

**Climate Change and Individual Obligations:
A Dilemma for the Expected Utility Approach, and The Need for an Imperfect View**

Julia Nefsky

1. Introduction

Unless global greenhouse gas emissions are rapidly reduced, the consequences of climate change will be catastrophic.¹ What is the significance of this for how ordinary people ought to live their lives?

Of course, government action is crucial for tackling the climate crisis. But it is also true that if enough individual people reduce their emissions, this would reduce global warming and its harmful consequences. We thus frequently hear calls to change our individual behaviour: take a bike or public transit instead of a car, forego air travel, choose vegan options over meat and other animal products, avoid single-use plastics, buy energy-efficient appliances, etc. The suggestion is that we ought to make these sorts of choices – we act wrongly if we do not.

Something like this surely must be right. Given the facts about climate change, surely we, as individuals, have moral obligations when it comes to our emissions-producing lifestyle choices. But there is a problem in understanding how there could be such obligations: while we together could reduce climate change harms if enough of us make low-emissions lifestyle choices, any one ordinary individual's actions typically do not seem to make a difference. It seems as though the consequences of climate change are not going to be different, for instance, depending on whether or not I drive to work, or whether or not I fly to Europe this summer. But

¹ We are already seeing the effects of climate change (extreme weather, droughts, wildfires etc.) For projections about what is likely to happen in the future see, for instance, the IPCC's Special Report on Global Warming of 1.5°, <https://www.ipcc.ch/sr15/>

if climate change harms will be just as bad give or take my individual emissions-producing activities (driving, flying, etc.), then how can I be said to be acting wrongly, on climate change grounds, in doing such things? This is the ‘inefficacy challenge’.²

The inefficacy challenge actually runs deeper than just described. If it really won’t make a difference whether I drive to work or take the subway, not only is it unclear how I could be acting wrongly by driving, it is unclear that there is any point at all (to do with climate change) in refraining. What reason would I have to inconvenience myself in this way, if doing so won’t make a difference for the better? Elsewhere I have focused on this more basic challenge at the level of reasons,³ but my concern in the present paper is with our moral obligations.

The aim of this paper is not to answer the inefficacy challenge for individual climate-change-related obligations, but rather to say something about the *sort* of answer we should be looking for. And I do this by way of raising a new problem for a popular approach to the inefficacy challenge: the expected utility approach. The first half of the paper focuses on the expected utility approach. I argue that it faces a problematic dilemma. In the second half of the paper I diagnose the problem. There we see that the dilemma arises from a very general feature of the view, and thus is shared by various other views as well. I then discuss what a solution to the inefficacy challenge needs to look like if it is to avoid the dilemma.

The inefficacy challenge is not specific to climate change. It arises in any ‘collective impact’ context: any context in which people can collectively cause, or fail to prevent, great

² A different kind of climate change-related obligation that individuals might have is to pressure governments and other organizations to change their policies. It is important to note, though, that the inefficacy challenge can arise with respect to activist activities as well. While very many signatures on a petition or participants in a protest could make a difference to government decisions, it’s hard to believe that these decisions will go differently give or take any one ordinary individual’s signature or attendance at the protest. The discussion in this paper applies to that context as well, but for simplicity I’m focusing my attention on obligations with respect to lifestyle choices.

³ See Nefsky 2017.

harm or injustice by acting in a certain way but in which any one such act does not seem to make a difference. While the central example of the paper is climate change, in the final section I discuss how the point extends to other collective impact cases.

2. A Dilemma for the Expected Utility Approach

Importantly, I am approaching the inefficacy challenge without assuming any prior commitment to a general moral theory like consequentialism. That is, I am approaching the problem directly, rather than from the perspective of someone committed to a particular general theory. What I am interested in is how we can give an intuitively plausible answer to the challenge – one that delivers certain intuitively important results. As an illustration, consider Ms. Lavish:

Ms. Lavish lives a high-emissions lifestyle. She has a big home, which she keeps very warm in the winter and very cool in the summer, and which is decked out with lots of fancy energy-consuming features, like an indoor swimming pool and gas heat lamps for her rooftop deck. She travels frequently for weekend getaways and vacations. She has a gas-guzzling SUV, which she loves driving. She uses it multiple times a day – to get to the grocery store (even though it is just down the street), to get to work (even though she could easily take public transit), and so on. Indeed she loves driving her SUV so much that on most Sundays she goes for a drive along the nearby scenic highway just for fun.⁴

I think a skeptical answer to the inefficacy challenge (i.e. one that says that individuals do not have climate-change-related lifestyle obligations) is really implausible when one thinks of someone like Ms. Lavish. Surely Lavish should not be living like this, given the facts about climate change. I am interested in how we can give a plausible solution to the inefficacy challenge that delivers this intuitively important result, among others.

To be clear: there might be other reasons, unrelated to climate change why Lavish should not be living like this. But for our purposes, we want to leave any such other factors aside. The

⁴ The example of going for a Sunday joy ride in one's gas-guzzling SUV comes from Sinnott-Armstrong 2005.

question is: how can we explain why her behaviour is morally problematic on climate-change-related grounds?

One popular way of answering this is by appealing to expected utility. This expected utility approach ('EU') proceeds in two steps⁵: first, it says that it is incorrect to say that your individual act won't make a difference, because there is always at least some chance that it will. In the case of climate change, this step goes something like this: as emissions climb higher, there are going to be various points – or, thresholds – at which higher greenhouse gas concentrations in the atmosphere result in further harmful outcomes (like floods, droughts, hurricanes, and so on, that would not otherwise have occurred, or would not otherwise have been as severe). While it may be very unlikely that your individual emissions-producing act (e.g. your drive to work) is the one that causes such an emissions threshold to be crossed, there is at least a chance that it does. There is a chance, in other words, that your small individual act pushes things over a tipping point, resulting in further harm that otherwise would not have occurred. So there is at least a small chance that you do make a difference for the worse by driving.

The second step is that even if the chance of making a difference is only tiny, this will be a tiny chance at making a large difference for the worse. We are talking about a difference to whether or not severe storms, floods, droughts and so on, occur, or to how severe they are. These things affect large numbers of people and animals. A small chance at making a large difference in harm amounts to a substantial expected harm. As a result, we can say, on expected utility grounds, that it is typically impermissible to do a higher emissions act when there are lower

⁵ Many philosophers have adopted EU as an approach to collective impact cases in general. For two excellent general defenses of this approach, see Kagan 2011 and Norcross ms. For the EU argument applied specifically to the climate change case, see Morgan-Knapp and Goodman (2015), Lawford-Smith (2016), and Broome (2019). Broome (2012, 76-77) also partially appeals to it, and Hiller (2011) does too.

emissions options available: it is impermissible *unless* there is some sufficiently weighty expected benefit to doing it over the alternatives.

Applied to Lavish's case, EU says: take any instance of the activities described. For instance, take one of her Sunday drives along the scenic highway in her SUV. This has a substantial expected harm, because there is a chance (even if only very small) that the emissions produced will push things over some threshold, resulting in harmful outcomes that would not otherwise have occurred, or would not have been as severe. This expected harm is going to clearly outweigh the expected benefit of this drive. After all, the only expected benefit is pleasure for her, and lower-emissions options could still be quite pleasurable, even if not as much. So, Lavish acts wrongly in driving her SUV along the scenic highway for fun that day. And, analogously, every other time she does it.

