

REPLY

The Demandingness of Individual Climate Duties: A Reply to Fragnière

Colin Hickey 

Utrecht University
Corresponding author. E-mail: c.j.hickey@uu.nl

Abstract

In this article, I respond to Augustin Fragnière’s recent attempt to understand the demandingness of individual climate duties by appealing to the difference between “concentrated” harm and “spread” harm and the importance of “moral thresholds”. I suggest his arguments don’t succeed in securing the conclusion he is after, even from within his own commitments, which themselves are problematic. As this is primarily a critical project, the upshot of this discussion is that if there is a defensible way to justify the intuition that the duty to reduce emissions can’t be overly demanding, it has to be found elsewhere.

In a recent article, Augustin Fragnière takes on the important task of trying to clarify the content of our climate duties. In particular, he attempts to justify the intuition that an individual’s duty to reduce emissions can’t be overly demanding, and must leave space for “personal life-projects”.

Fragnière situates his discussion among the standard-bearers of the demandingness literature, in order to motivate the general idea that morality can only demand so much of people before it starts cutting into a legitimate entitlement they have to live a life of their own (Singer, 1973; Williams, 1973; Scheffler, 1982). Given the structure of the modern world (and climate change in particular), because we are implicated in “a whole host of seemingly trivial daily activities that might contribute to harming distant strangers”, Fragnière worries that a straightforward application of the duty not to harm others would “doom” most of our personal projects (Fragnière, 2018: 648).¹

A number of climate ethicists, for instance, have tried to calculate the expected harm resulting from an individual’s emissions by taking their admittedly small share of the global emissions that lead to climate catastrophe and assigning the same share of the expected global consequences (costs, harms, etc.). Most famously, John Nolt suggests

¹From here on, all references to Fragnière will just include page numbers. It is worth flagging at the start that Fragnière doesn’t substantively engage with the prospect of offsetting. He mentions John Broome’s endorsement of offsetting in passing, but it isn’t clear whether he would accept something similar. Given the rest of the context, we can infer at least that he thinks offsetting doesn’t solve the key problem about demandingness.

that the average American's lifetime emissions amount to the "suffering and/or deaths of one or two future people" (Nolt, 2011).²

If true, the duty not to harm might justify a very strict duty to minimize carbon emissions; indeed, to a "painfully low" level that Fragnière thinks would likely undercut the viability of pursuing a life of our own, with our own projects.³

Fragnière thus embarks to find a way to "justify some limit to the applications of the no-harm principle . . . in the specific case of climate change" (648). While some have gestured at the lack of control individuals have over their emissions, or the high cost (interpreted as a setback in interests) of such demanding emissions reductions (650–51) in order to try to mitigate the stringency of the duty, Fragnière thinks there is a theoretical explanation about the nature of how different kinds of harming behaviors generate reasons for action and the duties that attach to them.

He argues that we can't apply the harm principle in a straightforward way because climate change is not a "paradigm case of harm-doing", by which he means a case where one causes the harm "alone and directly" (652–53). He thinks there are principled reasons why "contributing to unstructured collective harm is less wrong than causing the same amount of harm in the context of a paradigm case of harm-doing" (653).

The central argumentative move for justifying such a difference involves distinguishing between what he calls "concentrated" harm and "spread" harm. When dealing with harms in the statistical way Nolt and others do, linking an individual's share of emissions to a corresponding share of the global expected harm, the "amount of harm attributed to the individual agent [is taken to be] *concentrated* on a few victims" (653). But one could also distribute the individual's share of harm broadly, so that they bear a tiny share of each harm across all victims.⁴ In this way "the amount of harm attributed to the individual agent is *spread* over all the victims" (653).

Fragnière suggests, first, that given the nature of the climate problem we should interpret the harms associated with greenhouse gas emissions as "spread" harms. I will cast some doubt on this below. But second, and central for my purposes, he argues that causing spread harm is less wrong than concentrated harm, even if they look similar under a certain statistical lens.

