chemicals, C1-C3, and in each of the three pairs of them. Even if it cannot be said that our individual actions (of adding a single chemical into a bowl) are constitutive parts of the corrosiveness as an emergent property of that chemical composition, we might say that our individual actions are constitutive of the chemical composition as such. This shows that a non-emergent action could be a constitutive part of a whole that gives rise to an emergent property. I might thus be constitutively responsible for that. If someone is harmed by the corrosive chemical composition, for example, I would be blameworthy for contributing constitutively to its cause. The same seems to hold in the climate case: Even if my individual emissions would not be constitutive of any climate change harm as such, I may be blameworthy for acting in a way that is constitutive of the ‘cloud’ of emissions, as it were, that causes climate change harm.

6. Second objection: Non-threshold cases

There is a certain type of collective impact cases that appears to pose a problem for a responsibility-based account of reasons for action. These are called “non-threshold cases” (Nefsky 2017) or “non-triggering cases” (Tiefensee 2022), which are distinguished from so-called “threshold cases” or “triggering cases”. Nefsky explains the difference as follows:

In threshold cases, for each outcome of the morally significant sort in question, there is some precise number of acts of the relevant type needed to bring it about: any less will not be enough to bring it about, and any more will not change things with respect to that outcome. If a threshold is hit exactly, though—as in the case of a tie or a one-vote-win—then each act can make a difference. In non-threshold cases, on the other hand, there is no precise number of acts of the relevant kind needed for a given outcome. While acts of a certain type together cause (or are part of what cause) the outcome in question, there is no sharp boundary between enough such acts and not enough. So, in non-threshold cases you cannot have enough acts for a particular outcome without having more than enough such acts, and thus taking one away will never itself make a difference. (2017: 2746)

Tiefensee exemplifies this difference by saying that “whereas election wins are clear examples of triggering phenomena, in that victory is generally secured upon reaching the precise threshold of 50% of the votes plus one, no such precise threshold appears to exist in relation to air pollution [or] water contamination...” (2022: 3308).

The possibility of non-triggering cases assumes vagueness in the form of semantic or metaphysical indeterminacy. In the case of climate change this means, Tiefensee points out, that “while some amount of greenhouse gases is sufficient for global warming to be harmful, which exact amount this is remains vague” (2022: 3311). More precisely, she thinks that non-triggering cases requires metaphysical indeterminacy. In the climate case, this means that there is no fact of the matter as to which precise amount of greenhouse gases would be minimally sufficient to bring the collective harm of global warming about (see also Kingston & Sinnott-Armstrong 2018).

As Tiefensee mentions, however, this possibility hinges on a number of controversial assumptions. First, it assumes that there is in fact metaphysical indeterminacy. Second, it assumes that climate change is of such kind. Being aware of these controversies, she emphasizes that she will not commit herself to these assumptions, but rather investigate what would follow if they were true (2022: 3309).

She discusses two different interpretations of metaphysical indeterminacy in this respect. On the first, “what is metaphysically indeterminate is *when* the increasing amounts of CO₂ molecules become sufficient to cause harmful global warming” (2022: 3311-2, my emphasis). There is thus no precise threshold after which, but rather a range or interval within which, the relevant climatic effect may be caused. On the second interpretation, “metaphysical indeterminacy could be understood along the lines of *ontic* indeterminacy. More precisely, we could argue that the cloud itself is an ontically indeterminate object, such that there is simply no fact of the matter as to which molecules are part of it” (2022: 3321, my emphasis).

The problem with non-triggering cases, involving metaphysical indeterminacy of either of these types, is that the standard notions of causal and constitutive contribution, respectively, appear inapplicable. In such cases, an individual agent’s choice of action seems not to constitute any member of any set of actions that is minimally sufficient for the realization of the collective impact, and might not even be a constitutive part of the whole that causes it. If climate change harm is non-triggering, it would be hard to explain how individuals are contributing (causally or constitutively) to it (see, e.g., Wieland & van Oeveren 2020: 175-6). Since causal and constitutive contribution is a condition for prospective and retrospective moral responsibility, respectively, it seems that we thus have to accept that individuals lack any responsibility-based reasons to act in a specific way in non-triggering collective impact cases.