The exact same point applies to many of Lavish's other activities. She is constantly making choices like this: driving when she could easily avoid it, cranking up the air conditioner, heating her indoor swimming pool, flying to yet another holiday destination, and so on, all for the sake of gains in her own pleasure, comfort or convenience – gains that are minor when compared to the substantial expected disutility of each of these acts. Thus, the view says: Lavish is acting wrongly much of the time, on climate-change-related grounds.

So far so good. But I am going to argue that EU faces an unhappy dilemma. There are two possibilities for the view: either it cannot really get those non-skeptical verdicts *or* it yields implausibly strong verdicts. Either way, it does not deliver intuitively important results that we should hope our solution will deliver. So we should be interested in finding a different approach that can do better.

Beginning with the first horn of the dilemma: some have argued that the expected utility calculus, done correctly, will not actually give the desired non-skeptical verdicts. For instance, Budolfson points out that even if one's car trip is in fact the one that causes some emissions threshold to be crossed, you only make a difference for the worse if that threshold would not have been crossed, or would not have been crossed at approximately that time, had you not taken that car trip. But that, he argues, is essentially impossible. This is because there is "an entire planet of other people" constantly engaging in emissions producing activities. So, if some emissions threshold is crossed at t_1 with your car trip, it would still have been crossed, if not at t_1 then at an imperceptible amount of time later without it.⁶ But this means that there is essentially no chance that any of Lavish's acts makes a difference to climate change harms. So, since Lavish does at least gain something from these activities, the expected utility of doing them will actually be higher than that of refraining. So, she acts permissibly as far as expected utility is concerned. Kingston and Sinnott-Armstrong make a similar argument.⁷

If this argument is right, then, the expected utility calculus – done correctly – actually yields the skeptical conclusion in Lavish's case, and to the question more generally.

One reply, which Broome emphasizes, is that it is not true that an imperceptible fraction of a second means making no difference in harm. This is because the atmosphere is a chaotic system. Because of this chaos or instability, a miniscule fraction of a second difference in the timing of emissions *can* make a difference to what happens. For instance, it could make a difference to the location or severity of a storm. This is also known as the butterfly effect: the

⁶ Budolfson ms, p. 36.

⁷ Kingston and Sinnott-Armstrong 2018, pp. 180-181.

idea that a small disturbance in the atmosphere can have random, unpredictable influences on weather.⁸

I don't think this reply works. The chance of making a difference via the butterfly effect is no more of a chance of making a difference for the worse than it is a chance of making a difference for the better. So, simply delaying a threshold-crossing by an imperceptible fraction of a second is no more likely to make the harmful outcomes more severe than it is to make them less severe. In other words, the chances of individually making a difference for the better and for the worse *via the butterfly effect* should really cancel each other out in an expected utility calculation.⁹

But a better reply comes from the fact that it is possible that there will be a peak point for global emissions after which emissions levels start declining. Perhaps world leaders will make the necessary changes in policy for this to happen. If the timing of the peak happens to coincide exactly with an emissions threshold being crossed, then each individual drive will count as having made a difference for the worse. This is because it's not true in this scenario that had your drive not happened, the threshold would still have been reached at basically the same time anyway. This peak-emissions threshold-crossing coincidence scenario seems highly unlikely.

⁸ Broome 2019, pp. 112 -113.

⁹ Thank you to Mark Budolfson for a discussion of this by email. Morgan-Knapp and Goodman consider this sort of reply. They point out that a walk in the park could also affect the weather via the butterfly effect. But they say that in that case the probabilities cancel out because "you have no information to suggest that taking a walk in the park is any more likely to cause a hurricane than it is to prevent one." (p. 185) But they claim this is not so with a drive: "the carbon dioxide emitted by the car engine does affect the process by which the sun heats the earth, and does change the level of energy in the atmosphere. Here the distribution of effects is not symmetrical. It is slightly skewed towards worse outcomes." (p. 185) Kian Mintz-Woo has made a similar point to me. However, I think our discussion above suggests that this way of thinking about the butterfly effect and climate change is not correct. The difference one's drive could make with respect to the concentration of greenhouse gases in the atmosphere is, if anything at all, purely a difference of a tiny fraction of a second in timing. (This leaves aside one exception, which I discuss next.) I don't think there is any reason to think that that sort of extremely miniscule difference in timing would be more likely to make a difference for the worse than for the better in the consequences of climate change. In other words, if it makes a difference to harmful outcomes at all, this would be *purely* because of chaos/the butterfly effect, which is no more skewed one way than another.

But, assuming it is indeed possible, this mean that there is a tiny chance of making a difference for the worse in climate change harms via a single drive.¹⁰

Going forward, then, I am going to proceed on the assumption that what we know is that there *is* a chance of making a difference to climate change-related harms by a single drive, or by any other ordinary emissions-producing activity, but that this chance is extremely small.¹¹

¹⁰ I have focused on what I think is the best argument found within articles like Broome (2019), Morgan-Knapp and Goodman (2015), and Lawford-Smith (2016), for the claim that there is a significant chance of making a difference to climate change harms via a single ordinary activity like driving. But these authors also bring in other points to try to defend the claim. In this note, I briefly address one other such point from Broome (2019), namely that some of the harms of climate change happen through “continuous processes”, as opposed to discrete events. An example he gives of a discrete climate event causing harm is a typhoon. Examples he gives of continuous processes causing harm are: “as water tables sink, people have to do more and more exhausting work to get the water they need; as the sea level rises, arable land is gradually washed away and families grow more hungry.” (p. 112) Broome suggests that in the case of these sorts of continuous processes, it isn’t that each emissions-producing act has a chance of crossing a threshold for a big change in harm, but rather that each act does make a difference for the worse but just a tiny, imperceptible one. Because the tiny difference affects a lot of people and animals, though, we should count it as a non-trivial difference. (p. 122) But I do not think this argument works. Just because the process that causes harm is continuous rather than a discrete event, this does not show that each individual drive makes a difference to how the process unfolds. The continuous process could be causally overdetermined. Also, even if we did concede that each individual act makes some imperceptible difference to the process, I do not think this establishes that each act makes a difference with respect to what matters morally. Maybe, for instance, each additional ordinary act makes the sea a trillionth of a meter higher, but this does not make a difference to the harmful effects of the rising sea level. (See Nefsky 2011, pp. 372-377 for a detailed explanation of why an imperceptible difference to something that causes harm does not imply a morally relevant difference in harm). In any case, while Broome does spend time on this imperceptible harms argument, he also seems to acknowledge that his stronger argument is the expected utility argument appealing to chances of making larger differences. (See p. 125)

¹¹ A different way of getting a chance of making a difference might be through the idea of ‘signaling’: by driving less or not at all, by avoiding air travel, by reducing one’s energy use at home, and so on, one sends a political signal – one communicates to the government, and to other people, that one thinks climate action is important, and that one is eager to cooperate in a solution. (Cullity, 2015, p. 12) While the relevant people will not necessarily pick up on these signals, they might. This could make a substantial difference for the better. But is this really a significant possibility? The chance of my making a difference to government policy by way of a signal I send when I, e.g., take the subway or forgo my vacation seems extremely small or non-existent. As Cullity writes, “it is fanciful to think that simply reducing energy consumption itself sends an effective political signal. Joining a political lobbying movement can be a way of participating in sending such signals; simply turning the lights off is not.” (2015, p. 13) On the other hand, it does seem that there is a chance of influencing some other ‘ordinary’ people to change their habits by way of changing our own habits. If a friend comes over, and sees various eco-friendly products in my house, this might influence her to switch to similar products. But for many of us, the chance of having a large enough influence in this way to actually make a difference to climate change harms seems tiny. (See Nefsky 2019 section 4.3 and also Nefsky 2018 pp. 270- 272 for more on