The reason why we should accept the latter claim, Fragnière argues, is because "the seriousness of the harm imposed on one single victim, and therefore the wrongness of the act that brings it about, includes moral thresholds" (655). If we accept the existence of moral thresholds, among the ways in which we can inflict harm on others, there are certain lines that can be crossed, such as killing someone, causing permanent disability, or violating their basic human rights, in which "the moral wrongness of the act causing the harm increases abruptly, in a more than proportional way" (655). Fragnière opts not to give an account of which thresholds exist, where they are set, and how much they affect the wrongness of various harmful behaviors – which we will see below ends up

²I follow Fragnière to engage on his terms, but it's somewhat unfortunate the outsized attention Nolt's rough calculation gets, given that more careful and empirically informed work collecting harm estimates exists.

³It's less clear that Fragnière's worry is well motivated when the individual duty is interpreted as having *net zero emissions*, as John Broome suggests, rather than no emissions at all (Broome, 2012). But I won't pursue the point.

⁴If someone bears one two-billionth of all the emissions, one would bear one two-billionth of the responsibility for X's displacement, Y's starvation, Z's heat stroke, etc. for all victims.

a significant sticking point for the success of the argument. But at a conceptual level, if such thresholds are plausible, Fragnière suggests that:

it follows that for a given amount of harm done there is more risk that an important threshold is crossed when harm is concentrated on one victim than when it is spread across a large number of persons. As a result, all other things being equal, causing trivial harm to a lot of people (without crossing any threshold) is less wrong than causing very serious harm to one person. (655–56)

Given that he thinks we should interpret our responsibility for climate harms as spread harms distributed over a massive set of victims, Fragnière concludes that our duties are significantly less demanding than would be a duty to not kill two people (as Nolt might have it, if interpreted in a “concentrated harm” way). Fragnière thinks there is more room for personal projects to defeat or outweigh the reasons for reducing emissions, on account of their spread harm footprint, thereby preventing them from amounting to a strict duty (657).

Exactly how much is up for grabs? Fragnière doesn’t think we are totally off the hook because the spread harms of emissions still offer a pro-tanto duty (or at least reasons) against emitting, and some emissions reductions could be done without sacrificing much (or enough to outweigh). This pro-tanto duty might still require significant behavioral changes. Indeed, Fragnière suggests that our climate context indicates duties that are “in fact fairly demanding” (600). He thinks it is “unlikely that frivolous and trivial interests would suffice to outweigh the reasons to reduce one’s GHG emissions.” Probably that will only be done by interests that are central to one’s integrity and identity. Fragnière calls them “core interests”, but leaves unexamined what actually counts as a core interest (660).

While Fragnière’s arguments are an interesting, innovative, and challenging contribution to the growing literature that is grappling with our climate responsibilities, I think it is less clear they actually succeed in securing the conclusion he is after.

The issue rests on how the idea of harm thresholds and the spread-harm analysis interface with notions of responsibility in the case of climate change. Exactly how to interpret Fragnière in this regard isn’t always clear. But in order for the concept of thresholds to do the work he wants, it seems that Fragnière has to maintain that an individual’s emissions can push some victims over the harm thresholds he prioritizes. Usually when we are talking about responsibility for crossing thresholds, we are talking about causal thresholds. And indeed, the main principle Fragnière employs about thresholds and the difference between spread and concentrated harm is stated in terms of “causing trivial harm to a lot of people (without crossing any threshold) is less wrong than causing very serious harm to one person” (656).⁵ For this principle to apply to the climate case, we must still be talking about “causing” harm in the move from “trivial” to “serious”. Likewise, as the parenthetical caveat “without crossing any thresholds” indicates, it should at least be possible for an individual to cross thresholds in spread-harm cases. Otherwise, introducing thresholds is just a way of stipulating that one’s harm contributions can’t matter in a morally important way, which seems question-begging. Indeed, why should one think that crossing thresholds is the morally

⁵He also presents the general duty in causal language as a “pro-tanto duty not to cause harm through their contribution” (658).

important thing about my emissions when it is impossible regardless of how much I emit?

