

# What is Moral Actualism?

Marta Bax

Does philosophy figure into our discussions about political theory and policy making? For many the answer is two-fold: when philosophy is practical—i.e., when it is a direct enquiry into government practices and the relationship between states and their citizens, the discipline is a crucial means to legitimising policies, and grounding actions in a set of coherent and justifiable beliefs. On the other hand, thought experiments, logic problems, language riddles—philosophy of a more theoretical nature—is not typically consulted. In politics, we want to narrow down pragmatic answers that can intuitively appeal to the masses, not overcomplicate ideologies to such a linguistically technical point that they become unintelligible to most. This essay is an attempt to 'do' political philosophy in the second, less appreciated way. At the end of my arguments, I aim to explain how these kind of enquiries might inform our politics, particularly when it comes to political discourse which is, by nature, theoretical—such as debates regarding our impact on future generations.

## INTRODUCTION

Sofia intends to have a child. She has recently been offered a promotion in New York, which has left her with an important decision to make. She could choose to move, and pursue a fulfilling career, or she could stay in Manchester, where she knows her mother would help with the new-born. Suppose she chooses the latter. When asked by her colleagues why she passed on the opportunity, Sofia replies 'it was the best option for the people involved. I'm not too attached to this particular job, and I think it would benefit my daughter to have her grandparents around. My wife is also staying in Manchester, and although she supports me unconditionally, I know she'd prefer to avoid long-distance. I'd also like to stay close to her, if possible.' Most people would think this a reasonable response.

Kalani has also been promoted. She too would like to have children someday. If she chooses to stay in her home city, Liverpool, she and her partner will try for a baby. If she moves to New York, however, she has decided to refrain from conceiving for a few years. Suppose Kalani opts for the latter. When asked by her neighbour why she passed on the chance to conceive, she replies 'it was the best option for the people involved. This job will bring me financial stability, and living elsewhere is a great opportunity I might not get later in life. Waiting until I am back in Liverpool to conceive means my partner won't have to live far from his child, and that I get the childcare support I need.' Now suppose her neighbour retorts 'but you haven't chosen the best option for all the people involved—your decision not to conceive was bad for the baby that would have been born if you had stayed in Liverpool, who would have had a good life. Because of your decision, he will not be born!'

Both Sofia and Kalani explain their actions by appealing to the interests of the people affected by their decision. The concern Kalani's neighbour has however raises a question for justifications of this kind: who counts as a person affected? Why do the interests of Sofia's future child, in part, determine the outcome of her decision? Why should this not apply to Kalani's case?

Here is a relatively intuitive assumption: in order to be affected by an action, a person must first exist. Here is a more intuitive assumption: this person need not exist at the time the action first takes place. A sufficient condition for being affected by an action is that one exists whilst the effects of that action take place. When

Sofia chooses to stay in Manchester, her future child does not exist. However, the consequences of Sofia's choice to stay in Manchester will affect said child. Therefore, the child's eventual interests contribute to Sofia's decision. Kalani is not in the same position as Sofia. When she chooses to leave for New York, her future baby *does not and will never exist*. As a result, the effects of Kalani's actions cannot violate that child's interests. She is under no obligation to consider the interests of this child, due to the fact that they *will never be affected by her actions*.

If you share these intuitions, you might find yourself drawn to a particular moral doctrine: 'Moral Actualism', wherein the moral status of any action  $a_j$ —actual or not—is determined by whether its outcome is better or worse for the people who exist given action  $a_j$ , than the outcomes of other available actions.<sup>1</sup>

In other words, what makes an option morally better than its counterpart is that its outcome is better for the people who actually exist, given that action is performed.

In this essay I defend moral actualism from a particular attack launched by philosopher Caspar Hare. Hare's counter paints actualism in untenable colours—he argues that the position ultimately is not a reliable moral doctrine because it occasionally leads to deontic absurdity. We should reject it and look for a different way to evaluate the moral status of actions, particularly those concerning future generations.

I want to provide actualists with an alternative. Instead of abandoning what feels like an intuitive moral doctrine—particularly one which lends itself well to legitimising democratic policies (which are supposedly good 'for the people'), they need only revise their position to defend it from attack. This revision is made possible by the vagueness of the terms 'actual' and 'exist'. It is evident that in his characterisation of moral actualism, Hare assumes that to be actual is to exist in the *modal sense*: existence is satisfied by membership in the actual world. I am going to argue—contra Hare—that this alone is not a sufficient condition for existence. To be actual is also to exist in the *temporal sense*: existence is satisfied by *present* membership in the actual world.

This revision is not only prompted by a need to defend actualism against Hare's counter. In fact, before explaining how temporal actualism is invulnerable to Hare's argument, I motivate the position without reference to Hare. Ultimately, we should favour temporal

<sup>1</sup> Adapted from Hare (2007), pp. 502–503.

moral actualism not just because it is the version of actualism that overcomes Hare's counter, but because careful consideration shows that it aligns better with our general intuitions about morality.

A word of warning: this paper lies at the intersection of ethics and metaphysics. As such, many of the labels I use have appeared in different debates associated with both these areas of philosophy. 'Modal actualism' and 'Actualism' are examples (Hare 2007, p. 499; Cohen & Timmerman 2020). In the context of this essay, please treat any label with the word 'actualism' in it as meaning a version of *Moral Actualism*, the position outlined prior.

## SECTION 1: THE PROBLEM

### §1.1: HARE'S MORAL ACTUALISM

Let us begin by getting clear on some actualist essentials i.e., those beliefs that fall directly out of the definition provided in the introduction. The first concerns the concept of the good. Moral actualists think goodness is relational: an action is good if it is *good for a particular person* (or people), compared to alternatives (Hare 2007, p. 499). Often these approaches to morality are rights-based: moral duties are generated by agents' right to determine what actions are imposed on them. To act correctly is to do nothing more than respect these rights.

The second obvious credo to moral actualism is this: rights—such as the right to have one's interest count towards moral deliberation—are strictly reserved for actual people: individuals who exist in the world the relevant action takes place in (Hare 2007, p. 499). Existence is crucial supposedly because it is a prerequisite for having interests, which is in turn a pre-requisite for having rights (Hare 2007, p. 509). For an action to be good for someone, that someone must first *exist*.

What exactly is existence? Who counts as an actual person, according to actualism? The answer to these questions does not seem to fall directly out of *Moral Actualism*.

Caspar Hare thinks that existence for the actualist must extend to future people who will be born as well as current people who are already alive, given  $a_j$  is performed.<sup>2</sup> His concept of existence is modal: a person exists insofar as they are alive—at some point in time—in the real world as opposed to a merely possible world (Hare 2007, p. 498). Both present people who are alive at the time  $a_j$  takes place and future people who will live in the world created by  $a_j$  therefore have the right to have their interests contribute to that action's moral status. After all, these are the people who risk being harmed by  $a_j$ . For Hare, then, actual people and non-actual people are mutually exclusive. All actual people, like Sofia's child, exist (in the sense that they exist in the real world), and all possible people (like Kalani's child) do not.