This vindicates the first step of the EU argument. But it seems questionable now that the second step would go through. It could very well be that, because the chance of making a difference for the worse is so small, the expected harm is not actually going to be substantial enough to outweigh the expected benefits to Lavish of her scenic drive, or her weekend getaway in Napa. If so, then the expected utility calculus again really gives the skeptical verdicts – it would say that Lavish acts completely fine, as far as climate change is concerned.¹²

That, then, is the first horn of the dilemma: perhaps EU really just gives us the skeptical verdicts. But now let's assume that the expected utility calculus does work out as described.¹³ So, let's assume that this tiny chance does yield a significant enough expected harm to outweigh the benefits for Lavish of all or most of her luxurious emissions-producing acts in an expected utility calculation. The problem now is – and this is the second horn of the dilemma – that EU cannot get that result without also getting other implausibly strong verdicts.

Consider Ms. Green:

Ms. Green has a very low-emissions lifestyle. She lives in a tiny apartment, with minimal and efficient appliances. She does not use an air conditioner in the summer even though it gets quite hot where she lives, and she delays turning on her heat in the Fall until it is necessary. She gets around almost exclusively by walking, biking and taking public transit. She does own a car – an old Toyota Camry that her grandma passed down to her. But she barely drives it. She keeps it primarily so that she can drive every once and a while to the woods (about an hour away) to go hiking. She barely travels. Her last trip was two years ago, when she went to Italy with a friend.

this point.) So I think the conclusion remains: we can generally assume that if there is a chance of individually making a difference via an ordinary emissions-producing activity, it is a very small one.

¹² Sinnott-Armstrong and Kingston (2018, p. 180) and Cullity (2015, pp. 7-8) also make this point.

¹³ A different issue with the expected utility approach is that we do not and perhaps cannot know which side of these two possibilities it is: we do not, and perhaps cannot, know if the expected utility calculus comes out against or in favour of the drive. Even if we can know that there is a tiny chance at making a difference, and so some expectation of harm, it is arguable we cannot know how to quantify it and weigh it up against the expected benefits. See Cullity 2015, section III for an argument like this. See also Heal and Millner 2014. They argue that given the depth and nature of our uncertainty about various aspects of the climate change problem, the expected utility framework for decision making under uncertainty is *not* a good one for use in climate change policy.

The problem for EU is that if it is going to come out against Lavish's driving her SUV for fun along the scenic highway, it also must come out against Green's occasional drive to the woods to go hiking. And in general, various seemingly permissible actions of Green's would come out to be wrong on this model.

Take one of Green's rare drives to the woods to go hiking, and compare it to one of Lavish's many Sunday drives along the scenic highway. Things could easily be such that the numbers that go into the two expected utility calculations are the same, and thus that the two trips have the same expected utility. In particular, we can stipulate that the two trips produce the same amount of total emissions. Even if Lavish's SUV produces more emissions per mile, we can just stipulate that Green's round-trip to and from the woods is longer. Thus, there would not be any difference in the probability of making a difference for the worse, or in the magnitude of that supposed potential difference. So, the two trips would have the same expected harm. We can also stipulate that the expected benefit to each of going on their respective trips would be the same: the two trips have an equally high probability of giving their respective agents an equally fair amount of pleasure. But if these factors are the same, then the two acts have the same expected utility.

Of course the verdicts depend not only on the expected utility of the acts in question but also on the expected utility of the alternatives available to the agents. But there need not be any difference there either.

So, the point is: if we say that the calculus comes out against Lavish driving her SUV for fun, we also need to say that it comes out against Green's driving her car to the woods. So, on EU, if the former is wrong, then so is the latter.

The same point can be made with respect to other things that Green does. Take her flying to Italy two years ago. The fact that she only travels once every few years does not enter into the expected utility calculation for that individual trip.¹⁴ So, there is no reason that taking this trip couldn't get the same values in the expected utility calculus as one of Lavish's many trips. So, if Lavish acts wrongly in travelling yet again on the expected utility view, then Green acts wrongly too in traveling just that one time.

This, intuitively, is not the right conclusion. Given that Green keeps her emissions so very low, and given that any one drive or flight is highly unlikely to make a difference for the worse (via climate change), it seems implausible that Green is violating her climate change related obligations by taking that one drive or that one rare vacation.

Now, if you are starting from a point of having already subscribed, on independent grounds, to an expected utility view about what individuals ought to do in general, this probably is not going to be a convincing argument against your view. You might be happy to just follow the EU approach wherever it leads. But my starting point is different: it is one of looking for an intuitively plausible answer to the inefficacy challenge. And the point is that, if this is what you are interested in, EU does not do well – it is either too weak, yielding the skeptical conclusion, or it is implausibly strong.

But let's consider some replies.

First, advocates of EU might reply: 'it is true that on our view, Green's drive comes out wrong, just like Lavish's. But it is not as though the two people are therefore moral equals. Green does way fewer of these wrongful acts. Whereas, Lavish is acting wrongly on emissions-related grounds constantly. So, Lavish is deserving of much more criticism than Green.'

¹⁴ Other than if the rarity itself has a relevant effect on how much pleasure she gets from the trip, which it need not. More on this point shortly.

I do not think we should be satisfied with this reply. The claim that Green acts wrongly in occasionally driving to the woods for a hike, or in flying to Italy that one summer, is still implausible. This is not to say that there is no climate change-related reason for Green not to do such things.¹⁵ But to say that these things are impermissible for her is intuitively false. She doesn't seem morally criticizable, or to be going at all morally astray, in making those individual choices.¹⁶

But perhaps the reply is that this intuition is mistaken, and we should *not* regard the conclusion that Green acts wrongly on those occasions as implausibly strong. One might say: 'yes, it is demanding, but morality should be very demanding when it comes to stopping severe climate change harms.'

But this reply won't do. The claim is not that morality should not demand a lot of us when it comes to stopping climate change harms. If it turned out that each time one drives, or flies, one has a fairly high chance at making a substantial difference in harm, via climate change, then it would not be implausible – at least not in the sense that I'm talking about – for morality to demand that one not drive a car just for the sake of the pleasure of going hiking, or for it to demand that one never travel for a vacation. My point is, rather, this: *given* the assumption that

¹⁵ What reason there is depends on one's answer to the inefficacy challenge at the level of reasons. See Nefsky 2017 for my answer (which can be combined with other answers as well).