Now, in fact, I don't think the spread harm analysis actually fits the climate case and I don't think there is a plausible causal story to tell about individual emissions possibly crossing harm thresholds for individual victims, and so think climate morality should be looking in different directions, which I will say more about below. However, there is still a significant problem for Fragnière *even if we assume that some plausible story could be told to support this claim*.⁶

To see why, revisit Fragnière's attempt to show that, despite what he has said about spread and concentrated harm and moral thresholds, individual emitters are not off the moral hook. He explains that it is in part because "emitting behaviour is replicated by millions of other agents", which adds up to a great amount of harm, and "everybody can reasonably be expected to know that other people are emitting GHG, and that knowledge must be taken into account in our moral assessment" (656). These elements cause significant problems for the view, and its attempt to reduce the demandingness of our climate duties, even on the spread harm-centered analysis Fragnière favors.

As we saw Fragnière say above, "all other things being equal, causing trivial harm to a lot of people (without crossing any threshold) is less wrong than causing very serious harm to one person."⁷ However, once situated in the collective setting that Fragnière accepts we are operating in with respect to climate change, *other things are not equal*. Once we have knowledge of others' contributions and, per hypothesis, the total harm, one is not warranted in concluding that one is unlikely to cross the moral thresholds Fragnière touts as making the big difference deontically.

Given his assumptions, the probability that one crosses an important threshold for some victim, somewhere given the fact of others' contributions cannot be ignored. On the spread harm analysis – to work with the figures Fragnière draws from Nolt – I may only bear one two-billionth of the harm to any individual victim. But given that, on Fragnière's own view, I know there are manifold others each with a similar share, and I must be able to push victims over the harm-thresholds with my contributions,

⁶Fragnière could try to employ a "non-causal" interpretation of what it is for an individual's emissions to push victims over the relevant thresholds, where one is assigned a mathematical share of harm done to each victim. Sometimes Fragnière seems to talk this way, which may not be surprising given the difficulties of telling a plausible causal story and the general sense that these discussions are about "statistical victims". The problem, discussed further below, is that once you shift toward statistical/mathematical apportionment of shares of harm or responsibility, it doesn't make sense to talk about individual responsibility for *crossing (or not) thresholds* at all, and the whole view disappears. On a non-causal story, it is hard to see why we shouldn't simply (like Nolt) divide up the total harms among the contributors. And this isn't to assume that all of a contributor's harm is concentrated on one victim; it's a function of the whole story about responsibility being statistical rather than causal (where we're not talking about individual victims other than in a statistical sense, and instead in terms of aggregations and shares), which breaks down the distinction between spread/concentrated harm. One could, of course, claim that being assigned, for example, the 2,000,000,000th (and last) share rather than the 1st share of harm to victim X is what constitutes "crossing the threshold" and bearing responsibility for their death. But aside from that assignment being unacceptably arbitrary, the same probability problem I sketch below would arise. There is a dilemma for Fragnière: either go fully statistical and his criticism of Nolt doesn't hold and crossing thresholds can't be relevant for climate responsibility, or go causal in order to make thresholds relevant and my criticism below looms.

⁷It is worth noting that this comparative claim could be true without generating the conclusion Fragnière reaches with respect to climate duties. To show that, he would have to explain not just why it is *relatively* less wrong in the spread case, but that it isn't wrong enough in absolute terms to justify the demanding duty. Fragnière doesn't do this, and as evident below, doing so plausibly is a challenge.

I have to assign a non-zero probability that my contribution is the one that pushes them over the threshold. It isn't licensed to assume that my contribution would be one of the innocuous, trivial contributions before the threshold. It may well be, given the assumptions he has to work with to avoid begging the question and preserve the relevance of thresholds by not collapsing into Nolt, the contribution that triggers the threshold being crossed, even on the spread harm analysis.