I think this conclusion is too hasty, however. If metaphysical indeterminacy is real in the sense that there is no fact of the matter as to which precise amount of emissions

is minimally sufficient to bring about the collective harm of global warming, or which emissions end up as constitutive parts of the ‘cloud’ as a whole which causes this warming, then this plausibly means that it is indeterminate as well whether or not a specific agent’s action will be a member of any such set, or a constitutive part of such a whole. However, this does not imply that our emitting actions are *never* members of sets of actions that are minimally sufficient for the realization of harmful climate change. Nor does it imply that our individual emitting actions *never* end up as constitutive parts of the ‘cloud’ that in effect causes such harm. Rather, it implies that our emitting actions *sometimes* are members of such sets, as well as constitutive parts of such wholes. As Tiefensee puts it:

[A]t the moment of releasing CO₂ molecules, we do not know where these molecules will end up: Will they remain totally detached from the cloud, such that they have nothing to do with the cause of harmful global warming? Will they find themselves in the cloud’s centre, such that they determinately belong to the cause of this collective harm? Or will they end up in the shaded areas, such that there is no fact of the matter as to whether or not the molecules we release are part of the harmful cloud, and thus part of the cause of the collective harm? (2022: 3322)

This suggests that metaphysical indeterminacy – if real – implies epistemic uncertainty: If there is indeterminacy in the world, then we cannot know if or when (or which of) our actions belong to which of these categories. Even if climate change (or any other collective impact case) is non-triggering, an individual’s action *might* end up in a set of actions that is minimally sufficient for the realization of the undesired outcome, or become a constitutive part of the whole which causes this outcome. Although we will not be able to determine exactly *which* emitters make such contributions, we are able to establish that non-emitters certainly *do not* make such contributions.

Other things being equal, it is plausible to assume that an individual is prospectively responsible to not risk making contributions to harm. And this gives her reasons not to take such a risk. Since it can moreover be argued that taking such risks is blameworthy, she would in addition have retrospective responsibility-based reasons pointing in the same direction. Hence, individuals would have both prospective and retrospective moral responsibility-based reasons to not take part in collective harm cases – whether or not they involve thresholds or metaphysical indeterminacy. In the climate case, the only way in which the agent can make sure her emissions do not end up in the ‘cloud’ of emissions that causes harmful climate change, is to not emit.

7. Third objection: The problem of overriding reasons

The account of responsibility-based reasons might appear to yield too strong reasons for individuals to act specifically in collective impact cases. To see this, suppose that a construction worker is about to fall down from the top of a wobbly scaffolding, unless all of the five and only bystanders step in to stabilize it. You are one of these five bystanders. As you happen to know, however, none of the other bystanders will step in. Hence, you know that your decision to step in can make no difference with respect to the construction worker eventually falling. In this case, it might seem implausible to say that you should step in.

Nevertheless, since the only way in which you can abstain from contributing causally (or constitutively to the cause of) the fall of the construction worker is to step in, the account I have proposed implies that you do have a responsibility-based reason to step in. Indeed, that is the only way in which you can see to it that your action will not constitute a member of a set of actions that is minimally sufficient for the construction worker's falling, or not become a constitutive part of the cause of that fall. Do we hence have a *reductio* argument against the account of responsibility-based reasons?

No. What explains the intuition that you should not to step in in the wobbly scaffolding case, is not that you cannot make any difference by stepping in, but rather that it makes a difference in some other respect not to step in. In most real-world cases like this, stepping in would cost time and involve risks to oneself – which could be avoided by choosing not to step in. The mere fact that an action cannot make a difference in some respect can never in itself be a reason not to perform it unless there is some alternative action the performance of which can make a difference in some (perhaps other) respect.

We hence need to distinguish between *pro tanto* reasons and all-things-considered reasons, of which only the former may be overridden by other more weighty reasons (see, e.g., Wieland & van Oeveren 2020). In the wobbly scaffolding case, the responsibility-based reason you have to step in is a *pro tanto* reason that is overridden by the reasons you have to not step in. In a situation where the other four bystanders would have stepped in, however, the high moral value of saving the construction worker from falling would imply that the reason for you to step in overrides the *pro tanto* reason (regarding costs of time and risks to yourself) to not step in. The lesson to learn from this is that just because one should not do X does not mean that one has no reason to do X.