¹⁶ To this they might reply that we need to separate whether Green is criticisable from whether she is acting wrongly. It could be that because she is so good overall, and far better than the average person, it would be inappropriate to criticize her, but that she still acts wrongly. I agree that there are situations in which the fact that someone is generally very good, or much better than average, can make criticizing a relatively small moral failing of theirs inappropriate, even though it is still a failing. For instance, suppose Nessa is normally an exemplary mother, who is extremely patient, kind, and calm, but on one occasion she snaps at her kids – getting overly angry in response to something minor. One might feel as though one cannot criticize her for snapping that one time; it would not be reasonable to criticize her, given how calm and patient she generally is and given that this is a relatively small infraction. And yet, even though criticizing her may be inappropriate, it was still wrong for her to snap on that occasion. But while I recognize that there are cases like that, Green's is not one of them: it is not that while Green is acting wrongly in driving to the woods to go hiking we should not criticize her for it; it is that she isn't acting wrongly in doing so. This rare drive of hers to the woods does not seem like an example of impermissible conduct.

her drive, or her taking that flight to Italy, has no more than a tiny chance at making a difference when it comes to climate-change-related harms, and *given* that Green keeps her emissions so low overall, it seems unfittingly demanding for morality to forbid Green from ever doing those things.¹⁷

Perhaps a better way to reply, then, is to dispute my claims about the expected utilities. One might argue that the expected benefit for Green of any one of her trips to the woods must be much higher than the expected benefit to Lavish of any one of her scenic SUV drives, because Lavish is doing things like this all the time. Similarly, because it is rare for Green to travel, her Italy trip must have been so much more meaningful and pleasurable to her than any one of Lavish's trips would have been to her. So my stipulation that they get the same expected benefit from their individual acts should be rejected. And if the expected benefit for Green is high enough, this could outweigh the expected harm, making these activities permissible for Green.

But the first issue with this reply is that it assumes that scenic drives, hikes in the country, and world travels always have a decreasing marginal utility, and that is not true. For these specific types of activities, whether or not they have a decreasing marginal utility within the particular sorts of choice situations under discussion will depend on the details. For Lavish, it need not be that each of her weekly drives is less pleasurable than it would be if she did it less frequently. It could easily be the opposite: because she does it weekly, the route becomes very

¹⁷ A related point is that one cannot get out of the problem by opting for a less demanding version of EU. The dilemma arises just as much for a satisficing version, or a hybrid version that incorporates agent-centred prerogatives. On a satisficing version, the problem will be that wherever you set the bar for impermissibility, if some of Lavish's individual car trips (or travels) are going to fall below the bar, then so will Green's. So, for a satisficing view to say that Lavish acts wrongly sometimes (which it must if it is to give a non-skeptical answer to the inefficacy problem), it must also say that Green acts wrongly with her occasional drive or rare flight. Similarly, for a hybrid view, which allows an agent to give a certain amount of additional weight to their own projects and commitments: if Lavish's individual car trips are going to fall below the bar, such that they are impermissible *despite* the extra weight the agent-centred prerogative assigns to the benefits she gets from them, then so will Green's.

familiar and thus the drive becomes more stress-free and purely pleasurable. We can also imagine that the ‘ritual’ element of it adds a significant amount of pleasure. And for Green, just because her car trips to the woods are rare, we cannot assume that they have some particularly elevated importance for her happiness. Her hikes can just be one among many things she enjoys doing. So, the stipulation that Lavish and Green in fact get the same amount of expected benefit from their respective car trips on a given occasion is entirely fine.

Now EU can at least say that *were* the example different, such that there was some big life-changing benefit to Green of going to this hike, then it could be permissible for her to do so, even though Lavish’s SUV drive is impermissible. But the second, more important issue is this: the thought that it is only permissible for Green to drive to the woods on a rare occasion *if* she will get much more pleasure or fulfillment from doing so than Lavish gets from any one of her many trips in her SUV, and that otherwise it is impermissible for her to do so, is itself highly unintuitive. Even if nothing major is at stake for Green in whether or not she drives to the woods on a given weekend, it should be permissible for her to do so, *given* that she lives such a low-emissions lifestyle and *given* that any one such trip is extremely unlikely to make a difference for the worse. We can fill out the scenario to make the stakes for Green obviously low: imagine that there are good hikes close to home that Green can access without driving but that she likes to drive out to this further location occasionally just for a change in scenery. Or even, imagine that she could get to this particular hiking trail without using her car – she could do so by taking a two-hour trip on a public bus. So, it is just a matter of convenience that she drives there. Even so, it seems unfitting for morality to forbid Green from enjoying this convenience on rare occasions, given that her overall emissions are so low.¹⁸

¹⁸ One might propose that EU can avoid these troubles by focusing on overall patterns of activity or lifestyle choices rather than on individual actions. If you look at the expected utility of Lavish’s and

3. Diagnosis

EU, thus, faces a dilemma: either it is stuck with the implausible skeptical verdict that individuals do not have climate-change-related obligations with respect to their lifestyle choices, *or* it must commit to the implausibly strong verdict that even someone leading a very low-emissions lifestyle cannot permissibly do things like drive or fly on rare occasions, if it is only for pleasure or convenience. What is it about EU that yields this result?

It is, in fact, something very general about the approach – something it has in common with many other views. And this is that it approaches the question of our individual obligations in collective impact cases by trying to identify a consideration that, in effect, makes acting in the relevant ways *pro tanto* wrong.

By ‘the relevant ways’ I mean the ways that when enough of us so act, serious harm or injustice results; that is, the ways of acting in a collective impact case that generate the inefficacy challenge. Let’s call these the ‘collectively harmful’ ways of acting. By ‘*pro tanto* wrong’, I mean wrong on any given occasion unless there is a sufficiently weighty competing consideration speaking in favour of so acting on that occasion. Advocates of EU might not be

Green’s lifestyles as a whole, instead of focusing on each act separately, we can get that Lavish acts impermissibly while Green acts permissibly. This is because – the argument would go – while each act of Lavish’s has only a tiny chance at making any difference for the worse, her overall pattern of activity has a substantial chance because the risk is repeated so many times. Green, on the other hand, only takes the tiny risk infrequently, and so her overall pattern does not have a high expected disutility. Repeatedly taking a tiny risk of harm adds up to a substantial risk, whereas taking it only a few times does not. (See Spiekermann 2014, pp. 76-77 for a similar idea.) However, while there might be something to this idea (see note 29 for more on this proposal), this is not the EU proposal under discussion. The authors I have in mind apply EU to individual acts, and not just to overall lifestyle choices or sets of acts. They are, as far as I can tell, working with a standard conception of EU on which it can be applied to any choice situation an agent faces, from individual acts to lifestyle choices. On this conception, there needs to be consistency between what EU says about an act viewed as part of a set and what it says about that same act viewed individually. (For a formal statement of this condition, see “Separability”, McClennan 1990, p. 122). So, the EU proposal in question cannot condemn Lavish’s pattern of driving every Sunday, without saying of at least some of those individual drives that they come out as wrong on the expected utility calculus. Since any one of Lavish’s individual drives can have the same utility/disutility specifications as one of Green’s, it cannot condemn Lavish’s lifestyle or pattern over time without also condemning Green’s individual drives.

inclined to use the term ‘pro tanto wrong’ in describing their view, but it is a good way of capturing the very general feature of the view that I have in mind. In the case of climate change: EU claims that any ordinary emissions-producing activity has significant expected harm, and it says that therefore any such activity will be wrong, unless there is a sufficiently weighty expected benefit to doing it. This has the effect of saying that an emissions-producing activity is pro tanto wrong.