To see the point schematically, imagine there are one billion equal contributors and two billion victims (where the only relevant threshold is death). Per hypothesis, we know that the amount of global harm is fixed at the deaths of two billion people. On the concentrated harm analysis, like Nolt, I would be responsible for the death of two people. On the other hand, with the spread harm analysis favored by Fragnière, as an emitter, I have one billionth of the share of the harm done to victim 1, 2, 3 . . . 2,000,000,000.

However, what this means is that, for each victim, I would have a one in one-billion probability of being the contribution that pushes that individual statistical victim over the threshold. When we calculate the overall probability, then, the probability that I would be responsible for pushing at least *someone* over the threshold is indeed quite high (about .865).⁸ So the most fair comparison is not, as Fragnière wants it to be, between causing trivial harms to a lot of people for the "spread" analysis and causing serious harm (death) to two people for the "concentrated" analysis. Instead, it is a comparison between causing a possibly trivial amount of harm to a lot of victims but with a high overall likelihood of pushing at least someone over the serious harm threshold, and the straightforward concentrated harm analysis of causing serious harm (death) to two people. This comparison significantly undermines the idea that the spread-harm analysis is a plausible way to relax the stringency of the duty.

Readers might worry that this argument problematically assumes that all of the responsibility falls on the final causal straw. That concern is legitimate and I don't mean to endorse the idea. However, it is difficult to make sense of the relevance of thresholds for responsibility ascriptions without adopting it, and thus hard to see what other choice Fragnière has. If he accepted that *responsibility for crossing thresholds* should be mathematically/statistically distributed through all the contributors, then in the context of many victims, there isn't a good reason to prevent us from "summing" together all of the discrete contributions (qua shares of each collectively caused threshold crossing) across all of the victims we have a share of. This would generate the same number as the "concentrated" model, only now by distributing "shares" of the purported distinctive *wrong* of crossing thresholds and not just of *harms*.

In order to get to the conclusion he wants, Fragnière needs some mechanism to prevent us from calculating our share of responsibility for all of the *threshold-crossing*, across all of our contributions to all of the victims. The main candidate (implausible as it is) to prevent such summing is to suggest the "*all-or-nothing*" model, where *causally* crossing the threshold is *unique* with respect to responsibility ascription (rather than getting blended out in a generalized distribution among all the contributors). But doing so runs us straight into the probability argument developed above. Fragnière

⁸To see this, we calculate one minus the probability of crossing no thresholds (imagine casting a billion-sided die, where only one face indicates crossing the threshold, two billion times, i.e., once for each victim), that is: $P(\text{AtLeastOne}) = 1 - \left(\frac{999,999,999}{1,000,000,000}\right)^{2,000,000,000}$. Similarly, the probability is .594 for pushing at least two over the threshold (one minus the probability of no thresholds minus the probability of one threshold), i.e.: $P(\text{AtLeastTwo}) = 1 - \left(\frac{999,999,999}{1,000,000,000}\right)^{2,000,000,000} - (2,000,000,000) \times \left(\frac{1}{1,000,000,000}\right) \times \left(\frac{999,999,999}{1,000,000,000}\right)^{1,999,999,999}$.

could try to adopt some weighting principle so that responsibility isn't assigned on an all-or-nothing basis, but rather as some proportion between being the "threshold crossing contribution" and just being a member of the set, but doing so would seem ad hoc in order to get the values he wants. There may be another mechanism for Fragnière to prevent summing across victims, but as soon as we accept that we have to distribute responsibility in a mathematized way among the contributors, it will be hard to motivate.

The particular problem for Fragnière is that he would need to allow for just enough summing to maintain the strength of the duty he wants (where it is "in fact fairly demanding" and *only* outweighed by core projects), but not more. If we can't sum together shares of responsibility for *threshold-crossing* across an individual's contributions to various victims, then we can't generate individual duties stronger than any single contributive share to a victim (which Fragnière and others assume is virtually innocuous, or trivial).⁹ And if we can't sum across an individual, there is even less reason we could sum contributions across different contributors. Fragnière accepts some "responsibility gap" (659) between what individuals as unstructured contributors do and what ultimately needs to be accounted for, which would need to be bridged by an account of collective duties. But this approach would leave virtually *all gap*, which he doesn't want to do. So it seems he would have to allow some "summing" of this variety. But if he did, there's no good, non-ad-hoc limiting principle. And because he is trying to *explain* demandingness, he can't use demandingness as the limiting principle (e.g., "you can sum up until it would deliver responsibility that would require sacrificing core interests"), which would be question begging.