But what if your only reasons to step in are responsibility-based reasons, and where

all other reasons – e.g., difference-based, and self-interested, etc. – point against stepping in? If they are weighty enough, I think we should just bite the bullet and accept that you have all-things-considered reasons not to step in. But what if it holds for all collective impact cases – say, that all other reasons together carry heavier weight and jointly recommend something different than the responsibility-based reasons? Then, of course, we would have to accept that our responsibility-based reasons for action would yield no concrete normative implications in such cases.

However, I do not think that is the case. First of all, the main reasons against stepping in (or in other ways acting in a specific way) in collective impact cases seem to be self-interested reasons, since doing so often requires a personal sacrifice. But it is not set in stone that such reasons always carry heavier weight than responsibility-based reasons in such cases. Moreover, it is not clear that an agent will always have self-interested reasons not to step in in collective impact cases. Sometimes she will benefit more from stepping in than from not. For instance, there are well known co-benefits from eating vegetarian instead of meat, as well as from taking the bike instead of the car to work, and so on.

Second, what an agent is morally required to do in cases of collective impact (as in any other type of case) is what she has all-things-considered reasons to do. And what she has all-things-considered reasons to do is determined by the weighing together of all pro tanto reasons she has in that situation. In the introduction, I briefly mentioned some such reasons for participating in collective impact cases – such as expected utility-based reasons, fairness-based reasons, virtue-based reasons, and helping-based reasons. Even if none of these pro tanto reasons would in isolation be capable of yielding any moral requirement of individuals to act specifically in collective impact cases, they might together be able to yield such a requirement.

8. Fourth objection: The problem of non-generalizability

In her criticism of other accounts of reasons for individuals to act in collective impact cases, Nefsky appears to implicitly assume what Andrea Asker (2023: 2384) explicates in a number of “success conditions” for such accounts. First and foremost, Asker explicates a “Generalizability condition”, according to which “[t]he successful account should identify a weighty enough moral reason in all the collective impact cases of concern”. As this means, an account of an individual’s reasons to act specifically in collective harm cases should have something interesting to say in such cases. This condition seems to be implicitly assumed also by others in the debate (see, e.g., Kingston & Sinnott-Armstrong 2018).

The discussion in the previous section suggests that the account of responsibility-based reasons fails to meet the generalizability condition. For instance, if responsibility-based reasons are in some cases insufficient to generate moral requirements, perhaps due to the existence of overriding reasons, then it will not be able to “identify a weighty enough moral reason” in such cases. And if some collective impact cases involve metaphysical indeterminacy, and if the notions of causal or contributive responsibility does not apply in all of these cases, it means that the account of responsibility-based reasons might not apply to those cases either.

As Asker points out, however, it is not obvious that generalizability should be accepted as a condition for accounts of individuals’ reasons to act specifically in collective impact cases. As she says, the best approach might well be “...a pluralistic approach, one that employs different accounts to identify moral reasons for individual action in different types of collective impact cases...” (Asker 2023: 2395). My previous arguments point in the same direction. If what an individual should do in a certain situation is a matter of what all-things-considered reasons for action she has in that situation, and that responsibility-based reasons constitute one type of pro tanto reasons that together with other pro tanto reasons determine her all-things-considered reasons, then it is simply implausible to assume that only one type of reason should identify a weighty enough moral reason in all the collective impact cases of concern. While responsibility-based reasons might be most salient in some collective harm cases, virtue-based, fairness-based, expected utility-based, helping-based – or any other relevant – reasons might be more salient in other such cases.

We should therefore accept that there might be cases – e.g., some non-threshold cases where an individual actually does not contribute (neither causally, nor constitutively) – where there are no responsibility-based reasons to step in. We should also accept that, if the individual also lacks any other-based reasons to step in or if she has stronger reasons not to step in – then she actually should not do so. This also suggests that it would be a mistake to assume from the start that all collective impact cases are such that they involve weighty enough reasons for individuals to step in.

8. Conclusion

In this paper, I have argued that an individual has two types of responsibility-based reasons to act in a specific way in collective impact cases: (i) she has prospective responsibility-based reasons to act, if she will thereby not contribute causally to the presence of harm where avoidable or to the absence of good where desirable; and (ii) she has retrospective responsibility-based reason to act, if she will otherwise be blameworthy for making a (causal or constitutive) contribution to harmful outcomes in such cases.