Let's allow for now that this characterisation of existence is exactly what moral actualists are after when they call a person actual. If this is correct, they should exclusively worry about the interests of current people and future people *who will be born* when deciding between possible actions. Put a different way: only individuals who are presently alive or will be alive in the future have the right to demand their interests be considered when it comes to moral deliberation.<sup>3</sup>

### §1.2: HARE'S COUNTER

In compliance with his characterisation of moral

actualism, Hare launches into a critique of the position to conclude that it inevitably leads to deontic absurdity. His counter starts with a thought experiment:

Imagine there is a person Kate, who has to decide between two actions:

- $a_1$  – bring relentlessly miserable Jack into existence.
- $a_2$  – bring relentlessly miserable Jane into existence.

Assume both Jack and Jane would prefer not to exist than to exist miserably (Hare 2007, pp. 503–504).

Kate doesn't know what to do, since both options seem equally bad.

Here is actualism again:

(*Moral Actualism*) The moral status of any  $a_j$  is determined by whether its outcome is better or worse for *actual people*, than the outcomes of other available actions.

And here's a question: what should actualism say about Kate's dilemma?

Remember that moral actualism generates *rights-based* moral obligations: an action is better than its counterpart *if it respects the right of actual people to have their interests count towards deliberation*. The action that is better for actual people, therefore, is the action that Kate should perform.

As per the thought experiment,  $a_1$  and  $a_2$  will create separate worlds:  $a_1$  will create a world containing the actual individual Jack (call this  $S_{a_1}$ ), whereas  $a_2$  will create a world containing the actual individual Jane (or  $S_{a_2}$ ) (Hare 2007, p. 504). A world where Jack is actual hence is a world where Jane is merely possible; a world where Jane is actual is a world where Jack is merely possible. Actualism will therefore say that Jack has the right to have his interests contribute to Kate's decision only if Kate picks  $a_1$ ; vice versa for Jane. Kate, however, exists in both worlds. She retains the right to have her interests count towards the decision-making process, regardless of what action she ultimately picks.

The problem materialises the second we are told to assume that both Jack and Jane prefer non-existence over miserable existence. Imagine the actualist trying to help Kate decide what action to perform. She knows only to pick  $a_1$  if it is the best option available for actual people. The actual people given  $a_1$  are Kate and Jack. To call  $a_1$  a *better* option for Kate than  $a_2$  would be a stretch: both actions—from Kate's perspective—aren't great; both force her to create a miserable child. However, assuming Kate does not have a preference for gender, one might reasonably assume Kate is indifferent with regards to whether  $a_1$  gets performed over  $a_2$ . To recommend performing  $a_1$  therefore would not violate *her* interests. Things change when we consider the interests of Jack, however. He *does* have a specific preference between  $a_1$  and  $a_2$ , and it's for  $a_2$  to be performed, since this would mean escaping miserable existence.  $a_1$  therefore cannot be the best option for the people who exist in  $S_{a_1}$  (the actual people) since  $a_2$  is equally as good for Kate, and  $a_2$  is better for Jack. If Kate chooses  $a_1$ , then, actualism will say she has done the wrong thing; the interests of actual people were such that she *should have picked*  $a_2$ .

<sup>2</sup> I will not be focusing on past people for the purposes of the essay, although it is worth mentioning that for Hare, past people are also actual: they share the property of existing, at some point in time.

<sup>3</sup> The 'or' here is inclusive.

What happens if the actualist recommends  $a_2$ ? Once again,  $a_2$  can only be picked if it is the best option for actual people. Once again, Kate is indifferent about  $a_2$  being performed over  $a_1$ , so her rights are not violated when  $a_2$  is performed. However, *this time*, it's not the interests of Jack that contribute to decision-making because Jack ceases to be actual, given  $a_2$ . The actual people in  $S_{a_2}$  are Kate and Jane. But Jane, like Jack, *prefers non-existence*, i.e., the outcome of  $a_1$  is better for her than that of  $a_2$ .  $a_2$  is therefore not the best option for the people who exist in  $S_{a_2}$  (the actual people) since  $a_1$  is equally as good for Kate, and  $a_1$  is *better* for Jane. If Kate chooses  $a_2$ , actualism will *also* say she has done the wrong thing; the interests of actual people were such that she *should have picked*  $a_1$ .

This means that whatever Kate does, she does the wrong thing (Hare 2007, pp. 503–504). Moral actualists are committed to telling Kate not to pick either  $a_1$  or  $a_2$  when picking one of these options is required of her. In other words, their verdict—to not perform either action—is *deontically absurd*, i.e., it violates a deontic principle of obligation: *if given alternatives {x, y}, x ought not to be done, then, given alternatives {x, y}, y ought to be done* (Hare 2007, pp. 504–505). These are the grounds on which Hare abandons the position, altogether.

Another way of putting the issue is this: Hare's *Kate* case shows that moral actualism is internally incoherent. The two founding principles of actualism—(1) to have people's interests contribute to decision-making, and (2) to only consider the interests of actual people—contradict. Jack's interest not to exist only counts when he is actual. But the second Jack is made actual, his interest not to exist gets violated. In fact, for both Jack and Jane, they become actual when it is too late: Jack's interest to not-exist only counts when we bring him into existence, as does Jane's.

At this point one might wonder whether the principle of deontic absurdity is so convincing as to lead us to such a conclusion. Might there not be tragic situations in which one is faced with two options equally as bad as each other so that the correct verdict is that one ought to perform neither? And if the thought experiment demands a decision be however made, could it not be acceptable to perform an action whilst contemporarily realising that it should not have been performed? One might want to judge both actions as *equally* bad, since both are bad for the relevant people affected. Kate is therefore free to choose from either  $a_1$  or  $a_2$  for this reason. It's not that  $a_1$  (for example) should be picked because it is *best*, but because it is *no worse* than  $a_2$  (Hare 2007, p. 505).

As Hare notes, this solution, though superficially inviting, would spell ruin for the actualist. Say Kate does use this reasoning to pick  $a_1$ . When questioned about why she purposefully chose to perform an action that was worse for Jack, she justifies herself by saying that *not* performing it would have been *equally as bad* for Jane. This is a roundabout way of making  $a_1$  permissible because, although it harms Jack, it avoids a harm to Jane. The problem is that according to the actualist's own doctrine, the interests of merely possible people can never count towards moral deliberation. Jane (given  $a_1$  is chosen) does not and will never exist. She is not an actual member of  $S_{a_1}$ , and therefore does not have the right to have her interests count towards the enactment

of  $a_1$  (Hare 2007, p. 505). To justify picking  $a_1$  because of the equally tragic entailment of  $a_2$  is to allow for the interests of possible people to count towards deliberation. And that is enough to violate actualism irreparably.

## SECTION 2: THE SOLUTION

The fact that this big issue facing *Moral Actualism* is buried deep within a thought experiment might irritate political theorists. Unless we suddenly form the ability to predict the miserable lives and consequent wishes not to live of the unborn, and believe instances like Mary's will be common, we might be tempted to dismiss Hare's counter. Most people generally perfectly understand *Moral Actualism* as a doctrine, and are not confused when the actualist claims the important factor when it comes to decision-making is respecting the wishes of people who are alive and will live, given they are those who will be affected by contemporary policies. To call this simple credo 'incoherent' on technical grounds—though logically accurate—seems supercilious and practically ineffective.

