

## INEFFICACY, PRE-EMPTION AND STRUCTURAL INJUSTICE

NIKHIL VENKATESH

Many pressing problems are of the following kind: some collection of actions of multiple people will produce some morally significant outcome (good or bad), but each individual action in the collection seems to make no difference to the outcome. These problems pose theoretical problems (especially for act-consequentialism), and practical problems for agents trying to figure out what they ought to do. Much recent literature on such problems has focused on whether it is possible for each action in such a collection to make such a tiny impact on the world that it makes no expected difference to the outcomes with which we're concerned. I argue that even if this is impossible, there are cases in which each action makes no difference, not because it has such a tiny effect on the world, but because if it were not performed, a similar action would be. This recognition allows us to connect these problems with discussions of structural injustice.

### I

*Introduction.* Sometimes, some collection of actions will produce some morally significant outcome (good or bad), but each action in the collection seems to make no difference to the outcome. Such cases—call them ‘inefficacy cases’—include pressing issues such as:

*Voting.* A good candidate will beat a bad candidate if and only if a large number of people vote for her. But unless the election is decided by a single vote—which is unheard of in large-scale elections—how could one voter's actions make a difference to the outcome?

*Climate change.* If every person in the global north reduced their CO<sub>2</sub> consumption by a sufficiently large amount, this would significantly reduce the harm done by climate change. But how could one household's emissions make a difference?

*Meat consumption.* If lots of people became vegetarians, fewer animals would be reared and slaughtered. But would any farmer reduce their production simply because one more individual has become a vegetarian?

Inefficacy cases pose a difficulty for act-consequentialism. Act-consequentialists hold that the moral status of actions derives entirely from the difference they make to outcomes. If an action is inefficacious, it is difficult to see how act-consequentialist theories can condemn or condone it. And sometimes in such cases we intuitively think it appropriate to do so: we want to say that most people in the global north should reduce their CO<sub>2</sub> consumption, and that voters usually ought to vote. Many act-consequentialists advocate vegetarianism as a means of reducing harms done to animals (Norcross 2004; Singer 1980). Given the value of the outcomes at stake in these cases, act-consequentialism ‘appears to fail even in its own favoured terrain’ (Kagan 2011, p. 117), the production of the good.<sup>1</sup>

The problem is not just for theorists. Most of us have two intuitions, which inefficacy cases place in tension with one another. One is that we ought to participate in collective efforts to make things better. The other is that the difference an action would make is very important to whether we ought to do it. This tension is partly responsible for paralysis, in individuals and in society, about what to do about certain pressing issues such as those above.

Some suggest that these cases are illusory: actions such as voting, cutting one’s emissions, and going vegetarian do make a difference—or at least, an expected difference. I argue that even if this is true with respect to one set of seemingly inefficacious actions—those that make a very small impact on the world—it is not with respect to another—those that are pre-empted.

## II

*Are Inefficacy Cases What They Seem?* Some inefficacy cases are merely apparent: we overlook the difference individual actions make due to faulty reasoning. Examples of such reasoning include:

*‘The problem will still exist, whatever I do. Therefore, what I do makes no difference’.* Many problems can exist at different degrees of severity. For instance, climate change will occur whatever you do. But your actions could affect how much the climate changes, and therefore how much harm eventuates.

---

<sup>1</sup> Inefficacy cases pose problems for theories other than act-consequentialism too—but for reasons of space I will focus on act-consequentialism.

*'The problem is very big, and my impact upon it is comparatively small'*. Small effects are still effects. Moreover, often what appear to be small effects in the context of big problems may not be so small. The problem of global poverty is huge. But if my actions could help one poor child to get an education, this could have a big effect on their life. *'If very few people took this action, this would make no difference. Therefore, what I do makes no difference'*. Consider an election. The good candidate will win only if she gets the majority of the votes. If very few people vote for her, she will lose. But it's possible that very few votes could make a difference. It may be that exactly half of your fellow voters will vote for her, and exactly half for the bad candidate. Then your vote alone would tip the balance.

We might worry that in this last case that unless there is exactly one vote in it, every vote is inefficacious. If she's more than one vote behind, she loses even if one more person votes for her. If she's more than one vote ahead, she wins even if one fewer does. Sceptics of inefficacy here appeal to *expected difference*. Act-consequentialism will condemn pulling the trigger of a gun directed at a person, where only one of several chambers in the gun is loaded, even though it is unlikely that harm will occur. The chance of harm renders the expected difference negative. Similarly, if there is a chance that the election comes down to one vote, your vote might make a difference, that is, it makes an expected difference.