All of which is to say that Fragnière faces significant problems if he tries to distribute responsibility for crossing thresholds among the contributors. But the other way of introducing thresholds (the "all-or-nothing" sense we started with), as we saw from the calculations above, doesn't change the expected utility/responsibility calculation of my emissions' contributions enough to do the work Fragnière wants – given that we are operating with the known assumption that others are contributing and the global harm is fixed. Even if spread around, the risk of pushing someone over a threshold that would justify a strict duty not to cross would be non-negligible.¹⁰

Admittedly, things are more complicated once we include additional thresholds aside from death and given the fact that the share of harm amounting to two billion

⁹To return to the examples of fn. 4, we could generate an individual's responsibility for one two-billionth of X's displacement or Y's starvation or Z's heat stroke, but not the combination of them.

¹⁰I have, of course, only calculated the probability of pushing at least one person over the threshold, not a full expected utility calculation. I don't want to take a stand on quantifying the magnitude of the harm, but whatever value Fragnière attaches to the death threshold, it needs to be uniform. So, whereas the concentrated case would attach that negative value to the probability of 1, the spread case would attach it to the probability of .865. My contention is that this slight decrease in probability shouldn't relax the duty in the way Fragnière suggests. Further, this is assuming that crossing the death threshold is the *only* thing that matters. If, additionally, we have to sum together the expected disutility of our non-threshold-crossing contributions, it only makes Fragnière's case harder. The "all-or-nothing" interpretation mentioned above, problematic as it is, would prevent such summing. But without a plausible way of blocking the aggregate expected utility from all of our spread-harm contributions from mattering for our responsibility, whatever extra normative "oomph" there may be from a (slightly) higher probability of crossing a wrongness-magnifying threshold, in relative terms, it does nothing to undermine the case for a demanding duty on the basis of, in absolute terms, a high aggregate expected disutility *plus* a high (even if not 100%) chance of crossing a threshold.

death-equivalent harms might be spread over more than two billion statistical victims.¹¹ The first complication makes things worse for Fragnière because adding more thresholds that can potentially generate strong duties would increase the risk that one's fractional share of harm spread across the victims crosses a morally relevant threshold that non-linearly increases the wrongness of the harm. The second complication could potentially relax the estimation and serve the function Fragnière desires if, for example, the two billion death-equivalent harms are spread over four billion rather than two billion victims (e.g., one billion deaths and three billion non-death harms that, given some conversion factor, are equivalent to one billion deaths). And so the assumptions one makes would matter.¹²

However, before investing too much in this direction to try to salvage Fragnière's conclusions, by his own lights, it is worth saying, first, that without knowing the other thresholds and their conversion factors, we won't know if this helps or hurts Fragnière's cause. If the three billion non-death harms amount to significant maiming or otherwise violating human rights, which Fragnière himself floats as thresholds that likewise generate a strong duty (but slightly less so than crossing the death threshold), then things may have gotten *worse* rather than better if trying to relax the duty. Second, the main way to significantly increase the number of spread-harm victims would be to include non-human animal victims of relevant moral status. But, assuming it is unlikely that it would take the same number of spread-harm units to amount to the death of a person than, for example, a crow, this might also just make it more likely that we cross an important death threshold of a morally considerable being which might trigger stringent duties with smaller shares of harm per death.

Generating Fragnière's desired outcome would thus rely on significant assumptions about the number of projected victims, the distribution of harms among those projected victims, as well as the number of thresholds and their deontic weight. As those assumptions get plausibly filled in, it is not promising that things would fall Fragnière's preferred way to relax our duties.