Despite this, I am instead going to accept the nature of Hare's argument, and argue against it on his own terms. I do this mainly because I think the solution I am presenting has interesting repercussions for how we think about policymaking and reveals an insight into legislators may more effectively justify their policies so they are endorsed by more people. I'll come back to this thought towards the end of the essay, once my proposed counter is made clear.

The key to resisting Hare's counter is understanding that this modal characterisation of existence is a product of Hare's conviction that the actualist need not engage with.

With §2.1 I aim to provide an alternative reading of actualism, one which understands existence in more than simply modal terms. This is the proper way *Moral Actualism* should be understood. The first reason for this is that temporal actualism by itself seems an intuitive moral doctrine. To illustrate this, I consider an argument against adopting temporal actualism, but ultimately explain that the best this can do is slightly weaken the position; in fact, this weakening of actualism might actually work in its favour.

The second and more pressing reason to adopt temporal actualism over Hare's modal interpretation is to ensure that Hare's counter fails to get off the ground. §2.2 explains why it is that moral actualism—understood in temporal terms—is invulnerable to Hare's counter

### §2.1: TEMPORAL MORAL ACTUALISM

What does it mean to exist? Hare provides us with one option: to exist is to be a member of the *actual world*, at some point in time. But there's another option: to exist could also mean to be a member of the actual world, *in the present*. This kind of existence is determined by one's temporal status as well as their modal status. An actual person—under this interpretation—is a *present person who exists in the actual world*. Possible people, then, are people who do not presently exist in the actual world. Temporal actualism implies that modally actual people and possible people are *not* mutually exclusive. Some possible people *will become* actual, like Sofia's child. Others, like Kalani's, will not.

This would make it the case that people only gain

the right to have their interests contribute to moral decision-making *once they are born*.<sup>4</sup> Moral agents only have a duty to consider the interests of future people *once those people exist in the present*.

Why think the temporal interpretation is what moral actualists are after when they call a person actual?

Here's one reason: temporal actualism parallels how rational people tend to think about when duties should come into effect, in real life. I might think that my future soulmate is already out there in the world, and that *once I notice them*, I will be completely monogamous. Does this generate a duty in me now to be monogamous to that person who I will *eventually meet*, pre-encounter? Most people would think not: I do not owe that person any particular loyalty until I get into a relationship with (or at least meet) them. Similarly, just because someone *will eventually* become a moral agent, does not mean that they are to be treated like one in the present.

Furthermore, in order to understand the preferences of future people and therefore act in a way which fulfils those preferences, we must first have some source of information which tells us what they are. The obvious problem that is not made so clear in the *Kate* case is that, in real life, we *don't* know who the future children are, or what they want. When Sofia makes a decision she thinks will be better for her future child, she is not respecting the wishes of a future individual—her daughter—but her own idea of what her future daughter will want. This process is more accurately explained by saying that we justify our actions which affect future generations *not* as duties *we owe them*, but duties *we owe ourselves*, in which the rationale is that they will someday benefit a future person. But the fundamental mistake here is to believe that there is some other person who is committing us to perform in one way over another; there cannot be, because that future person does not *presently* exist.

To use a highly politicised current example—often, arguments against the depletion of global resources contributing to global warming share a morally favourable rhetoric of conservation because it satisfies the interests of future people (people who are actual in the modal sense), specifically. But there are other ways to justify these kinds of policies. There are authors who argue that the decision to conserve resources as opposed to depleting them stems more from the aesthetic interests of temporally actual people more so than it does the interests of future people (see [Bennett 1978](#), pp. 61–73). Secondly, even if it is performed with the sole benefit of future generations in mind, the decision to conserve could still be explained by appealing to the interests of current people only. To commit to something with someone's benefit in mind is not necessarily to act in accordance with their interests. Instead, one might reasonably hold themselves to a general moral standard to do what is best for future generations in the same way they might believe one should always act with kindness, where possible. There are *general moral standards* that are not strictly relational. They are not generated by the interests of any individual, but by a collective interest, or a pre-constructed moral code. There is no particular person that can claim the general right to conserve the planet's resources, or the right to act with kindness, but current people might still believe and do act as though these standards are of genuine moral concern. The concern with conservation need not be justified by

the actual interests of future people, but by the perceived interests of current people regarding future people.

Parallel this thought to the *soulmate* example. What causes me to want to be monogamous to my soulmate is not generated by an understanding of what's good *for them*: for all I know, my soulmate could want an open relationship. The duty I think I have is instead generated by an *accepted* general social standard for good relationships: monogamy.

Importantly, this does not mean that our new understanding of what it means for something to be 'good' fails to conform to what had been previously outlined as a central credo of actualism: the idea that something is good only if it is good *for a specific person, or group of individuals*. The *good* is still relational, in the sense that conservation remains good if it is valuable to a specific group of people—it is not *inherently* good—however, those people may not be the set of future temporal generations which will eventually be affected by conservation policies. The goodness of conservation derives from the value current temporal people assign it—we ought to conserve because we think it will work out best for future people, and we collectively (or those who wish to conserve, at least) think that looking out for future generations is something we should do.

What it does mean, however, is that sometimes our rationalisations of why we ought to perform certain actions over others are slightly misguided. I originally framed Sofia's future child's interests as a partial motivator towards her decision to stay in Manchester. This new version of actualism appears to deny that the future child has any right to moral consideration. Whilst this is true—and, I believe, reflective of real life—Sofia's assertion can still be actualist if what she really intends to say is that she wishes to enact an absolute moral rule she endorses, in that the moral rule to do what is best for your future children, whomever they might be. It is not that a particular future individual has a right to moral consideration which creates an obligation in Sofia to stay in Manchester, but instead Sofia's own belief that a particular course of action is best which results in her final decision.

Finally, nothing stated so far directly clashes with any of Hare's later solutions to his own problems (those which depart from Moral Actualism). Hare's own solutions to Non-Identity problems is not routed in a moral doctrine, but rather is tentatively sketched out in the second section of his essay:

Mary does wrong because she has a certain kind of impersonal responsibility—a responsibility to nobody in particular. Some argue that this is a responsibility that any person has—e.g., to avoid bringing about suffering. Others argue that this is a special kind of responsibility that only parents have—to avoid creating children whose lives will have certain features. This seems to me broadly the right way to approach the problem. ([Hare 2007](#), 513–514)

I see little difference between Hare's 'impersonal responsibilities' and the alternative I have been talking about. Hare suggests that these are truly what motivate people like Sofia when a decision is claimed to be made 'for the benefit of a future child'. This, in principle, is not incompatible with actualism, once we allow that

<sup>4</sup> This mirrors how rights-based approaches in ethics tend to distribute rights anyway. Rights are the kinds of things reserved for people *interacting* with people in the present, actual world.

there can be relational duties (those current people owe to other current people, generated specifically due to individual interests) and 'absolute' or 'general' moral duties (those current people owe to each other, generated by communal interests, such as that to maintain human flourishing).