Several philosophers suggest that all apparent inefficacy cases can be given a similar analysis (Kagan 2011; Norcross 2004; Singer 1980). There are some thresholds (ten thousand votes, five hundred parts per million of CO<sub>2</sub>, twenty-five chickens purchased) at which the outcome changes (the better candidate wins, a hurricane destroys Barbados, a chicken farm is extended). It is possible that the actions of other people will take us right up to a threshold that your action will trigger. If it does, it is true of your action that if it hadn't been performed, the outcome would have been different.

The most prominent opponent of this argument is Nefsky (2011, 2015, 2017, 2019). She appeals to cases in which individual actions have such a tiny impact on the world, that it seems they could not trigger a threshold with respect to the outcome in which we're interested. One is Parfit's *Drops of Water* case (Parfit 1984, p. 76), which she adapts as follows:

Imagine that there are ten thousand men in the desert, suffering from intensely painful thirst. We are a group of ten thousand people near the desert, and each of us has a pint of water. [A] cart will be driven into the desert, and any water in it will be evenly distributed amongst the men. If we pour in our pints, the men's suffering will be relieved. The problem is, though, that while together these acts would do a lot of good, it does not seem that any individual such act will make a difference. If one pours in one's pint, this will only enable each man to drink an extra ten thousandth of a pint of water. This is no more than a single drop, and a single drop more or less is too minuscule an amount to make any difference to how they feel. (Nefsky 2017, p. 2744)

Nefsky claims there is no chance that your pint will push us above a threshold that causes the thirsty men to be better off. If there are thresholds, they are vague such that an additional pint will never be the difference between them being triggered or not: one drop per man simply makes too tiny an impact. Proponents of the expected difference view respond that if we allow that an extra drop makes no difference at all, we will end up with a Sorites paradox (Arntzenius and McCarthy 1997; Kagan 2011). Nefsky replies that the Sorites paradox is a paradox to be confronted, not a *reductio* to be wielded (2011, 2019).

Nefsky's view might lead to other paradoxes (Barnett 2018). But the expected difference view, even if it can avoid Nefsky's criticisms, has its limits. Hedden (2020) has shown that it fails in cases involving infinities. And we are often fairly sure that an action will not trigger a threshold (Budolfson 2019), or that the threshold it's likely to trigger (say, taking someone from an extremely thirsty to merely very, very thirsty state) makes a very small difference (Nefsky 2019, §4.2). Then, the expected difference account may give us too little reason for or against actions we think are morally significant. But perhaps our intuitions about such significance ought to be revised (Norcross 1998).

I will not adjudicate this debate. Instead I want to discuss another kind of inefficacy case, which is real even if Nefsky is wrong about tiny impacts.

### III

*Pre-emption.* In pre-emption cases, if one fails to perform an action, another person will perform a similar action with equivalent

consequences. Thus one's action would be inefficacious: it makes no difference to (the value of) the outcome. This need not be because it has a tiny impact on the world. The action might make a big impact on the world, but in its absence a similar impact would be made by a different action.

Williams (1973, pp. 97–8) provides an illustration of pre-emption. George is offered a job conducting research into chemical warfare. He thinks chemical weapons ought not to be developed—let's suppose this is true, and that it would be best if no scientist worked in such areas. George is told that if he doesn't take the job, another scientist will. Assuming that this other scientist will conduct equivalent research, and bracketing other considerations, act-consequentialism cannot condemn George for taking the job. George's action is inefficacious: the work will go on just the same whether it is done by him or the other scientist.

Expected difference cannot help us. This job triggers a threshold: if someone takes it, that will significantly affect the outcome with which we're concerned (the development of chemical weapons). But George's taking it does not make it more likely that someone takes it; that is, by stipulation, certain. If he doesn't trigger the threshold, someone else will. So the expected difference of his taking the job is nil.<sup>2</sup>

Consider climate change. One might think that one's individual emissions make no difference to climate change because they have such a tiny impact on the levels of CO<sub>2</sub> in the atmosphere, too tiny to cause significant climate-associated harms (Kingston and Sinnott-Armstrong 2018). The expected difference view counters that your tiny increase might push CO<sub>2</sub> levels over a threshold at which such a harm is caused (Broome 2019). Therefore, emitting any amount of CO<sub>2</sub>, however small, makes a negative expected difference—and one ought, other things being equal, to reduce one's emissions.