Many responses to the inefficacy challenge have this same general nature, and as a result they face the exact same problem: if they can get non-skeptical verdicts at all, it is at the cost of having to say, implausibly, that even someone leading a very low-emissions lifestyle cannot occasionally drive a car for the sake of pleasure or convenience, or cannot even take the rare vacation.

Take, for example, responses that appeal to notions of complicity or participation. Such views approach collective impact cases by saying that even if it makes no difference for the worse, your act can be wrong because it amounts to participation in a system, or way of life, or set of acts – the details vary – that causes great harm or injustice. By doing a collectively harmful act (e.g. driving, or boarding a plane), you are taking part in a system that causes these bad consequences. You are thus complicit, even if you do not yourself make a difference.¹⁹ Presumably the claim is that complicity, in the relevant sense, is pro tanto wrong, rather than always all-things-considered wrong. Otherwise, these views wouldn’t be able to say that it is permissible to drive someone to the hospital to save them from a life-threatening emergency. So, these views have that same general nature that I attributed to EU: they say that an emissions-

¹⁹ See, for example, Kutz 2000, chapter 6, Parfit 1984 chapter 3, and Driver 2015. The point I am about to make also goes for a view that says that emissions-producing activities involve complicity in the sense of cooperating with someone else’s wrongful plan – for instance, cooperating with oil companies or the airline industry. For this understanding of complicity, see McPherson 2016, section 3.

producing act is wrong on any given occasion, unless there is some sufficiently weighty consideration in favour of doing it on that occasion.

They, thus, face the same problem: by stipulation, there are no weighty considerations speaking in favour of Green's driving to the woods on the given occasion, or at least, nothing that is not also there for Lavish's driving her SUV along the scenic highway and so nothing that this view should count as 'sufficiently weighty' if it wants non-skeptical verdicts. So, if it is going to get non-skeptical verdicts, this view must say that Green acts wrongly. This is the exact same result as with EU.

4. The Imperfect Approach

What is the alternative to having a view of that nature? It might seem as though any non-skeptical answer would have to be like this: if it is going to get obligations not to act in the collectively harmful ways, these would have to take the form of obligations not to do so unless there is a sufficiently strong countervailing consideration in favour of doing so. What else would a non-skeptical answer to the question look like?

But there is an alternative, and one that I think is much more plausible, at least in the case of climate change. The alternative is to argue for what I am going to call an 'imperfect obligation': an obligation to avoid acting in the collectively harmful ways – the ways that when enough of us so act, great harm or injustice results – *enough* of the time. So, rather than arguing that there is a consideration that makes acting in such ways *pro tanto* wrong (that is, wrong in any instance unless there is, in that instance, sufficiently strong countervailing considerations), this alternative type of approach is to argue that there is a consideration that makes acting in that way *too much of the time* wrong. On such an approach it can be permissible to act in a collectively harmful way even when there is no weighty moral consideration in favour of doing

so. But one must not do so too much. What counts as ‘too much’ would depend on the particular account; I am just describing a general nature that many different accounts could have.²⁰

I am using the term ‘imperfect obligation’, but nothing hangs on this choice of terminology, and I am not claiming that my usage matches how others use the term. There is a relationship between my sense of imperfect obligations for collective impact contexts and what most others have in mind when they talk about imperfect duties. In particular, if we have an obligation in a collective impact case that is imperfect in the Kantian sense, this will also be imperfect in my sense. It will have the feature that I am talking about.²¹ But there can be views that are imperfect in my sense, without being imperfect in the Kantian sense. In any case, though, this is not important for my present purposes.

When it comes to climate change, an imperfect view – in my sense of the term – would say that what one is obligated to do is *keep one’s emissions low enough*. It is an obligation, in other words, to reduce your emissions (if they aren’t already low enough). This sort of obligation concerns your lifestyle choices taken cross-temporally. You must make these choices so that you do not emit too much over time.

On such a view, the wrongness of Lavish’s activities is found in what she does over time, and not in what she does on any one occasion taken in isolation. So, her driving her gas-guzzling

²⁰ Whether emitting ‘too much of the time’ should be understood as a pro tanto wrong or an all-things-considered wrong will depend on the details of the account. There are ways of understanding ‘too much’ such that overriding considerations are built in, and ways of understanding it such that they aren’t. Especially on the former type of account, the content of what counts as ‘too much’ will need to vary depending on the individual. e.g. it might be different for someone poor than for someone wealthy, or for someone who is an ambulance driver than for someone who is a teacher at a school reachable by public transit.

²¹ At least, on most interpretations of a Kantian ‘imperfect’ duty. The same goes for various other accounts of imperfect duties that are not specifically Kantian. For instance, see Hanser 2014. On Hanser’s account, “an imperfect duty is a second order reason, with requiring force, to act in accordance with first-order justifying reasons of a certain kind often enough over time.” (p. 311) A duty that qualifies as imperfect on this account will also be imperfect in my simpler, less robust sense. For another example, see Schroeder 2013.

SUV along the scenic highway a given Sunday need not be wrong. What is wrong is that she does this sort of thing regularly.²² And on such a view, Green can permissibly drive to the woods on a given occasion, even if there is no weighty reason to do so, because she can do this without violating her obligation to keep her emissions low enough.

Of course, what counts as ‘low enough’ will have to be filled out. But again, this is the job of the particular account of our obligations. I’m just describing a general nature that many different accounts could have.²³

The possibility of an imperfect view has often been overlooked in discussions of collective impact cases. Authors often see the task as one of looking for a consideration that makes an instance of the relevant act type wrong (in the absence of strong competing considerations), and resultingly they think that if they can’t find such a consideration then we are left with the skeptical conclusion.²⁴ But there is at least one important view in the literature that

²² Her individual drive might be wrong, but one would not see this if one looked at it in isolation. If it is wrong, on an imperfect view, this would be derivative on its place in the overall pattern (e.g. it is past the point when one has already emitted too much overall.) Alternatively, though, there can be imperfect views on which *no* individual act is wrong, even though the overall pattern over time is wrong. This would be the moral equivalent of Tenenbaum and Raffman’s claim that there can be ‘top-down irrationality’: certain patterns of activity can be “irrational without any of the momentary actions that compose those patterns being irrational”. (2012, p. 22) One could have an imperfect view in which there is top-down wrongness: a certain pattern of activity (e.g. Lavish’s) can be wrong without any of the momentary actions that compose it being wrong. I give an example of such a view in note 29 below.

²³ One might wonder whether the notion of an imperfect obligation that I am invoking really marks off something distinctive. Is a so-called ‘imperfect obligation’ not to act in a collectively harmful way really just a *perfect* obligation not to act in that way too much? I am not going to delve into the details of whether or not this recasting works, because even if it does work, this is fine for me. It is not important for the point I am making that there is a deep, irreducible distinction between these two types of obligations. What is important is rather that there are two different ways a non-skeptical answer to the inefficacy challenge could go: (1) it could say that acting in collectively harmful ways is *pro tanto* wrong, or (2) it could say that while there are restrictions for individuals on acting in such ways, such acts are not *pro tanto* wrong – it can be permissible to do such things, even when there is no weighty moral consideration in favour of doing them. These are two different ways an answer to the problem could go, even if there is no deep category difference in the kind of obligation invoked.