## §2.2: THE APPLICATION OF THEORETICAL PHILOSOPHY TO POLITICAL THEORY

It is with this present sentiment that I want to highlight the place of theoretical discussions like these in political discourse and policy making. What temporal moral actualism essentially recommends is to stop justifying altruistic policies (such as those to conserve resources in order to reduce global warming) as being good for some future people that do not yet currently exist, but to instead route the rationale behind such policies in a common ground shared by present people, now—a sort of 'we're doing this for us, not (only) them'. This argument is illustrated by Samuel Scheffler in his book *Why Worry About Future Generations?* He writes:

When we ask why we should care about future generations, we are not simply asking why we should care whether people exist in the future or how those people fare. We are asking why we should care that the chronological succession of generations, which has delivered each of us here, should extend into the future under more rather than less favourable conditions. The difference between these questions is important. If we ask why we should care about future people, for example, or what our responsibilities toward them are, we may be tempted to suppose that the only thing that is at issue is the weight we should give to their interests or welfare. We may fail to consider the possibility that the importance to us of future generations lies partly in the fact that they are our *successors*, that their existence extends the chain of generations in which we ourselves are participants. (Scheffler 2018, pp. 15–16)

There is something powerful and all-encompassing in the idea that pursuing policies which protect future generations is something that we owe not necessarily to them, but 'generally'. If future generations from now were to unanimously agree that they would have preferred to not exist, this would likely not alter current sentiments that protecting humanity was, *is*, the right thing to do. Scheffler grounds this general value we attribute to carving out a future for people who do not exist as a product of knowing that this is what our ancestors did for us—but importantly, ensuring a future for future people is not something we owe to them, either.<sup>5</sup> Rather, it is a moral duty we give to ourselves. It is more intimate and existential than doing someone in the future a favour—it is to carry on the legacy which we (present temporal people) value, the legacy that others valued before us.

If politicians promoting conservation efforts (reminder that the conservation/depletion example is just one of many) switched from building policies on this more personal ideal than a classic utilitarian picture, it would be interesting to see what the public reception would be like. More generally, theoretical discussions like these *do* have practical applications, ones that can be discovered when thinking about thought experiments and other non-traditional methods of political enquiry. Finally, this counter (if it works) offers a solid

solution to Hare's counter, one which ultimately allows us to maintain *Moral Actualism* as a political doctrine without the background knowledge that it is actually incoherent. As mentioned in the introduction, the idea that goodness is relational, and policies should be made with people's interest primarily in mind, is critical for systems of government like democracy. Theoretical discussions like these help us question and re-affirm the foundation these power structures are built on.

## §2.3: THE TEMPORAL RESPONSE TO HARE'S COUNTER

So, how exactly does *Temporal Moral Actualism* solve Hare's counter? Remember that temporal actualism says that *if* the moral status of an action is determined by a relational duty—that is, a duty owed to specific people born of their specific interests—*then* that duty is owed only to *actual* people; people who presently exist in the actual world.

$a_1$  ought only to be realised if it is the best option available for the actual people in  $S_{a_1}$ .  $a_2$  similarly ought only to be realised if it is the best option available for the actual people in  $S_{a_2}$ . But here is the difference between Hare's actualism and temporal actualism: whereas Hare's actualism judges the actual people in  $S_{a_1}$  to be Kate and Jack, the actual person in  $S_{a_1}$  under temporal actualism is just Kate. Hare's actualism judges the actual people in  $S_{a_2}$  to be Kate and Jane, but under temporal actualism, the actual person in  $S_{a_2}$  is, again, just Kate. It is therefore only Kate's interests that must be considered under temporal actualism. We know that Kate is indifferent about whether  $a_1$  or  $a_2$  gets performed (remember that both options to her just mean that she will birth a miserable child). Neither  $a_1$  nor  $a_2$  is the *better* option for Kate—both are equally bad as each other because both are equally bad *for her*. Therefore, actualism will say that she is free to choose either option.

This means that *whatever Kate does* under temporal actualism, *she does a permissible thing*.  $a_1$  and  $a_2$  truly are arbitrary:  $a_1$  is better/worse for no one in  $S_{a_1}$ , and  $a_2$  is better/worse for no one in  $S_{a_2}$ .  $a_1$  and  $a_2$  therefore have equal moral status, which makes either option permissible for Kate to choose. Therefore, the actualist verdict—to pick from  $a_1$  or  $a_2$ —escapes deontic absurdity.

What *Temporal Moral Actualism* does is deny a key assumption of Hare's argument: when Kate chooses to bring Jack or Jane into existence, this is meant to be morally significant because the interests of two *future* actual people conflict in a way which generates a dilemma. But Hare never once considers that for actualism to survive, the interests of future people need not be relevant to decision-making. By ignoring the temporal interpretation of 'existence', Hare creates a strawman. He assumes freely that rights-based duties are owed to future people when they need not be.

## CONCLUSION

With these final points I conclude my defence of actualism with respect to Hare's counter. I hope to have at least shown that temporal moral actualism is an available refuge for any actualist threatened by Hare's arguments. A stronger possible conclusion is that temporal actualism is a more *intuitive* moral theory than Hare's actualism, even without thinking about his counter from absurdity. Temporal actualism matches our intuitions about the kinds of people who can claim

<sup>5</sup> Some people might actually argue that it *is* a duty we owe to past people as opposed to present people. Similarly to how we might think we owe people who have died the right to fulfilling their wills, or not desecrating their character, we owe it to those people to build a future for generations to come. I think this is interesting but incorrect, but will not delve into a whole discussion of the topic in this essay.

to have rights, *people who exist in the present*. Finally, I have suggested that actualists might need to make space for both strictly relational and general conceptions of the good to account for instances where we think the right action need not necessarily be what is right *for specific individuals*. I imagine the next steps for temporal actualism will involve pitting actualism against other moral doctrines to further investigate any compelling alternatives.

If there was a key take-away from the discussion, I would say it was the power of carefully crafted thought

experiments. With little information and an imagined scenario, Hare manages to convince us that future people, under actualism, are deserving of the same kinds of rights reserved for present people. When it comes to thought experiments generally—but especially in the literature on future generations—we need to be more careful. Ask yourself what assumptions the author is making, and whose perspective the experiment relies on. Otherwise, we risk generating conclusions that lead us to abandon a perfectly viable moral doctrine.

