Against Sinnott-Armstrong

by Chad Vance
excerpted from
"Justifying Subsistence Emissions: An Appeal to Causal Impotence" (2023, *Journal of Value Inquiry*)
and
"Climate Change, Individual Emissions, and Foreseeing Harm" (2017, *Journal of Moral Philosophy*)

Causal Impotence and the Group Harm Principle

(from Vance, 2023)

Now, either our individual emissions cause harm, or they do not. Some believe, as John Broome does, that we have "no reason to doubt that every bit of emission that you do cause is harmful." (2012, 77) But, what if that is mistaken? For, there is, I think, a compelling argument to the contrary.[1] It goes as follows: The problem of climate change is so massive, it is said, that our emissions are merely 'a drop in the bucket', so to speak—or more aptly, a drop in a *flood*. Imagine a destructive flood averaging one meter deep, and covering an area over 40 times larger than Vatican City. As the flood is washing away homes and taking lives, you, using an eye-dropper, add a single drop of water to that flood. How much additional harm do you think this would cause? The most plausible answer is: *None*.

But, in terms of quantities, this is precisely analogous to putting one joyride's worth of CO_2 into the atmosphere.[2] In other words, while an act of joyriding clearly adds some non-zero amount of *greenhouse gases* into the atmosphere, it makes no difference at all to the *amount of harm* that results from climate change. The result is a collective action problem, where, though our individual actions make a perceptible difference along *some* dimension (e.g., the amount of CO_2 in the atmosphere), they do not make a perceptible difference along any dimension that is *morally significant* (e.g., harm). Therefore, we have no moral obligation to reduce our individual greenhouse gas emissions after all.

That is how the causal impotence argument generally goes. To further illustrate its central claim, perhaps it will be helpful to consider the following case (adapted from Parfit 1984, 80):

> **Harmless Torturers** An innocent person is hooked up to a device, which shocks them with increasing amounts of electric current as a dial is turned. Turning the dial once increases the amount of current delivered to its victim by such a small amount that its effects are imperceptible. (Turning it 1,000 times, however, results in severe pain.) You are one of 1,000 torturers, each of whom turns the dial only once. The result is that the victim suffers severe pain.

Clearly, the torturers *collectively* cause significant harm. But, what is the causal impact of any particular individual's contribution? The answer cannot be that each turn of the dial causes some very small amount of harm. For, as the case is described, turning the dial only once causes *no perceptible difference*.[3]

---

[1] The most famous instance of this argument is found in Walter Sinnott-Armstrong's 2005. See also his 2018 follow-up (with Ewan Kingston).

[2] These numbers represent the use of one gallon of fuel, and are derived from Vance (2017, 563n), plus the size of the Vatican City (0.44 km^2).

[3] One might propose that there can be imperceptible increases in the amount of pain that one is experiencing, but this seem to me an oxymoron. As Kagan notes, "Indeed, it isn't even clear that it makes any *sense* to say that pain has been increased imperceptibly." (116) Nevertheless, perhaps it is still the case that one can be *harmed* imperceptibly, as

We might instead claim that there are thresholds[4]; e.g., perhaps the victim notices no change from zero to one, or two, or three, or four—but *does* perceive an increase in pain when the dial is turned a *fifth* time). But, again, this is *not* how the case is described. To clarify, let us stipulate that, in Harmless Torturers, for *any* two adjacent dial settings, n and n+1, the difference in the amount of current administered is so small that there is *no perceivable difference* to the victim. If there were determinate thresholds, however, there *would* sometimes be a perceivable difference—namely, at every threshold.[5]

So, how much harm *does* each individual turn of the dial cause? Admittedly, given the case as described, the only possible answer is quite odd: *None*. In Harmless Torturers, no single torturer causes any harm whatsoever. What is more, there *is no fact of the matter* about when the victim experiences an increase in pain. Rather, it is *indeterminate*. Julia Nefsky describes such cases in the following way:

> What is distinctive about nontriggering cases is that no single act serves as a trigger. So, the structure is not that of a tipping point: there is no precise point at which a limit is hit and the next act triggers a change in morally relevant outcome. Instead, the boundaries between one morally relevant outcome and another are vague, and so the difference between *n-1* and *n* acts of the relevant type can never, no matter what *n* is, make the difference between one morally relevant outcome and another. (2011, 377-378)

Applying this observation about the causal structure of certain collective action cases to SUV joyrides, Kingston and Sinnott-Armstrong write, "a plucked hair makes no difference to a person's baldness. The same approach should be taken with regard to joyguzzles." (2018, 180) They conclude that the harm resulting from climate change must be an *emergent* phenomenon, writing:

> In our view, climate change is emergent in this way. Just as individual molecules of oil do not cause parts of sensations of sliminess (or yellowish color), so individual molecules of greenhouse gas do not cause parts of dangerous climate impacts. Instead, as with the sliminess and color of oil, what increases the dangerous impacts of climate change is larger groups of molecules of greenhouse gases. (2018, 175)

The above constitutes, I think, a compelling argument for the conclusion that no single individual's greenhouse gas emissions cause any harm. But, perhaps the reader is still unwilling to accept this

---

Derek Parfit believes. (1984, 79) I find this more general claim implausible too. (Though I do believe that one can be *wronged* imperceptibly—for example, by a Peeping Tom whom the victim never finds out about.)

It is worth noting that Parfit qualifies his position as follows, writing, "Some people disagree. Even if we believe that there can be imperceptible harms and benefits, it may thus be better to appeal to what groups together do. This appeal is less controversial." (ibid., 82) He mentions this alternative again in a later unpublished work, writing, of an environmental collective action problem which he calls 'Harmless Polluters', "since the effect [of my pollution] on each will be … imperceptible, it may be hard to think of my act as seriously wrong. … It may help to remind myself that I am a member of a group who together do great harm." (1988, 28-29) Indeed, this alternative strategy is exactly the one that I champion in §5.

[4] Also called 'triggers' (e.g., Kagan, 2011); small-scale versions of 'tipping points' (see, e.g., Broome, 2012, 34). (I should also note that, strictly speaking, the previous suggestion is a thresholds scenario too—namely, one taken to its limit, such that *every* increase in the number of dial turns triggers an increase in the amount of pain.)

[5] Note that *some* instances of collective action *do* clearly have a thresholds structure. For instance, take voting: Most of our votes make no difference. But, there *is* some particular number of votes, n, such that one additional vote (n+1) will "tip the scales" in favor of the other candidate.

pessimistic picture of individual causal impact, on the grounds that it is too "metaphysically odd".[6] After all, the suggestion here *does* seem to be that zero plus zero plus zero plus… adds up to some very large number—and that is absurd! But, I should like to point out that, even if some of our individual emissions *do* cause harm, there is independent reason to believe that they nevertheless *never make any perceptible difference*. The argument for *that* conclusion comes from Mark Budolfson, who writes:

> greenhouse gas levels are currently accumulating, and will continue accumulating into the foreseeable future. To see why this is a problem for the argument, imagine that this week, as you engage in some emissions-generating activity, your emissions cause a catastrophic tipping point to be crossed. Nonetheless, even if you had avoided those emissions and thus hadn't tipped the scales yourself, it is certain that someone else's emissions would have tipped the scales at essentially the same time, because an entire planet of other people would still have been emitting at the same time even if you had not been emitting. This shows that, given the empirical facts, there is no chance that you could delay a catastrophic tipping point from being crossed today or in the foreseeable future by reducing your emissions, and thus there is no good reason for reducing emissions that arises from the possibility of tipping points being crossed now or in the foreseeable future. (ms, 36)[7]

Essentially, there is an overdetermination problem here, such that, even if some of our individual emissions *do* cause harm—whether in a linear way where *every* emission makes some small difference, or in a way involving thresholds, such that only *some* emissions make a (larger) difference—it would still turn out that no one's individual emissions make any *perceptible difference*. For, at best, our emissions only bring about the harms that they cause a mere fraction of a second sooner than they otherwise would have occurred, had we refrained from emitting them. …

[Now consider] the following case (from Sinnott-Armstrong 2005, 289):

> **Car Push** Four people are pushing a car off of a cliff with an innocent person trapped inside of it. It takes the strength of three people to push the car. You decide to help them push. The car goes over the cliff, and the person inside of it dies.

Here, it is clearly wrong to help the others to push the car. But, notice: Your individual contribution *makes no difference*. That is, with or without your contribution, the exact same amount of harm occurs either way. So, right away, we can see that it *is* at least sometimes morally wrong to contribute to a harmful collective action, even when one's contribution makes no difference to the amount of harm that occurs (i.e., when one's contribution is causally impotent with respect to the harm).

We might think that this opens up a route to morally condemning the production of luxury emissions. That is, we might wonder whether joyriding in a gas-guzzling SUV is morally wrong for the same reason that pushing the car in Car Push is wrong. Sinnott-Armstrong puts this suspicion to rest, however, by pointing out that pushing the car is morally wrong because you clearly *intend the harm*—a feature which is absent in the case of joyriding in an SUV. But, elsewhere (Vance, 2017) I

[6] See Hiller (2011, 349).