²⁴ This is implicit in the possibilities that many authors consider, and the way they talk about them. But for a more explicit example, see Sinnott-Armstrong 2005. There Sinnott-Armstrong says that to investigate whether we have global-warming-related obligations as individuals with respect to our

does yield imperfect obligations, in my sense of the term. This is the view that our individual obligations in these cases take the form of obligations to do one's 'fair share' of a collective obligation. Whether or not one does one's fair share of our collective obligation to reduce harmful climate change is a matter of whether one keeps one's overall emissions low enough, rather than being about what one does on any given occasion taken in isolation.²⁵ One can drive for fun along a scenic highway, or to the woods to go hiking, without violating this obligation, and this is so even if there is no weighty consideration in favour of such a drive.²⁶

I give the fair share view, though, just as an example of an existing view that has the general nature that I am talking about.²⁷ I am not specifically advocating for it, and indeed, I

lifestyle choices we can focus on a test case: is it wrong to drive a gas-guzzling SUV just for fun on a particular sunny Sunday afternoon? He looks at many different attempts to explain why this drive would be wrong, and argues against each one, concluding that it is not wrong. He draws from this the skeptical conclusion: individuals do not have obligations with regard to their lifestyle choices and emissions. It is governments not individuals who have an obligation to do something about climate change. (pp. 311-312) But once one recognizes the possibility of an imperfect view, one sees that this skeptical conclusion about individual obligations does not follow from the arguments about the test case. Even if it is true that an individual joy ride in a SUV is not wrong, this does not mean that individuals are not obligated to reduce their emissions. An imperfect obligation could even require drastic reduction in emissions, without making any one such drive taken in isolation wrong.

²⁵ One author who explicitly points out this difference between EU and the fair share approach is Hiller. Hiller is an advocate of EU, and notes explicitly that on his view a Sunday drive is pro tanto wrong, whereas on a 'equal shares' approach it is not. (Hiller 2011, p. 360) Hiller does not see this as a disadvantage of his view, but I am arguing that it is a feature that makes it implausible.

²⁶ Various authors take a fair share approach to obligations in collective impact cases. Baatz (2014) is one such author who explicitly talks about how the fair share duty he is advocating for is imperfect in the Kantian sense. Baatz argues that there is a collective duty to limit global emissions, and that as a result each person has a duty to not exceed their fair share of the global emissions budget. But how much is one's fair share depends on the details of one's circumstances, and all we can really say in general is that each person is morally required to reduce their emissions "as far as can *reasonably* be demanded of them." (Baatz, p. 10) This is an imperfect duty in the sense that it is not clear what exactly it requires, and so there is latitude for the individual to decide. As I mentioned, an imperfect duty in this Kantian sense is *also* an imperfect duty in my sense: on such a view an individual emissions-producing activity need not be wrong, even if there is nothing substantial speaking in favour of it. Baatz notes this feature of his account. However, he also argues that it is very unlikely that, in the actual world we live in, a drive in a gas-guzzling SUV would not involve exceeding one's fair share. (See p. 12)

²⁷ A related but different example can be found in Schroeder 2013, pp. 26-27.

myself am inclined to take a different approach.²⁸ But since my goal in this paper is only to argue that we should aim for a view that yields an imperfect obligation, and not to argue for any one particular such view, I will not get into this here.²⁹

5. Extension to Other Cases

My focus has been on the nature of our obligations as individuals with respect to our lifestyle choices and climate change. I have argued that we should aim for a view in which such obligations are imperfect. I think the same conclusion is correct in various other collective impact contexts as well. Suppose, for instance – as might be the case – that a large-scale change in demand from conventional to fair trade coffee would greatly reduce the exploitation of coffee farmers in developing countries, but that any one purchase won't, or at least is extremely unlikely to make a difference. The relevant decisions in the supply chain (about labour practices, who is employed on various farms, and so on) will almost certainly not go differently depending on any one ordinary individual's coffee purchases. Supposing these claims are true, this is a

²⁸ Very briefly: I've argued elsewhere that an act can make significant progress toward changing an outcome even if it cannot (or is extremely unlikely to) individually make a difference. I argue, for instance, that by refraining from driving I make small but real progress toward less harmful climate change outcomes, regardless of whether the harms of climate change will be different depending on my having done so. (See Nefsky 2017). If I am right about this, then I think this can ground an obligation. I ought to refrain from driving *because* doing so will make progress toward preventing severe climate change harms. So, in my view, being highly unlikely to make a difference does not undercut having an obligation grounded in what you individually can do toward changing the outcome – in the instrumental progress you can make. But what it does do is affect what it makes sense for this obligation to be – namely, it makes it an imperfect obligation. Because things do not, or are extremely unlikely to, depend on any one individual such act of yours, the obligation is imperfect: it is an obligation to make such choices 'enough' of the time, rather than to do so each time. On my view, fairness to other agents who are doing something toward preventing climate change harms might be a relevant factor in setting how much counts as enough, but – contrary to the fair share approach – fairness to such agents is not the primary grounds for the obligation to reduce one's emissions.

²⁹ A different possibility for how one might get an imperfect view is that one could argue that in collective impact cases we should apply expected utility to patterns of activity over time, or overall lifestyles, but reject the idea that the expected utility analysis works when applied to individual acts in these cases. On such an approach, one could say that Lavish's pattern over time is wrong because she is taking a miniscule but real risk of harm frequently, whereas Green's is not because she is only taking the risk very rarely. But it would say that if you take any individual act of Lavish's or Green's it is not wrong. This would be a major departure from the typical EU proposal, as explained in note 18.

collective impact case. A solution to the inefficacy challenge that yields an imperfect obligation here would say: each of us ought generally to opt for fair trade coffee, but the occasional purchase of non-fair-trade coffee can be permissible (even if there is no significant moral consideration in favour of it). Occasionally having a cup of coffee at a place that just sells ‘coffee’ (no further information available) is okay. But making a daily habit of it is not. It seems to me that this is the right sort of response to aim for. A response that condemns even the occasional purchase of conventional coffee (in the absence of some unusual weighty reason for the purchase) seems unreasonable, given that such a purchase almost certainly won’t make a difference for the worse to anyone. But a skeptical view which says consumers are under no moral constraints when it comes to such purchases seems like it cannot be right either, *if* it is really true that collective consumer choices in this domain have a major impact on the lives of thousands of workers. An imperfect account would be ideal.

Consider next the idea that most middle-class individuals in wealthy countries have an obligation to donate to efforts to aid the global poor. While this is often overlooked, the inefficacy challenge comes up here. It is often assumed that if one took the money that one would otherwise choose to spend on, say, a new pair of pants or movie tickets, and donated it instead to a famine relief fund, one would thereby be “preventing another person from starving.”³⁰ But this is doubtful. As Cullity says, faced with a famine, a relief agency would likely:

make a large-scale calculation of the size of the overall need, the amount of money they are likely to be able to raise from various sources to pay for it, and the extent to which it makes sense to draw on contingency funds in the light of other likely calls on them. Guided by these large-scale considerations, they then arrange the shipment of a large quantity of food.³¹

³⁰ As Singer (1972, p. 235) puts it.

³¹ Cullity 2004, pp. 58 - 59.

With this sort of decision-making process in mind, it seems highly unlikely that \$30 more or less in the agency's large pool of funds will make a difference to what aid is provided to people.