- Anglin, W. (1978). 'In defense of the potentiality principle', in Sikora, R. I., and Barry, B. (eds.) *Obligations to future generations*. Temple University Press, pp. 31–37. ISBN: 9780877221289.
- Bennett, J. (1978). 'On maximising happiness', in Sikora, R. I., and Barry, B. (eds.) *Obligations to future generations*. Temple University Press, pp. 61–73. ISBN: 9780877221289.
- Cohen, Y., and Timmerman, T. (2020). 'Actualism, possibilism, and the nature of consequentialism', in Portmore, D. W. (ed.) *The Oxford handbook of consequentialism*. Oxford University Press, pp. 139–161. <https://doi.org/10.1093/oxfordhb/9780190905323.013.5>
- Hare, C. (2007). 'Voices from another world: must we respect the interests of people who do not, and will never, exist?', *Ethics*, 117(3), pp. 498–523. <https://doi.org/10.1086/512172>
- Narveson, J. (1978). 'Future people and us', in Sikora, R. I., and Barry, B. (eds.) *Obligations to future generations*. Temple University Press, pp. 38–60. ISBN: 9780877221289.
- Roberts, M. A., and Wasserman, D. T. (2019). 'Dividing and conquering the non-identity problem', in Liao, S. M., and O'Neil, C. (eds.) *Current controversies in bioethics*. Routledge, pp. 81–98. ISBN: 9780367872755.
- Scheffler, S. (2018). *Why worry about future generations*. Oxford University Press. <https://doi.org/10.1093/oso/9780198798989.001.0001>
- Schwartz, T. (1978). 'Obligations to posterity', in Sikora, R. I., and Barry, B. (eds.) *Obligations to future generations*. Temple University Press, pp. 3–13. ISBN: 9780877221289.
- Scott, R., Jr. (1978). 'Environmental ethics and obligations to future generations', in Sikora, R. I., and Barry, B. (eds.) *Obligations to future generations*. Temple University Press, pp. 74–90. ISBN: 9780877221289.
- Sikora, R. I. 'Is it wrong to prevent the existence of future generations?', in Sikora, R. I., and Barry, B. (eds.) *Obligations to future generations*. Temple University Press, pp. 112–166. ISBN: 9780877221289.
- Tomberlin, J. E. (1996). 'Actualism or possibilism?', *Philosophical Studies*, 84, pp. 263–281. <https://doi.org/10.1007/BF00354490>
- Warren, M. (1978). 'Do potential people have moral rights?', in Sikora, R. I., and Barry, B. (eds.) *Obligations to future generations*. Temple University Press, pp. 14–30. ISBN: 9780877221289.

# Collective Impact: A Sceptical Approach

Conor Walsh

Enough people acting in a certain way can produce significant, collective harm. But if my individual act makes no perceptible difference to this harm, do I have a moral obligation to change my behaviour? Previous literature argues yes, individuals do have a moral obligation to change their behaviour in these cases; collective harm requires individual responsibility. But in this article I argue against this tradition by presenting a sceptical solution to problems of collective impact. This article will comprise four parts. First, we consider a definition of problems of collective impact. Second, we examine a consequentialist solution from Shelly Kagan, before, thirdly, presenting Julia Nefsky's *non-superfluous contribution* argument. Finally, I argue for a sceptical solution to these problems, claiming that, although individuals can make a non-superfluous contribution towards change, individuals do *not* have moral obligations in cases of collective impact problems. Instead, moral obligations fall to both the collective, either governments or regulatory-bodies, and those who profit from collective harms. In essence, this article argues that collective harms require collective responsibility.

Issues of collective harm, such as climate change, are more salient amongst younger generations ([The Economist 2023](#)), and individuals are also willing to make sacrifices in order to alleviate these harms. US consumers spent nearly \$1.9 billion on plant-based milks in 2018, while the number of vegans in the US grew from 1% of the population in 2014 to 6% in 2017 ([Kateman 2021](#)). A recent study also found that consumers are willing to pay a 9% premium for environmentally friendly food ([The Economist 2023](#)). But what moral reason do we have to make these kinds of sacrifices—to switch from beef and dairy to Quorn and soy? Individually, it seems as if our actions make no perceptible difference to rising global temperatures or the cruelty of factory-farmed animals. The collective harm of climate change will remain whether I drive or take the bus, and the scale in which factory farms operate means that my refraining from buying bacon is unlikely to save the life of a factory-farmed pig. Therefore, it appears that I have no reason to change my behaviour.

This disconnect between individual actions and collective harm is often referred to as the *problem of collective impact*. These problems are found not only in climate change and animal cruelty, but also in everyday consumer choices, and even elections.

Traditional literature on this subject has predominantly sought a theory for why individuals ought to change their behaviour. However, in this article I will seek a sceptical solution ([Nefsky 2018](#)) to problems of collective impact. This approach is sceptical since it argues that individuals do not act wrongly in these cases, deviating greatly from recent, notable work in this area by Shelly Kagan and Julia Nefsky. Instead of finding moral reasons for individuals to change their behaviour, I will argue that the responsibility for reducing collective harm falls *upon the collective*. Although I accept that individuals, in the words of Julia Nefsky (2019, p. 11), can make a non-superfluous contribution towards change, I do not believe this is sufficient to generate a moral obligation for individuals to change their behaviour.

This article will comprise of four parts. In the first section, I will outline the problem of collective impact and why it matters. In the second section, I will present a consequentialist solution to this problem, as argued by Shelly Kagan (2011), who argues that individuals

*might* make a difference by being part of a triggering cohort. Following this, the third section will present Julia Nefsky's non-superfluous contribution argument ([Nefsky 2019](#)), which argues that individuals may not make a perceptible difference, but can still make a non-superfluous contribution towards change. Finally, in the last section, I will present my own sceptical approach to this problem, arguing that the moral obligation for solving these harms falls only upon collections of individuals, either in the form of governments, regulatory bodies or those who profit from collective harm, not individuals in their own right. In essence, I will argue that collective harms require collective responsibility.

## I. PROBLEMS OF COLLECTIVE IMPACT

We begin by defining the problem of collective impact in more specific terms. This is where individual actions, taken collectively, produce harmful consequences, even though (a) no *single* act appears to make a difference and (b) had any individual acted differently, the collective harm produced would have remained the same. Since our individual actions produce no perceptible harm, nor make any difference to the collective harm, it appears that they cannot be wrong. Therefore, it is difficult to say that any individual *ought* to have acted differently ([Nefsky 2011](#), p. 364).

This conflict between individual actions and collective harm creates a moral dilemma—it appears unclear how we can solve these problems of collective harm if individuals have no moral obligation to act differently. The crux of the problem lies in the aggregation of individual actions, leading to uncertainty as to who is responsible for these collective harms. The aggregative nature of this problem will be fundamental to the sceptical solution presented in Section IV of this article.

Multiple examples of this problem exist; having provided a definition for problems of collective impact, I will now present three examples to demonstrate the importance of these issues for both the study of politics, and society more widely. The first example is found in *consumer behaviour*. Collectively, consumer decisions can have significant implications for global poverty, worker exploitation, animal rights, and the environment ([Nefsky 2019](#), p. 2). Many individuals buying clothes from a fast-fashion brand can result in the exploitation of many thousands of workers; yet my individual

decision to refrain from such a purchase is unlikely to have any perceptible impact. The scale at which such brands operate means that my purchase is insignificant.

A second example of a collective impact problem can be found in *elections and referendums*. The crux of the problem lies again in aggregation. Enough individuals voting in a large election for a 'bad' candidate or policy, could result in tremendous harm, and yet no individual vote is able to make a difference to this outcome (Nefsky 2019, p. 1). Even in cases where no harm is produced, what incentive would one have to vote in an election where popular support means the outcome is almost guaranteed prior to polling?

In cases where polls are close, collective impact problems are less relevant—take the United Kingdom's Brexit referendum, for example. National divisions meant the stakes were high. Every vote counted, and there appeared no issue of collective impact. But in states with non-partisan issues, and clearly defined executive parties, such as Singapore or Japan, or where outcomes appear guaranteed prior to polling, as in the case of recent Hungarian elections, the issue of collective impact is very real.