One could question this by invoking vague thresholds and Sorites. But perhaps our personal carbon emissions are to a great degree

<sup>2</sup> Perhaps George cannot be sure that a rival chemist will take the job if he refuses. Then it looks as though the expected difference of his taking the job is negative. However (as in Williams's original case), perhaps other scientists would pursue even more dangerous research in this role than George would—and perhaps he is more in need of the salary. These possibilities might cancel out the possibility of rival chemists also refusing the job, such that the expected difference remains nil.

pre-empted.<sup>3</sup> In that case, your individual emissions would not add to the quantity of CO<sub>2</sub> emitted—they just mean that more is emitted by you and less by others. Here are two ways in which emissions could be pre-empted. One is if all carbon that can be emitted will be emitted. There's a finite supply of fossil fuels. Unless some is left in the ground, the peak of CO<sub>2</sub> in the atmosphere is wholly determined by the size of this supply. If it is to be burnt anyway, my emitting will not incentivize any more burning. The second mechanism is if carbon emissions are regulated by a cap-and-trade scheme. Under such a scheme, firms buy permits to emit a given amount of CO<sub>2</sub>. If I reduce my emissions, the firms that would have supplied me with carbon-intensive products will buy fewer permits. This will make those permits available to other firms, who will buy them and emit equivalent amounts of CO<sub>2</sub>.<sup>4</sup>

It's unlikely that pre-emption will render individual emissions wholly inefficacious. Not all emissions require permits, and there is a chance that some fossil fuels will be left in the ground. Slowing the pace of emissions may reduce harm. But probably, owing to pre-emption, reducing your emissions by one tonne will reduce total emissions by less than a tonne, and thus reduce the harm done by climate change by a smaller amount than Broome and others who ignore pre-emption estimate. And here, given the personal costs of reducing emissions, act-consequentialism may not endorse individual reductions, even when it would be best if we all reduced our emissions. Broome estimates that going for a drive in an SUV does expected harm worth \$1 (2019, p. 111). If we account for pre-emption, that expected harm will be even less.

Pre-emption also arises in cases of structural injustice, an area that has seen much recent discussion (relatively little of it from consequentialists). Take the canonical example of structural injustice, Young's (2006) analysis of the global clothing industry. Many workers suffer serious injustices, including overwork, low pay and unsafe conditions. These harms are the direct effects of those who employ or manage them—usually small enterprises in poor countries, dependent on larger exporters, who are in turn dependent on large multinationals. However,

<sup>3</sup> The closest I can find to this argument in the literature is in [Maltais \(2013, pp. 591–7\)](#) and [Cripps \(2013, pp. 123–5\)](#). They do not closely distinguish between tiny impacts and pre-emption (and use the language of overdetermination), which I think leads [Broome \(2019\)](#) to mischaracterize and dismiss their arguments.

<sup>4</sup> Some environmentalists buy permits and retire them, showing awareness of this problem.

each of the links in the chain believes itself to be operating close to the margin in a highly competitive environment, and usually is under heavy pressure to meet orders at low cost by firms higher up the chain. (Young 2006, p. 110)

If a manufacturer improved the condition of the workers, they would be out-competed by a rival who would treat workers more harshly. If a multinational paid manufacturers more, allowing them to improve working conditions, they too would be out-competed. The same harm would go on, then, whatever they do. So this is a pre-emption inefficacy case.

Young recognizes a similar tension to that facing the act-consequentialist: not wanting to blame agents for actions that made no difference to the outcome, but wanting to condemn that outcome and provide moral obligations to change it. Young, and the literature following her (see, for example, Wollner 2019; McKeown 2016; Zheng 2021), emphasizes the importance of groups, structures and politics for individual morality, and distinguishes different species of moral responsibility.

#### IV

*Conclusion.* Act-consequentialists cannot ignore inefficacy cases simply by invoking the expected difference view and the Sorites paradox. Such cases could arise through pre-emption and pose the very same problem: we want to condemn or condone actions that make no difference to outcomes. And pre-emption is widespread: as Pinkert says, cases in which agents ‘know that were they to stop participating, others would step in and compensate by taking over their transactions’ are very common in market economies (2015, p. 978).