Suppose that is correct. If so, then a view which says that, even if someone is donating a large portion of their own income, they act wrongly if they spend some money on movie tickets or a nice dinner out (*unless* there is a weighty reason for them to do so, e.g. their own life projects depend on it) is highly implausible. And as with climate change, it is doubtful we can avoid that implausible result without falling into a skeptical conclusion, unless we adopt an imperfect view.

Many philosophers have argued for views of beneficence that count as imperfect in the relevant sense: these include both Kantian-style imperfect duties of beneficence (in which there is some indeterminacy involved), and also 'cap' or 'fair share' views of beneficence.³² But advocates of these views often assume that when you spend small amounts of money rather than donating it, you are passing up an opportunity to prevent some distant person from suffering or dying. They argue that, nonetheless, because of – for example – considerations of integrity or fairness or self-ownership, you do not need to donate every time that you could without significant cost to yourself. You just need to do so 'enough' overall. Wherever you stand on those arguments, my point is simpler and weaker. It is not about whether it is okay to pass up a

³² For an example of the second sort of account, see Murphy 2003. For an argument for the first sort of account, see Noggle 2009. See also Portmore 2012, especially pp. 35-36. Noggle raises a similar 'either too weak or too strong' dilemma with various moderate principles of beneficence to support the need for the sort of Kantian imperfect view that he describes. But I am making quite a different case here. Crucially, I am appealing to the claim that an individual contribution has a low probability of making a difference. This provides quite a different (or, if you want, an additional) basis that the 'too strong' side is indeed too strong. It no longer is (just) a point about how much sacrifice morality can demand of you. It is instead about what a fitting format for demands look like *when* the chances of individually making a difference by contributing to a collective pool are very low. Note that Noggle objects to the 'fair share' or 'cap' views on the basis that they allow someone to not save a person drowning in a pond in front of them as long as they do their share of beneficent contributions overall. (The idea is that an 'ends-based' imperfect conception of the sort that Noggle endorses avoids that result.) However, this argument against fair share or cap views does not have the same force insofar as we think that a cap or fair share duty is restricted to contexts in which one's individual contributions are highly unlikely to make a difference themselves (but in which very many such contributions by very many people will). This is why I am open to its being either the cap/fair share-style view or the Kantian sort of imperfect duty.

clear opportunity to prevent death or suffering. Instead, it is restricted to what we should say about our duty to make monetary contributions that are such that, while very many would make a large difference, any one is highly unlikely to make a difference itself. The point is that at least if that is the sort of context we are in, an imperfect obligation (in the broad sense specified) has great intuitive plausibility.

But while an imperfect view is highly intuitive in various collective impact cases, our intuitions might not go the same way in certain other cases. One case in which one might not find the imperfect view attractive or intuitive is that of meat consumption and factory farming. This is a collective impact case because, while a large shift in demand away from factory-farmed meat would result in many fewer animals being tortured on factory farms, it seems that any one purchase is not going to, or is very unlikely to, make a difference to the suffering of animals. The factory farm, which operates at a massive scale, is almost certainly not going to be making different production decisions in virtue of a single purchase.³³ But even so, for many of us, it might seem as though the facts about the suffering of animals on factory farms should make it impermissible to ever purchase factory-farmed meat – at least unless there is some unusual, strong reason to buy it (e.g. your life depends on it). An imperfect view might, thus, seem too weak.

There are two possibilities: this is right and we need a view that yields a perfect obligation (saying purchasing factory-farmed meat is *pro tanto* wrong), or this isn't right and it should be an imperfect obligation there too. Begin with the first possibility. So, suppose we should get a perfect obligation: buying factory-farmed meat is always wrong, in the absence of strong moral considerations in favour of doing so. Even if that is right, this is not an argument in

³³ Or at least, so it seems. I will not argue for this here. See Budolfson 2018 for an argument that this is indeed correct. Whether or not one agrees with the full argument in his paper, I think it at least shows that the chance is very small. However, see Hedden forthcoming for a reply to Budolfson.

favour of EU, or any other view (e.g. complicity) that is only capable of yielding a perfect obligation or none. After all, in the climate change case and others, a perfect obligation remains implausible. What we should aim for instead is a view that can *differentiate* the two kinds of cases: one that can yield an imperfect duty in the climate change case and a perfect duty in the meat case. EU cannot do that.

There are two ways one could go about differentiating the two kinds of cases. One is to come up with a single source of our obligations that yields an imperfect obligation in some cases, like the climate change case, but a perfect obligation in others, like the meat case.³⁴ The second is to argue that there are different moral factors at play in the two cases, and thus different sources for the obligations, such that in the one case there is a perfect obligation and the other an imperfect one. Neither of these paths is available within EU.

But now consider the other possibility: could it be that the imperfect view is right in the factory-farming case too? Keep in mind that just because we have an imperfect view, it need not be very permissive. How much counts as refraining ‘enough’ will depend on the details of the account. It could be quite restrictive, such that one must not eat factory-farmed meat almost all the time. What an imperfect view denies is that one can never permissibly purchase factory-farmed meat unless there is a strong morally relevant consideration in favour of doing so. A view

³⁴ I’m not sure whether or not the fair share approach can get a perfect obligation in the meat case, but it does provide an example of a single account of our obligations that can sometimes yield perfect obligations and sometimes imperfect. The fair share view can, for example, explain why the obligation is perfect in ‘one-off’ cases: cases in which each person only has one chance to contribute or not in the relevant way, rather than repeated chances over time. Voting in an election is like this. The fair share view can say that because we have an obligation to elect the better/less harmful candidate, each of us has a duty to do our share of that. And because each of us has only one chance to either vote or not in a given election, doing one’s share requires voting (for the better/less harmful candidate), even though doing so almost certainly won’t make a difference to whether or not that candidate wins. Thus, you are obligated to vote for the better candidate (or, perhaps, to abstain if you cannot come to a sufficiently informed decision), provided there is no weighty competing consideration. This counts as a perfect obligation in my sense. See Brennan 2009 for the fair share reply to the inefficacy objection to voting, including the idea that it might make abstaining obligatory for some people.

could deny that, and thus count as imperfect in the relevant sense, while still saying that it is wrong to buy factory-farmed meat any more than very rarely. Of course, it could be more permissive. But the point is that exactly how restrictive or permissive it is depends on the particular account, and is not settled by the claim that the duty is imperfect.

With this in mind, I don't think we should find it necessarily implausible that our obligations in the factory-farming case are imperfect too. Or at least, I don't think we should find it implausible *if* it is indeed true that any one purchase is extremely unlikely to make a difference to the suffering of animals (and to other bad consequences, like climate change.)

Many people who are vegetarian, and vegetarian on broadly consequentialist grounds having to do with animal suffering and wanting to prevent it, regard it as forbidden to ever purchase meat, and yet are totally comfortable taking at least occasional car trips just for convenience, or taking occasional airplane flights just for the sake of pleasure or career advancement. They are comfortable doing so, despite knowledge of the looming catastrophic threats to humans and animals of climate change. Why is that? Why do they have this combination of attitudes?