The third problem of collective impact that I will present is *climate change*. When enough people drive, fly, heat their homes, or even boil their kettles, the harmful consequences of climate change will occur. But individual choices to take the bus instead of driving, or to only turn one's heating on for an hour a day, instead of six, does not lead things to go differently. Climate change will occur regardless of whether or not I choose to make sacrifices, or act in a more environmentally friendly manner. If I flick the switch on my kettle, I will (arguably) be contributing to the climate crisis. Yet will my act make a significant difference? With such a small action, it appears unlikely.

Having presented three contexts in which problems of collective impact matter, the prevalence of these issues for both the study of politics and society more generally should now be apparent. Collective impact problems pose a real obstacle for electoral participation in non-partisan issues or systems that predict near-inevitable outcomes. Individuals appear somewhat powerless in these situations, but what of collectives—especially in the case of governments? Individuals on their own might not be able to act, but the collections in which we organise ourselves (nation-states, regions, political parties, activist campaigns, and so on) could have some influence. Asking questions about the roles and obligations of these groups is important and will be considered in Section IV of this article. As for society more generally, these problems are purely man-made, and oftentimes are issues that we wish to solve. If society is to attempt to solve global inequality, climate change, worker exploitation, political apathy, and other issues that fit the collective impact structure, we must ask serious questions about our individual responsibility in solving these harms, and the role that collectives can play.

Now that we have defined the problem of collective impact, provided three examples of the problem in action, and explained the significance of these problems for both the study of politics and society more generally, we will proceed to evaluate some solutions to these problems.

The following two sections, Section II and Section III, consider two traditional solutions, both of which aim to find moral reasons for individuals to change their behaviour. Section IV goes on to present a sceptical solution to these problems which deviates from the traditional literature. This sceptical solution will follow my thesis that collective harm requires collective responsibility.

## II. I MIGHT MAKE A DIFFERENCE

The first traditional solution we shall consider is from Shelly Kagan. In his paper 'Do I make a difference?' (Kagan 2011), Kagan seeks to prove that consequentialism is sufficient in solving problems of collective impact. Initially, the nature of these problems appears troublesome for consequentialist theories since individual actions appear to make no perceptible difference—it is difficult to argue that individuals should change their behaviour because of the consequences of their actions when these consequences are irrelevant. Kagan attempts to overcome these problems and to solve these issues within the consequentialist framework.

Kagan argues that problems of collective impact can be limited to *triggering cases*. These are cases in which most individual acts make no difference at all, but for some act—the triggering case—a substantial difference can be made. This is the triggering act which brings about the collective harm (Kagan 2011, p. 119). Without this, the rest of the acts are unable to bring about a collective outcome, even when aggregated. It is only the triggering act that can bring into effect this harm. For example, one or two individuals leaving banana peels on the ground is unlikely to be considered a tripping hazard, but if enough individuals do this the ground could pose a real danger. Since the first two or so individuals dropping their banana peels are not sufficient to cause this tripping hazard, there is a triggering individual—let us say the third or fourth individual—who brings about this hazard. This individual is the triggering case. Two things must be extracted from this example. First, triggering cases may not be a specific number in every circumstance, it is not always the hundredth person who drives instead of taking the bus that is the triggering case, but a rough range which is open to interpretation. Second, triggering cases are only 'triggering' because they are part of a wider triggering cohort. Without the other cases that come before it, triggering cases are not able to be the trigger of anything. Therefore, all cases in a triggering cohort carry some causal weight in the collective outcomes which they produce.

This idea of a triggering cohort is Kagan's main premise for explaining why individuals can make a difference in cases of collective harm. Kagan argues that individuals can still act wrongly in instances of the collective impact problem since, in our consumer society of mass production, there is still a triggering number of acts, let us call it  $T$ , such that I have a  $1$  in  $T$  chance of being part of a triggering cohort. Limiting his arguments to cases of factory farming, it is clear that Kagan's cohort can have an impact: I can have a  $1$  in  $T$  chance of triggering a change in demand sufficient to reduce supply by level  $T$  (Kagan 2011, p. 127). This is assuming negative net utility—that the suffering produced in production is greater than the pleasure received from consumption (Kagan 2011, p. 124).



We can illustrate this further by considering an example of factory-farmed chickens. Let us imagine the triggering amount,  $T$ , is 100, so that if the sale of chickens falls by 100, the farm will produce 100 fewer chickens the following month. I only have a 1 in 100 chance of being part of a cohort which triggers a sufficient change in demand to reduce supply, but when I am part of this cohort I can make a big difference—my refusal to buy a chicken corresponds *exactly* to saving one chicken's life, so long as I am in the triggering cohort (this is inside knowledge that we are unlikely to have), because the fall in demand equals the fall in production. Therefore, as long as I am part of a triggering cohort, it seems that my actions *might* make a difference after all.

On the face of it, this seems a satisfying solution to problems of collective impact. Kagan has used a consequentialist framework to show that I *might* make a difference after all. This solution is not without its flaws and there are two major objections relevant to this example.

The first challenge facing Kagan's argument is that he assumes a fall in demand will directly correspond to a fall in supply (Nefsky 2018, p. 274). In essence, the fall in demand of 100 chickens achieved by the triggering cohort will result in exactly 100 chickens *not* being killed. In reality, this assumption seems foolish; there are many strategies that businesses can utilise in order to realign supply and demand. Of course, cutting supply, thus reducing the number of chickens killed, is one strategy. But slaughterhouses could also seek new markets in which to sell their chickens, devise a new marketing strategy to attract new customers, or simply lower their prices to realign demand without changing supply. So even when demand drops by a triggering amount there is no guarantee that supply will fall by an *equal* amount, if at all.

The second issue for this consequentialist solution is that the scale considered by Kagan is unrealistic—factory farms operate in the millions, not hundreds. As Julia Nefsky observes (2019, p. 8), this renders the chance of being in a triggering cohort *negligible*. So not only is my abstention from purchasing one chicken unlikely to save a corresponding chicken, the chance that I will be in a cohort producing any impact at all is statistically irrelevant.

Kagan is not likely to take these challenges lightly. He would likely accept the first objection, and admit that little can be done to predict the practices of large-scale corporations—after all, that is why the exact number of the triggering cohort is information unbeknownst to the consumer. But in light of the second objection, the size of  $T$ , the scale of the triggering amount, does not matter to Kagan. What matters for Kagan is the ideal that individual actions *might* make a difference. Yet, this does little to solve our problem because this 'might' is so small that it is negligible—our actions remain imperceptible, and so the core issue of the collective impact problem remains to be solved.

### III. I CAN MAKE A DIFFERENCE

In the previous section, Shelly Kagan's expected utility argument was unable to prove that my actions *might* make a difference. Kagan's proposed solution to problems of collective impact appears unrealistic in

its assumption that a fall in demand will correspond directly to a fall in supply, and the scale in which triggering amounts occur proves so large that they are insufficient in contradicting the imperceptible nature of individual actions. In this section, we shall examine an alternative solution from Julia Nefsky—the idea that I *can* make a difference. Like Kagan, Nefsky also seeks to prove that individuals *ought* to act differently in instances of the collective impact problem.

Nefsky argues that even if our actions do not make a perceptible difference, this does not mean they are superfluous. Change is still *possible*. Individual actions *can* play a non-superfluous part in changing the outcome of a collective impact problem, even if they are unable to materialise change in their own right (Nefsky 2019, p. 10).