Those of us who are sympathetic to act-consequentialism, then, should be investigating whether the theory can be revised to account for such cases. And in doing so we should engage with the burgeoning literature on structural injustice. The solution to inefficacy cases for the act-consequentialist may turn out to be, like Young’s solution to cases of structural injustice, political, collective and revisionary of both traditional moral concepts and social conditions.<sup>5</sup>

---

<sup>5</sup> I would like to thank Véronique Munoz-Dardé, Ulrike Heuer, Joe Horton, Han van Wietmarschen, Peter Railton, Sarah Buss, Ariana Peruzzi, Luke Devereux, Kacper Kowalczyk, and an anonymous reviewer for their guidance in this research, as well as the

*Philosophy, Logic and Scientific Method*  
 London School of Economics  
 Houghton Street  
 London WC2A 2AE  
 UK  
[n.venkatesh2@lse.ac.uk](mailto:n.venkatesh2@lse.ac.uk)

## REFERENCES

- Arntzenius, Frank, and David McCarthy 1997: 'Self Torture and Group Beneficence'. *Erkenntnis*, 47(1), pp. 129–44.
- Barnett, Zach 2018: 'No Free Lunch: The Significance of Tiny Contributions'. *Analysis*, 78(1), pp. 3–13.
- Broome, John 2019: 'Against Denialism'. *The Monist*, 102(1), pp. 110–29.
- Budolfson, Mark 2019: 'The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response'. *Philosophical Studies*, 176(7), pp. 1711–24.
- Cripps, Elizabeth 2013: *Climate Change and the Moral Agent: Individual Duties in an Interdependent World*. Oxford: Oxford University Press.
- Hedden, Brian 2020: 'Consequentialism and Collective Action'. *Ethics*, 130(4), pp. 530–54.
- Kagan, Shelly 2011: 'Do I Make a Difference?' *Philosophy and Public Affairs*, 39(2), pp. 105–41.
- Kingston, Ewan, and Walter Sinnott-Armstrong 2018: 'What's Wrong with Joyguzzling?' *Ethical Theory and Moral Practice*, 21(1), pp. 169–86.
- McKeown, Maeve 2016: 'Global Structural Exploitation: Towards an Intersectional Definition'. *Global Justice: Theory Practice Rhetoric*, 9(2), pp. 155–77. Online at <https://www.theglobaljusticenetwork.org/index.php/gjn/article/view/116/96>.
- Maltais, Aaron 2013: 'Radically Non-Ideal Climate Politics and the Obligation to at Least Vote Green'. *Environmental Values*, 22(5), pp. 589–608.

---

audiences at the 2022 Joint Session of the Aristotelian Society and the Mind Association and at Michal Masny's transatlantic ethics network. This research was funded in part through the London Arts and Humanities Partnership, Longview Philanthropy, and the Forethought Foundation for Global Priorities Research. Moreover, I would like to take the opportunity to thank the Aristotelian Society (in particular, Holly de las Casas and Guy Longworth) for all that they have given me over the years.



- Nefsky, Julia 2011: 'Consequentialism and the Problem of Collective Harm: A Reply to Kagan'. *Philosophy and Public Affairs*, 39(4), pp. 364–95.
- 2015: 'Fairness, Participation, and the Real Problem of Collective Harm'. In Mark Timmons (ed.), *Oxford Studies in Normative Ethics*, 5, pp. 245–71. Oxford: Oxford University Press.
- 2017: 'How You Can Help, Without Making a Difference'. *Philosophical Studies*, 174(11), pp. 2743–67.
- 2019: 'Collective Harm and the Inefficacy Problem'. *Philosophy Compass*, 14(4), p. e12587.
- Norcross, Alastair 1998: 'Great Harms from Small Benefits Grow: How Death Can Be Outweighed by Headaches'. *Analysis*, 58(2), pp. 152–8.
- 2004: 'Puppies, Pigs, and People: Eating Meat and Marginal Cases'. *Philosophical Perspectives*, 18(1), pp. 229–45.
- Parfit, Derek 1984: *Reasons and Persons*. Oxford: Oxford University Press.
- Pinkert, Felix 2015: 'What If I Cannot Make a Difference (and Know It)?' *Ethics*, 125(4), pp. 971–98.
- Singer, Peter 1980: 'Utilitarianism and Vegetarianism'. *Philosophy and Public Affairs*, 9(4), pp. 325–37.
- Williams, Bernard 1973: 'A Critique of Utilitarianism'. In J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against*, pp. 77–150. Cambridge: Cambridge University Press.
- Wollner, Gabriel 2019: 'Anonymous Exploitation: Non-Individual, Non-Agential and Structural'. *Review of Social Economy*, 77(2), pp. 143–62.
- Young, Iris Marion 2006: 'Responsibility and Global Justice: A Social Connection Model'. *Social Philosophy and Policy*, 23(1), pp. 102–30.
- Zheng, Robin 2021: 'Moral Criticism and Structural Injustice'. *Mind*, 130(518), pp. 503–35.