The responsibility-based account has the advantage not only of avoiding some of the problems to which other accounts are vulnerable, but also to answer some of the remaining issues observed by others in the debate. For instance, Tiefensee argues that, due to the possibility of non-triggering cases and metaphysical indeterminacy,

[w]e must be able to show that individual agents have a reason to act in a specific way in view of a morally relevant aggregate effect E, even though their actions make *no difference* to E and they are *uncertain* whether or not there is a *fact of the matter* as to whether or not their actions are *partial causes* of E. (2022: 3322)

The account of responsibility-based reasons does just that. Moreover, Wieland and van Oeveren (2020: 185), defending the account of participation-based reasons, say that one remaining question related to their account is this: “why is participation morally significant [...]?” The account of responsibility-based reasons answers this question as follows: Participation is morally significant because it lets the agent fulfill prospective and/or retrospective responsibilities.³

References

- Asker, Andrea S (2023) ‘The Problem of Collective Impact: Why Helping Doesn’t Do the Trick’, *Philosophical Studies* **180**: 2377–97. doi:10.1007/s11098-023-01995-7.
- Baatz, Christian (2014) ‘Climate Change and Individual Duties to Reduce GHG Emissions’, *Ethics, Policy & Environment* **17**: 1–19. doi:10.1080/21550085.2014.885406.
- Beebee, Helen & Alex Kaiserman (2020) ‘Causal Contribution in War’, *Journal of Applied Philosophy* **37**: 364–77. doi:10.1111/japp.12341.
- Braham, Matthew & Martin Hees (2009) ‘Degrees of Causation’, *Erkenntnis* **71**: 323–44. doi:10.1007/s10670-009-9184-8.
- Braham, Matthew & Martin van Hees (2012) ‘An Anatomy of Moral Responsibility’, *Mind* **121**: 601–34.
- Cullity, Garrett (2000) ‘Pooled Beneficence’, in Michael J Almeida, ed., *Imperceptible Harms and Benefits*: 1–23. Springer Netherlands.
- Goodin, Robert E (1995) *Utilitarianism as a Public Philosophy*, Cambridge Studies in Philosophy and Public Policy. Cambridge University Press.

³ I would like to thank Andrea Asker, Krister Bykvist, Daniel Ramöller, and the audience at the *Second Internal Workshop* on May 6 2024 at IFFS, for helpful comments.

- Goodin, Robert E (2018) 'Constitutive Responsibility: Taking Part, Being Part', *Analysis* **78**: 40–45. doi:10.1093/analys/anx146.
- Hourdequin, Marion (2010) 'Climate, Collective Action and Individual Ethical Obligations', *Environmental Values* **19**: 443–64. doi:10.3197/096327110x531552.
- Jamieson, Dale (2007) 'When Utilitarians Should Be Virtue Theorists', *Utilitas* **19**: 160. doi:10.1017/s0953820807002452.
- Kagan, Shelly (2011) 'Do I Make a Difference?', *Philosophy & Public Affairs* **39**: 105–41. doi:10.1111/j.1088-4963.2011.01203.x.
- Kingston, Ewan & Walter Sinnott-Armstrong (2018) 'What'S Wrong with Joyguzzling?', *Ethical Theory and Moral Practice* **21**: 169–86. doi:10.1007/s10677-017-9859-1.
- Mackie, J L (1965) 'Causes and Conditions', *American Philosophical Quarterly* **2**: 245–64.
- Nefsky, Julia (2017) 'How You Can Help, without Making a Difference', *Philosophical Studies* **174**: 2743–67. doi:10.1007/s11098-016-0808-y.
- Poel, Ibo van de, Lambèr Royakkers, & Sjoerd D Zwart (2015) *Moral Responsibility and the Problem of Many Hands*. Routledge.
- Smiley, Marion (2023) 'Collective Responsibility', in Edward N Zalta and Uri Nodelman, eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2023/entries/collective-responsibility/>.
- Tiefensee, Christine (2022) 'Indeterminacy and Collective Harms', *Philosophical Studies* **179**: 3307–24. doi:10.1007/s11098-022-01827-0.
- Wieland, Jan Willem & Rutger van Oeveren (2020) 'Participation and Superfluity', *Journal of Moral Philosophy* **17**: 163–87. doi:10.1163/17455243-01702002.
- Williams, Garrath (2010) 'Moral Responsibility', *Oxford Bibliographies Online*.