By this logic, we might have moral reasons to make individual sacrifices after all: not driving my gas-guzzling car might not make a perceptible difference to greenhouse gas emissions, but it will make a non-superfluous contribution towards reducing emissions; voting in an election with a predictable, harmful outcome might not make a perceptible difference to this outcome, but it will make a non-superfluous contribution towards reducing the chance of the predicted victory; and not buying clothes from a fast fashion brand might not make a perceptible difference to the exploitation of vulnerable workers, but it will make a non-superfluous contribution towards reducing this exploitation. Therefore, it seems my actions *can* change something, or at least contribute towards bringing about change, contrary to what the collective impact problem might lead us to believe.

My action being imperceptible is no longer a problem—what matters is that my action *can* contribute towards change which can occur if the circumstances are right. The question of what these circumstances are poses a challenge for Nefsky's argument. Having considered the merits of Nefsky's arguments, we shall now go on to consider one objection facing her non-superfluous contribution solution, in addition to a fundamental limitation of the traditional approach employed by both Nefsky and Kagan.

We begin by considering an objection to Nefsky's argument. For my actions to make a non-superfluous contribution towards change, we must have good reason to believe that others are also willing to contribute towards change, otherwise our actions will be redundant. This is because if no other individual is willing to change their behaviour, it means that our individual sacrifices cannot contribute towards anything, since there is nothing to contribute towards. If this is true, my voting and not-voting in the case of a harmful candidate is irrelevant, since even if my vote is non-superfluous, it can only contribute towards change if there is a potential change to contribute towards. If I am certain that no other individual will act, perhaps for fear of violence, torture or disenfranchisement, then my act remains irrelevant. This time it is irrelevant because there is no change for it to contribute towards, and thus it cannot make a difference despite remaining non-superfluous.

Nefsky may claim that this objection misunderstands her argument. According to Nefsky, my action is not dependent upon others also acting, and its non-superfluous nature does not result from a *belief* in actual

change, but that I can make a step towards creating change—much like my donation towards Shelter, a British housing charity, is unlikely to *solve* homelessness, but *can* take a step towards creating change. The creation or potential for change is not dependent on others, but something which I can create, if not contribute towards, in a non-superfluous way (Nefsky 2019, p. 11).

Putting this aside, however, the argument that individuals refusing to make the sacrifices outlined in the perception argument would be morally wrong is unconvincing. In my view, Nefsky falls victim to the fallacy that individuals—in isolation—can have any impact. This fallacy exposes a fundamental limitation in the traditional literature on problems of collective impact: both Kagan and Nefsky overlook why collective action problems are problematic in the first instance. To solve them, we must ask what is at the core of these problems. The answer to this question lies in the aggregation of individual actions, the collective consequence of which can have harmful effects.

Since the aggregation of individual actions is why collective impact problems are problematic in the first instance, no moral reason can be sufficient in motivating individual behaviour change that can aggregate a collective which is sufficiently large to make a difference. As a result, it must be shown that collective harm requires a collective solution—to claim otherwise is naïve. Even if Nefsky's argument convinces some individuals to change their behaviour, it will not convince a sufficient number of individuals for any significant change to materialise. In order to overcome this limitation, we must seek a solution to problems of collective impact which is not limited to the confines of individual action, but which considers the wider scope of collectives and those who profit from collective harm. This is what I intend to do in the final section of this article.

#### IV. WE CAN MAKE A DIFFERENCE

In the previous two sections, this article has considered two traditional solutions to problems of collective impact. These solutions are traditional in the sense that their authors wish to find reasons why individuals ought to change their behaviour. The first considered was Shelly Kagan's expected utility argument which uses a consequentialist framework to prove that individuals *might* make a difference. We then considered Julia Nefsky's recent work on non-superfluous contributions, the idea that I can make a difference, even if it is non-perceptible. I will now defend a sceptical response to problems of collective impact, presenting an argument which aims to prove that we, collectively, can make a difference.

Before presenting this sceptical argument, I will address why such an approach is appropriate, and how it might help us provide solutions to problems of collective impact. First, the way in which we organise ourselves, in communities, societies, nations and nation-states, must be acknowledged. It is this organisation of individuals which I believe is key to solving problems of collective impact. We are, as proven by these groups, social beings. Therefore, it is possible to imagine a collection of individuals who might be able to assume the responsibility of solving collective harms.

Second, I believe a sceptical solution of this sort—reducing responsibility to the collective, not the individual—is a more realistic approach to problems

of collective impact. Collective harm is bad because of the many millions of individuals acting in a particular way which, on their own, produce no perceptible harm. If we reverse engineer this issue, the collective good that is possible through collective action seems the only realistic option to solve collective harms, given the scale at which they occur.

Finally, this second point raises the need for an important clarification: I would like to distinguish reducing *responsibility* to collections of individuals and reducing *action* to collections of individuals. In presenting this sceptical approach I am not denying that individuals will have to change their behaviour, nor am I claiming that individual actions are superfluous—in fact, I strongly agree with Nefsky that individual actions can make a non-superfluous contribution towards change, even though this is insufficient in producing solutions on the scale required to combat collective harm. What I am attempting to do is to limit the *responsibility* of this harm to the collective, so that no individual *ought* to change their behaviour for moral reasons. Instead, with responsibility in the collective domain, it is up to collections of individuals—most likely governments, international agencies, regulators, local communities, and corporations—to incentivise change on an individual level, as well as changing the behaviours of those corporations who profit from collective harm.

Now that I have addressed the relevance of sceptical solutions in solving problems of collective impact, how such solutions might help reduce collective harms, and distinguished collective responsibility from collective action, I shall proceed to outline my sceptical argument.

This argument is two-pronged. First, it seems that given the scale of our economies, individuals are powerless unless they act as a collective—individual actions considered in isolation cannot make a difference. This was evident in the examples listed in Section I, such as buying clothes from a retailer which treats its workers well instead of one which exploits them; voting in an election with a certain harmful outcome instead of abstaining; and taking the bus instead of driving my car. The negligible act of the individual was also exemplified in the second objection to Kagan's arguments in Section II. Consequently, individuals acting alone *cannot* be morally wrong. Instead, collections of individuals have a responsibility (moral obligation) to prevent and solve collective harm in these instances, since it is only these collectives that have sufficient power to make a difference. Therefore, the collectives in which we organise ourselves—governments, both local and national, international organisations, corporations, charities, and many others—carry the burden of responsibility for solving collective harms when no individual act is sufficient in making a difference.

Secondly, in our consumer-centric society, every product has *some* value—there will always be a market for any good produced. As such, if corporations (firms or businesses) profit monetarily from goods or services which create or contribute towards collective harm—either in the production, use or disposal of these goods or services—then they also have a moral obligation to reduce these collective harms.

Therefore, individuals do not act wrongly in problems of collective impact, since the moral responsibility for

preventing harm falls on collections of individuals—like governments—and those firms who profit from collective harm. As previously stated, this is not to say that individuals should not be expected to change their actions, but only that the reasons for changing their actions should not result from individual moral obligations. Instead, a change in actions should be decided and incentivised by that of the collective in which they are organised.