[7] Kingston and Sinnott-Armstrong voice a similar concern, writing, "given that we're constantly over-emitting—your refraining from joyguzzling only causes the threshold to be reached a fraction of a second later than it would have, had you joyguzzled. But, that's not morally significant." (2018, 177)

have argued that a causally impotent contribution to a collective harm can be morally wrong even when there is *not* malicious intent. For instance, consider the following variant of the case above:

> **Car Push (Light Exercise)** Four people are pushing a car off of a cliff with an innocent person trapped inside of it. It takes the strength of three people to push the car. They truthfully claim to be pushing the car only for the purpose of getting some light exercise—though foreseeing, of course, that together they will collectively cause the death of one person as a side-effect of their efforts. You are a bystander who happens to be jogging by just then. Correctly seeing that your individual contribution will make no difference to the amount of harm done, you help to push the car, also merely for the purpose of getting some light exercise. The car goes over the cliff, the one inside of it dies, and the rest of you are all a bit more fit.

Once again, no single individual's contribution makes any difference to the amount of harm that occurs. But now it is also stipulated that none of the individuals pushing the car intends any harm. Rather, the harm that occurs is merely an unintended (though foreseen) *side-effect* of the group's collective action. Yet, clearly it is still morally impermissible to push the car in this case. In my earlier work, on the basis of this sort of case, I argued at length for the conclusion that it is prima facie morally wrong to contribute to a harmful collective action, even when that contribution makes no difference, and even when one intends no harm, but only *foresees* that harm will result from the group's action. My conclusion was a modified version of what Sinnott-Armstrong called 'The Group Principle' (2005, 298), which roughly states:

> **The Group Principle:** It is prima facie morally wrong to perform an action if this action makes us a member of a group whose actions together cause harm.

Specifically, I argued for the conclusion that, *even if* your action makes no difference to the amount of harm that occurs, it is nevertheless prima facie morally wrong to:

(a) perform an action which makes you a member of a harmful collective activity, where
(b) the harm caused by the group's action is severe, and
(c) you either intend the harm that the group's action causes, or at least foresee it, and
(d) refraining from the action is not very costly (i.e., you will not suffer great harm by refraining).

This conclusion from my earlier work entailed that your contributions in both the original Car Push case as well as its Foreseeing Variant are morally wrong. But, more importantly, it also entailed that the production of *luxury emissions* is morally wrong—and that these are wrong *even if* our individual emissions cause no harm.[8] [9]

[8] The following worry for my proposal may have occurred to the reader at this point: What if I do *not* foresee that the collective action toward which I am contributing will cause harm – e.g., because I am a climate change denier? Are the luxury emissions of climate skeptics morally permissible for this reason? I address this worry for my proposal at length in my (2017). But here I should simply like to reiterate that I am in agreement with Avram Hiller when he writes, "even if some individuals are ignorant of the expected effects of their actions, individuals ought not be ignorant" (2011, 353). For, I believe that there is at least some minimal duty to become informed about the effects of one's actions, or the collective actions to which one is contributing. Climate skeptics act wrongly because they have failed in this respect.

[9] Here is another potential worry: What if joyriding in SUV's is very important to me, and I would experience great sadness without it? In short, what if sacrificing my luxury emissions *is* very costly? Are they permissible in this case?

The Noisy Airport Objection
(from Vance, 2017)

This principle delivers the intuitive verdict that the agents act wrongly in the Foreseeing Variants of Car Push and Harmless Torturers. However, if modified in this way, the principle no longer fails to deliver the verdict that we have no moral obligation to reduce our individual greenhouse gas emissions. For, emitting such gases makes one a member of a group whose actions are together causing harm (namely, the harms that result from climate change). Furthermore, though, e.g., a Sunday joyrider does not *intend* harm by driving, it is foreseen with near certainty that harm will occur as a result of the group's emissions. In short, according to [The Group Principle], it is morally wrong to go for a Sunday joyride in an SUV. …

However, Sinnott-Armstrong will likely reject this [Group Harm Principle], since it seems to have unacceptably strong implications elsewhere. Consider this case:

> **Noisy Airport** Suppose that everyone in an airport is talking loudly. If only a few people were talking, there would be no problem. But the collective effect of so many people talking makes it hard to hear announcements, so some people miss their flights. Suppose, in these circumstances, I say loudly (but not too loudly), "I wish everyone would be quiet." (298)