I should reiterate that I agree that there are potential problems with statistical harm-based approaches like Nolt's to individual climate morality. So I am not trying to defend them in criticizing Fragnière. But I don't think he has quite pinpointed the problems, nor has he developed a viable alternative to adequately justify a relaxation of our duties to emit less on that basis. In fact, while space constraints prohibit me from saying more, I think the very complications just sketched help reveal part of the reason why statistical harm-based approaches, whether spread or concentrated, can be problematic and depend too much on mathematized assumptions that don't always fit well with responsibility ascriptions.¹³

¹¹Note also that the unequal contribution will tend to make things worse for high emitters like those Nolt is concerned with and whose duties Fragnière's argument purports to relax.

¹²Of course, if one isn't convinced by Fragnière about thresholds or how they intersect with responsibility, and more inclined toward a mathematized model which assigns responsibility for shares of total harm in relation to shares of total contribution, then the expected utility/responsibility calculations won't differ between the two billion deaths and the two billion death-equivalent harms (spread between some deaths and some less serious harms) because the conversion factor between non-death and death harms already does all the work.

¹³John Broome (2019) and Garrett Cullity (2019) criticize some sloppy alternatives and both make strides in thinking about the causal-harm pathways of individual emissions. Broome depicts a "damage function" of our emissions that is jagged and chaotic, but which clearly has a negative expected value. This may be the

This leads to a final point. Even if Fragnière's arguments were successful in limiting the no-harm principle's reach, the actual takeaway with respect to relaxing an individual's duties is quite limited, and particularly toothless if one thinks that the main explanation of our climate duties rests elsewhere. Despite his more ambitious language otherwise (e.g., 648), it was only ever going to be a partial defense against over-demandingness targeted at harm-based accounts of our climate duties. But there are many other, and indeed I think more generally plausible and accepted, accounts of the duty to reduce emissions. Arguments about distributive justice, fair shares, and not freeriding are likely more plausible bases for individual climate duties to begin with, and they are not touched by the distinction between "spread harm" and "concentrated harm" or any threshold thesis. So even if Fragnière had defeated some harm-based demands, it wouldn't tell against demands that are differently sourced or justified, like fair burden sharing. Given all of this, Fragnière can't be seen to have granted individual agents the "measure of moral freedom in the pursuit of their most important life-projects" that he aspired to.¹⁴

References

- Broome, John.** 2012. *Climate Matters: Ethics in a Warming World* (New York, NY: W.W. Norton & Company).
- Broome, John.** 2019. Against Denialism, *The Monist*, **102**: 110–29.
- Cullity, Garrett.** 2019. Climate Harms, *The Monist*, **102**: 22–41.
- Fragnière, Augustin.** 2018. How Demanding is Our Climate Duty? An Application of the No-Harm Principle to Individual Emissions, *Environmental Values*, **27**(6): 645–63.
- Nolt, John.** 2011. How Harmful Are the Average American's Greenhouse Gas Emissions? *Ethics, Policy and Environment*, **14**(1): 3–10.
- Scheffler, Samuel.** 1982. *The Rejection of Consequentialism* (New York, NY: Oxford University Press).
- Singer, Peter.** 1972. Famine, Affluence and Morality, *Philosophy & Public Affairs*, **1**(2): 229–43.
- Williams, Bernard.** 1973. A Critique of Utilitarianism, in *Utilitarianism: For and Against*, ed. by J. J. C. Smart and B. Williams (Cambridge: Cambridge University Press), pp. 75–150.

right way to think about causing climate harm, but regardless it is very plausibly not *all* we need to know for individuals' climate responsibility ascriptions.

¹⁴Thanks to the reviewers, Travis Rieder, and the Fair Limits team for their helpful feedback.

Cite this article: Hickey C (2021). The Demandingness of Individual Climate Duties: A Reply to Fragnière. *Utilitas* **33**, 353–360. <https://doi.org/10.1017/S0953820820000461>