

2015

Consequentialism, Climate Harm and Individual Obligations

Christopher Morgan-Knapp
c.morgan-knapp@binghamton.edu

Charles Goodman
Binghamton University--SUNY

Follow this and additional works at: https://orb.binghamton.edu/philosophy_fac

 Part of the [Philosophy Commons](#)

Recommended Citation

Morgan-Knapp, Christopher and Goodman, Charles, "Consequentialism, Climate Harm and Individual Obligations" (2015).
Philosophy Faculty Scholarship. 27.
https://orb.binghamton.edu/philosophy_fac/27

This Article is brought to you for free and open access by the Philosophy at The Open Repository @ Binghamton (The ORB). It has been accepted for inclusion in Philosophy Faculty Scholarship by an authorized administrator of The Open Repository @ Binghamton (The ORB). For more information, please contact ORB@binghamton.edu.

CONSEQUENTIALISM, CLIMATE HARM AND INDIVIDUAL OBLIGATIONS

Christopher Morgan-Knapp and Charles Goodman
Binghamton University

Steven Gardiner recently proposed a do-or-die test for moral theories. If a moral theory does not acknowledge that failure to address a pressing, anthropogenic, and tractable global threat is a serious criticism of it, he says, then that theory is “inadequate and must be rejected.” (Gardiner 2011, 218)

We might wonder what moral theory could possibly fail such a test. It is, after all, an awfully low hurdle. To pass the test, a theory doesn’t need to successfully address the global threat: it merely needs to acknowledge that failing to successfully address it is a problem. And how could a serious moral theory simultaneously recognize that something is a significant moral problem, and then imply that doing nothing about it is just fine?

According to some ethicists, though, this is just what act-consequentialism could do in response to anthropogenic climate change. In particular, they think that even if widespread, voluntary reductions in carbon emissions are both necessary and sufficient to avoid a climate catastrophe, act-consequentialism will counsel against making them.¹ Their first point is that, in consequentialist terms, anthropogenic climate change should be avoided: the world we are in for if nothing is done to avert (or at this late date, diminish) climate change is much worse than worlds we could bring about by choosing to emit less carbon into the atmosphere. Their second point is that act-consequentialism implies that the vast majority of us should not reduce our emissions. And this will be so, they say, even if such reductions are needed to avoid the awful consequences of climate change. Here is how one such critic, Ronald Sandler, makes the case for this second point:

Almost any action performed by almost any agent will have a vanishingly small effect on [climate change]. Many of these same actions will have burdens for the agent and those close to her (i.e., family and friends), in terms of, for example, time, economic costs, social costs, and professional

¹ Gardiner himself says that “standard utilitarian thinking (such as cost-benefit analysis) might well fail the ... test.” (241) But his worries primarily concern the use of act-consequentialism as a decision procedure and hence are different from the more theoretical worries we will be discussing.

costs. In such cases the local utility of actions that contribute to [climate change] ... will outweigh the inconsequential global utility of those actions. Therefore, act utilitarianism cannot explain why we ought to act or live in ways [needed to mitigate climate change] when doing so has costs or sacrifices associated with it. (Sandler 2010, 170-1)

Sandler takes his argument to be reason for rejecting consequentialism, but other theorists are willing to bite the bullet: they take the same reasoning to show that we aren't really morally obligated to emit less. After characterizing the problem of climate change as a tragedy of the commons (which he abbreviates 'T of C'), Baylor Johnson asks whether taking unilateral, individual steps to lower emissions is morally required. His answer:

The only reason to adopt unilateral restraint is to avert a T of C. So if unilateral restraint cannot reasonably be expected to achieve its purpose, there is no reason, and hence no moral reason to adopt it. I claim that averting a T of C is the only reason for adopting unilateral restraint because in a T of C there is nothing wrong with any one person's use of the commons. No one person's use is large enough to harm the commons. Harm results only from the aggregate level of use. (Johnson 2003, 277)