Here, you do not *intend* harm, though you do (I will assume) foresee with some non-negligible degree of probability that harm will result from the group's action. Furthermore, the benefit is minor (complaining to your friend), while the harm is not (passengers missing flights). The principle that I have advanced seems to deliver the verdict that one acts wrongly in this case. However, Sinnott-Armstrong explicitly rejects this conclusion, writing,

> My speech does not seem immoral, since it alone does not harm anyone. Maybe there should be a rule (or law) against such loud speech in this setting (as in a library), but if there is not (as I am assuming), then it does not seem immoral to do what others do, as long as they are going to do it anyway, so the harm is going to occur anyway. (298)

This objection reminds me of a criticism of John Arthur's, which Peter Singer responds to in the context of his argument that we ought to give up our luxury goods and donate to famine relief (1972, postscript). (To be honest, it also reminds me a little of a spoiled teenager who, phone privileges having just been revoked, insists loudly, "You're ruining my life!") In response to Singer's weaker moral principle – *If we can prevent something very bad from happening without sacrificing anything of moral significance, then we ought to do so* – one might object in the following way: "I could save a drowning child at the expense of ruining my very expensive pants, but losing these pants would be morally significant to me. So, I have no duty to save the child."

In reply, Singer expresses that he had hoped to rely on the reader's intuitive notion that luxury goods are of no moral significance, without offering a fully-developed theory of moral significance. Here, I have hoped for the same, with respect to luxury emissions. I should also like to point out though, as Singer does, that at the very least, my principle uncontroversially rules out the luxury emissions of those who, *by their own admission*, readily accept that the goods obtained via their emissions are of no moral significance.

What is more, it seems to me that there also exists some minimal duty to cultivate in one's self the sorts of desires and goals that provide happiness or fulfillment *without* being "eco-gluttonous" – i.e., without producing excessive greenhouse gases or other environmental pollution. If there is such a duty – and I believe there is, though I do not have the space to argue for it here – then we might reply to the individual who has developed a deep and significant attachment to SUV joyriding that they really ought to wean themselves away from such a desire and cultivate instead some desires and interests that contribute less significantly to harmful collective actions.

I take issue with this verdict, however. For, if the case were described such that the benefit was truly minor or non-existent, the harm was truly significant, and the degree to which one foresaw that significant harm would occur was truly near certainty (in short, if Noisy Airport was like the Foreseeing Variants of Car Push and the Harmless Torturers in all of the morally relevant respects), then we should easily agree with the verdict that one acts wrongly in Noisy Airport. To illustrate, let us fill in some of the details so that Noisy Airport clearly has these features:

> **Noisy Airport (With Further Details)** You are in a crowded airport, where everyone is talking. It is so loud, that attempts by airport staff to page one of the passengers on the intercom cannot be heard. As it turns out, the passenger being paged is trying to deliver a life-saving kidney to an individual in need of a transplant. If she fails to hear the page, the deliverer will miss an important announcement about a gate change. This in turn will result in her missing her flight, and the kidney recipient will die. Airport staff have handed out fliers, which inform members of the crowd that harm will likely result from their noise. You read the flier, and, amidst the noise, say loudly to your friend, "I wish everyone would be quiet!" The deliverer misses her flight, and the recipient dies while awaiting transplantation.

My proposal clearly delivers the verdict that one acts wrongly in this case—but, I think, rightly so. When the benefit is negligible, the harm is significant (someone's death, in this case), and it is foreseen with near certainty that the group will cause harm, it seems morally wrong to participate in the harmful collective action, even when one's individual contribution is neither necessary nor sufficient for that harm. In Sinnott-Armstrong's original case, we are likely imagining a scenario which lacks these important features. Catching or missing a flight is not usually a life or death situation—not to mention, airports typically announce gate changes on screens found throughout the airport, rather than over intercoms. Therefore, we typically have little reason to believe that a group's noise will cause significant harm (i.e., we reasonably foresee with near-zero probability that harm will result from the collective noise of many voices). For this reason, it remains unclear whether one acts wrongly in Noisy Airport until further details are supplied.[10]

[10] Note that this is an admission of the presence of some vagueness. I am claiming here that, the less severe the harm is, and the less reasonably certain one is that harm will occur, the weaker one's duty to refrain from the action (and therefore, the more easily this wrongness is overridden, even by minor benefits). But, that seems right. Such vagueness is to be expected, since both foresight and the severity of a harm come in *degrees*.