This argument builds upon the work of Walter Sinnott-Armstrong (2010), who argued that our individual moral obligation in problems of collective impact is to get governments to do their job. Instead of selling my heavily-polluting sports car—or making other significant sacrifices, like switching to oat milk instead of dairy—I should continue to drive my heavily-polluting sports car whilst campaigning the government to change policy so that driving such sports cars would be illegal (Sinnott-Armstrong 2010, p. 304). After all, it is the collective that has the power to create such significant harm, and so it is the collective—in the form of, or represented by, the government—which has the power to do something about it. One could argue that expressing views through our actions is of equal importance, and so I should not continue to drive my sports car. However, the impact of driving my sports car is so insignificant and imperceptible that my individual action does not matter. What matters is my ability to change the agenda and successfully campaign for a change in government policy, that is, to successfully campaign for collective action.

This collective approach, argued by Sinnott-Armstrong and developed by my own argument, finds strength in humanity's natural ability to aggregate, a feature which both Nefsky and Kagan overlook; one that is at the core of collective impact problems. The responsibility of corporations who also profit from these collective harms is an important nuance which I believe crucial to my sceptical approach.

To illustrate this sceptical approach, let us consider what it means for the examples given in Section I. In the first example, *ethical consumerism*, it means, firstly, that governments and collective institutions (such as regulators) have a responsibility to promote clothing production which does not take advantage of its workers. This could be done through labour protection policy, a minimum standards policy for goods, trade restrictions on countries whose labour laws allow such exploitation, or policy which promotes transparent supply chains, an example of which can be found in the *Modern Slavery Act 2015* (Home Office 2018). Secondly, those corporations who profit from exploitation of workers—the collective harm in this instance—must also share the responsibility of reducing exploitation. This could be done through improving their environmental, social, and corporate governance (ESG) practices: changing suppliers, improving pay and conditions of workers, or ensuring transparent supply chains.

The second example, *elections and referendums*, could be seen as more problematic for this sceptical approach. However, the sentiment remains the same. It is up to institutional arrangements, drawn up by the collective of individuals, the government, to prevent candidates from enacting potentially dangerous policies—perhaps through checks and balances—or to

increase competition in electoral systems to raise the stakes and incentivise voter participation. In the first instance, preventing harmful candidates from obtaining office or harmful policies from coming into effect would not be the responsibility of individuals, but the institutional safeguards enacted by the collective. In the second instance, individuals should be incentivised to participate by improvements in electoral competition.

Finally, in the example of *climate change*, it is again the responsibility of collectives—in this instance multiple governments across the world and the international institutions in which they organise, as well as those corporations who profit from the harmful effects causing climate change—to act. In many ways we have already seen these collectives taking responsibility through climate conferences, like the Conference of the Parties of the UNFCCC (COP); investments into state infrastructure and public transport; as well as corporations switching to renewable energy sources and investing in climate-friendly methods of production. Individuals can be incentivised by governments to walk or use public transport instead of driving, but this can only be done if the right infrastructure has been delivered by the government, the collective. A similar point can be made about heating. Individuals can be incentivised by governments to insulate their homes or use more environmentally-friendly heating techniques, but this can only be done if the government has reduced the costs of these technologies and made them easily accessible.

These three examples highlight the ability of collectives to solve problems of collective harm in instances where individual actions can make no difference. Many of these examples are evident in real life, as with the *Modern Slavery Act 2015* and COP, two government initiatives that have already been mentioned. Therefore, we can see that the collective approach to solving these problems is already embodied in our everyday lives. This must add to the validity of this approach.

Reducing responsibility to the collective in these instances is not the same as claiming individuals are not responsible for their actions. The responsibility for solving and preventing collective harms falls only upon the collective in instances where individual actions make no difference. Furthermore, individuals may still be required to change their behaviour, as explained earlier in this article, but the reasons for doing this come not from a moral argument of obligation at an individual level, but from the incentives of the collectives in which we organise ourselves.

## V. CONCLUDING REMARKS

In this article we have considered various solutions to problems of collective impact. We first defined these problems of collective impact, before presenting three examples of these problems: ethical consumerism, voting and climate change. I then argued that problems of collective impact are important in the study of politics, and society more widely, if we wish to solve some of the world's most pressing issues, such as: global inequalities, food poverty, and climate change. We then considered three solutions to these problems. The first two, from Shelly Kagan and Julia Nefsky, were from the traditional literature on this subject, and tried to find moral reasons for why individuals ought to change their behaviour. The third solution was sceptical; it sought

to prove that the responsibility for solving problems of collective impact falls not upon individuals, but the collectives in which we organise ourselves—primarily governments—and those corporations who profit from collective harm.

This article argued that collective harm requires collective responsibility. This does not contradict the idea that individuals are responsible for their actions, or that individuals can be required to change their behaviours in order to reduce these harms. Instead, the collective is responsible only in instances where no single act can make a difference, and individuals will be required to change their behaviour, not for

moral reasons, but through incentives provided by the collective.

So, should I, as an individual, seek to act in a more ethical way—to vote in an election to prevent a harmful political actor from gaining power, or, more simply, to use oat milk instead of soy? This article says yes, if you wish to. Your individual actions can, in the words of Julia Nefsky, make a non-superfluous contribution towards change (Nefsky 2019, p. 10). But you should feel no obligation towards performing these actions. Instead, the responsibility for preventing problems of collective impact falls upon the collective, not the individual.

Home Office. (2018). *Modern Slavery Act 2015*. GOV.UK, 30 July. Available at: <https://www.gov.uk/government/collections/modern-slavery-bill>

Kagan, S. (2011). 'Do I make a difference?', *Philosophy & Public Affairs*, 39(2), pp. 105–141. <https://doi.org/10.1111/j.1088-4963.2011.01203.x>

Kateman, B. (2019). 'Vegan restaurants are on the rise', *Forbes*, 21 August. Available at: <https://www.forbes.com/sites/briankateman/2019/08/21/vegan-restaurants-are-on-the-rise/?sh=56d5fdc51e80>

Nefsky, J. (2011). 'Consequentialism and the problem of collective harm: a reply to Kagan', *Philosophy & Public Affairs*, 39(4), pp. 364–395. <https://doi.org/10.1111/j.1088-4963.2012.01209.x>

Nefsky, J. (2018). 'Consumer choice and collective impact', in Barnhill, A., Budolfson, M., and Doggett, T. (eds.) *The Oxford handbook of food ethics*. Oxford University Press, pp. 267–286. <https://doi.org/10.1093/oxfordhb/9780199372263.013.23>

Nefsky, J. (2019). 'Collective harm and the inefficacy problem', *Philosophy Compass*, 14(4), Article e12587.

<https://doi.org/10.1111/phc3.12587>

Sinnott-Armstrong, W. (2010). 'It's not my fault: global warming and individual moral obligations', in Gardiner, S. M., Caney, S., Jamieson, D., and Shue, H. (eds.) *Climate ethics: essential readings*. Oxford University Press, pp. 332–346. ISBN: 9780195399622.

The Economist. (2023). 'How the young spend their money', 16 January. Available at: <https://www.economist.com/business/2023/01/16/how-the-young-spend-their-money>