Walter Sinnott-Armstrong, too, argues against an individual moral obligation to reduce our emissions grounded in the claim that individual emissions make climate change worse. "The problem with this argument," he says, is

that my act of driving does not even make climate change worse. Climate change would be just as bad if I did not drive. Global warming and climate change occur on such a massive a scale that my individual driving makes no difference to the welfare of anyone. (Sinnott-Armstrong 2005, 301)

Similar reasoning can also underlie the idea that consequentialists should use something other than actions as their primary evaluative focus. Take, for instance, Dale Jamieson's defense of the idea that when moral problems like climate change are at issue consequentialists should focus on character traits rather than actions. There he claims that "Joy-riding in my '57 Chevy will not in itself change the climate, nor will refraining from driving stabilize the climate..." and takes this to be an instance of a general problem; namely that if we look only at the consequences of actions in "large-scale collective action problems ... it appears that both morality and self-interest demand that 'I get mine,' since whatever others do, it appears that both I and the world are better off if I fail to cooperate." (Jamieson 2007, 167)

Act-consequentialists will be happy to hear that these authors are all mistaken: as we show in what follows, embracing act-consequentialism should not lead to complacency, much less a sense of righteousness, regarding our individual emissions. Quite the contrary, it implies that most of us should be making immediate and significant reductions in our contributions to the global stock of greenhouse gases. Act-consequentialism does not fail Gardiner's test.

But it is not only act-consequentialists who should be heartened. Johnson and Sinnott-Armstrong both seem to find their own conclusions regarding the moral status of frivolous emissions counter-intuitive, if not morally dubious. It seems they embrace them only because that is where their arguments take them. Understanding why spending one's Saturday speeding up and down the river on one's powerboat is wrong because of its contribution to climate change should thus help them bring their gut intuitions in line with their commitment to rational analysis.

Deontologists, too, should welcome this result. As we will detail below, at the root of these challenges to act-consequentialism is the idea that individual emissions are harmless. If that were right, a deontological duty to reduce one's emissions that appeals to a constraint against harming, or against risking harm, could not even get off the ground: a constraint against harming can't prohibit doing something harmless. But while deontological approaches may well face other difficulties in this context, our analysis will show why this is not one of them.

And finally, the reasoning at the core of this challenge to act-consequentialism is hardly esoteric. We suspect that many outside of academic philosophy think of their own lifestyle choices in similar terms. That is, though they acknowledge that their carbon footprint is heavier than it needs to be, many also conclude that this does not matter. If they were to give up their powerboat, or their exotic vacations, or their air conditioning, they think, this would not change anything regarding the climate. All it would really do is diminish their own contentment. Even from a moral perspective, then, it seems to them pointless. To the extent that such thinking is common, understanding where it goes wrong is an important step towards helping people make the choices that are necessary to avert a climate catastrophe.

§1. Preliminaries

Our core thesis is that act-consequentialism implies that luxury greenhouse gas emissions are typically wrong. Before mounting our defense, though, let us begin with some clarifications.

To say that some type of action is typically wrong is ambiguous. Following Parfit, we can distinguish two senses of ‘wrong’ that are germane to the current discussion. First, we might say that an action is ‘wrong in the fact-relative sense’ just when it would be morally impermissible to choose the action if we knew all the relevant facts. By contrast, we might say an action is ‘wrong in the evidence-relative sense,’ just when it would be morally impermissible to choose the action if what the available evidence gives us reason to believe were true.²

To see the difference, consider one of Parfit’s examples: You know that 100 miners are trapped underground, with floodwaters rising. You know that they are all in one of two shafts, but have no reason to believe they are in one rather than the other. You can choose to close one of three flood-gates, and you know that the consequences of your choice would be as follows:

		THE MINERS ARE IN	
		SHAFT A	SHAFT B
YOU CHOOSE	GATE 1	You save 100	You save 0
	GATE 2	You save 0	You save 100
	GATE 3	You save 90	You save 90