I am doubtful that it is because they know that the expected harm of a single meat purchase is much higher than the expected harm of an individual car trip or airplane flight. Nor, do I think it is because they are confident that the expected benefit of their driving or taking an airplane in each instance is substantial enough to outweigh the expected climate change harm. The difference in attitude is likely instead because of a phenomenological difference. The idea that it is very unlikely that one will individually make a difference is phenomenologically salient in the case of climate change, and is *not* so in the meat-purchasing case. Buying some chicken from the butcher does not feel like the kind of act that is very unlikely to make any difference at

all to chickens. If, however, one internalizes the idea that a single purchase of factory-farmed chicken is extremely unlikely to make a difference for the worse to any chickens, but *also* that many such purchases by very many people make a huge difference for the worse, then I think an imperfect obligation becomes intuitive here too.³⁵

Acknowledgements

Many thanks to Zach Barnett, Mark Budolfson, Matthew Hanser, Kian Mintz-Woo, Theron Pummer, Alex Rennet, Hamish Russell, Andrew Shroeder, Sergio Tenenbaum, and Jordan Thomson. I am also very grateful to audiences at The Centre for Ethics, Philosophy and Public Affairs at the University of St. Andrews; the Popper Seminar at the London School of Economics; the 2017 Meeting of the Pacific APA; the 2017 Meeting of the Canadian Philosophical Association; and, the 2018 Halbert Network Workshop on Collectively and Temporarily Extended Rights and Wrongs. Thanks also to the participants in the 2016 Bled Conference on Ethics, where I gave a talk that was very different from this paper but a precursor to it. My work on this paper was supported by an Insight Grant from the Social Sciences and Humanities Research Council of Canada; I am very grateful for this support.

References

Arntzenius, F. and McCarthy, D. (1997). "Self Torture and Group Beneficence." *Erkenntnis* 47: 127-144.

Baatz, C. (2014). "Climate Change and Individual Duties to Reduce GHG Emissions". *Ethics, Policy & Environment* 17(1): 1-19.

Brennan, J. (2009). "Polluting the Polls: When Citizens Should Not Vote." *Australasian Journal of Philosophy* 87(4), 535-549.

³⁵ Note that if the obligation is imperfect, abstaining completely is (most likely) morally even better – it is just not required. Also, for many individuals it might be psychologically easier to abstain entirely than to moderate their consumption. For individuals who are like that, the most promising strategy for meeting an imperfect obligation not to eat factory-farmed meat would be to treat it as if it were a perfect obligation.

- Broome, J. (2019). "Against Denialism". *The Monist* 102: 110-129.
- Broome, J. (2012). *Climate Matters: Ethics in a Warming World*. New York: W. W. Norton.
- Budolfson, M.B. (2018). "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequence Response," *Philosophical Studies*, <https://doi.org/10.1007/s11098-018-1087-6>
- Budolfson, M.B. (ms). "Collective Action, Climate Change, and the Ethical Significance of Futility", version 5.2.
- Cullity, G. (2004). *The Moral Demands of Affluence*. New York: Oxford University Press.
- Cullity, G. (2015). "Acts, Omissions, Emissions". *Climate Change and Justice*. Ed. Jeremy Moss. Cambridge: Cambridge University Press, pp. 148-164.
- Driver, J. (2015). "Individual Consumption and Moral Complicity". *The Moral Complexities of Eating Meat*. Eds. Ben Bramble and Bob Fischer. New York: Oxford University Press.
- Hanser, M. (2014). "Imperfect Aiding". *The Cambridge Companion to Life and Death*. Ed. Steven Luper. Cambridge: Cambridge University Press.
- Heal, G., and Millner, A. (2014). "Uncertainty and Decision Making in Climate Change Economics". *Review of Environmental Economics and Policy* 8(1): 120-137.
- Hedden, B. (forthcoming). "Consequentialism and Collective Action". *Ethics*.
- Hiller, A. (2011). "Climate Change and Individual Responsibility". *The Monist* 94(3):349–368.
- Intergovernmental Panel on Climate Change. "Special Report: Global Warming of 1.5°C", <https://www.ipcc.ch/sr15/>
- Kagan, S. (2011). "Do I Make a Difference?" *Philosophy & Public Affairs* 39(2): 105-141.
- Kingston, E. and Sinnott-Armstrong, W. (2018). "What's Wrong with Joyguzzling?" *Ethical Theory and Moral Practice* 21: 169-186.
- Kutz, C. (2000). *Complicity*. Cambridge: Cambridge University Press.
- Lawford-Smith, H. (2016). "Difference-Making and Individuals' Climate-Related Obligations". *Climate Justice in an Idealised World*. Ed. Clare Hayward and Dominic Rosers.
- McClennen, E.F. (1990). *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge: Cambridge University Press.

- McPherson, T. (2016). "Why I am a Vegan (and You Should be One Too)." In *Philosophy Comes to Dinner*, edited by Andrew Chignell, Terence Cuneo, and Matthew Halteman. New York: Routledge, 73-91.
- Morgan-Knapp, C. and Goodman, C. (2015). "Consequentialism, Climate Harm and Individual Obligations" *Ethical Theory and Moral Practice* 18(1): 177-190.
- Murphy, L. (2000). *Moral Demands in Nonideal Theory*. Oxford: Oxford University Press.
- Nefsky, J. (2011). "Consequentialism and the Problem of Collective Harm: A Reply to Kagan." *Philosophy & Public Affairs* 39: 364 -395.
- Nefsky, J. (2017). "How You Can Help, Without Making a Difference." *Philosophical Studies* (174): 2743-2767. <https://doi.org/10.1007/s11098-016-0808-y>
- Nefsky, J. (2018). "Consumer Choice and Collective Impact". *The Oxford Handbook of Food Ethics*. Eds. Anne Barnhill, Mark Budolfson, and Tyler Doggett, Oxford University Press, 267-286.
- Nefsky, J. (2019). "Collective harm and the inefficacy problem". *Philosophy Compass*. DOI: 10.1111/phc3.12587
- Noggle, R. 2009. "Give Till it Hurts? Beneficence, imperfect duties, and a moderate response to the aid question." *Journal of Social Philosophy* 40(1): 1-16.
- Norcross, A. (ms). "More Mistakes in Moral Mathematics".
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Clarendon Press.
- Portmore, D. (2012). "Imperfect Reasons and Rational Options". *Nous* 46(1): 24-60.
- Regan, D. (2000). "Perceiving Imperceptible Harms: With Other Thoughts on Transitivity, Cumulative Effects, and Consequentialism". In *Imperceptible Harms and Benefits*. Ed. M. Almeida. Dordrecht: Kluwer Academic Publishers, 49-73.
- Shroeder, A. "Imperfect Duties, Group Obligations, and Beneficence", *Journal of Moral Philosophy* 11:5 (2014)
- Singer, P. (1972). "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1(1): 229-243.
- Sinnott-Armstrong, W. (2005). "It's Not My Fault: Global Warming and Individual Moral Obligations". In *Perspectives on Climate Change: Science, Economics, Politics, Ethics*. Ed., Walter Sinnott-Armstrong and Richard B. Howarth. Amsterdam: Elsevier, 285-307.

Spiekermann, K. (2014). "Small Impacts and Imperceptible Effects: Causing Harm with Others." *Midwest Studies in Philosophy XXXVIII*: 75 – 90.

Tenenbaum, S. and Raffman, D. (2012). "Vague Projects and the Puzzle of the Self-Torturer." *Ethics* 123 (1): 86-112.