In the fact-relative sense, it would be wrong to close Gate 3. This is because if you knew all the relevant information, you would know where the miners were, and hence could save more lives by choosing to close either Gate 1 or Gate 2. But in the evidence-relative sense of ‘wrong,’ it would be wrong to choose either Gate 1 or Gate 2. This is because given the available evidence, they could be in either Shaft A or B, and so by closing one of them you have only a 50% chance of saving 100, whereas if you close Gate 3 you are sure to save 90. (Parfit 2011, 159)

We will have something to say about what act-consequentialism implies about the wrongness of emissions in the fact-relative sense. But when we say that act-consequentialism implies that luxury emissions are typically wrong, we mean this in the evidence-relative sense. This is because the evidence-relative sense is the sense of wrongness that is relevant when we are trying to decide what to do. When

² This distinction roughly corresponds to the distinction others have labeled with the terms ‘subjective rightness’ and ‘objective rightness.’ As we will explain, it is this distinction that lies behind the shift some utilitarians have made from actual to expected utility as the target choices are aimed at, morally speaking.

making decisions, we often do not – and indeed cannot – know all the relevant facts. We are making decisions under uncertainty, and hence are really choosing which risks to run rather than which outcome to bring about. This is especially clear when it comes to greenhouse gas emissions, which may stay in the atmosphere for centuries: we cannot have justified beliefs concerning precisely what the particular molecules of CO₂ we put into the atmosphere now will cause 200 years from now. So the pressing question prospectively is what we should and shouldn't decide given the evidence that is now available. And since the authors with whom we disagree are all interested in understanding what we should now decide to do about greenhouse gas emissions, they too are, or at least should be, focused on what is wrong in the evidence-relevant sense.

The second point that needs clarifying is what we mean by 'act-consequentialism.' Common to all versions of act consequentialism is the idea that the wrongness of an action is solely a function of the impersonal value of its outcome. But different variations on this core idea can be more and less plausible analyses of different senses of 'wrong.' One straightforward version of act-consequentialism holds that actions are wrong if there is some available alternative action whose outcome would be impersonally better. This version is a plausible consequentialist analysis of wrongness in the fact-relative sense. But it is not a plausible consequentialist analysis of wrongness in the evidence-relative sense. Consider again, the case of the miners. This theory would imply that your choosing Gate 3 is wrong, because even given only the available evidence, you can be sure that there is an available alternative that has a better outcome. As we suggested above, though, choosing Gate 3 does not seem wrong in the evidence-relative sense.

A far better act-consequentialist analysis of evidence-relative wrongness holds that an action is wrong just in case there is an available action whose outcome has a better expected value. The expected value of an action's outcome is the sum, over all possible outcomes, of the value of each possible outcome multiplied by the chances of that outcome occurring.

A final comment is in order regarding 'luxury' emissions. Here we follow many in making a distinction between increasing atmospheric carbon concentrations by, say, heating one's private swimming pool and by eating enough vegetables to stay healthy. The former is a luxury emission, while the latter is what is often called a 'subsistence' emission. (Shue 1993) There are several ways of generalizing this distinction. In what follows, we will use these terms to refer to the position any particular emission occupies on a spectrum that ranges from the most beneficial to the most frivolous emissions. On one side of the spectrum – the subsistence side – are emissions that make a very significant contribution to people's quality of life,

while on the other side – the luxury side – are those whose contribution is comparatively trivial.³

So to sum up, when we say that act-consequentialism implies that most luxury emissions are wrong, we are saying that when most of us increase atmospheric carbon in ways that do relatively little to improve our lives, there are alternative choices we could make whose outcomes have a higher expected value, and consequently that causing such emissions is morally wrong in the evidence-relative sense.

§2. The Initial Analysis and Two Challenges

We start with an assumption that would be granted by all the authors we are engaged with: things will not go as well as they could unless emissions are reduced.

Given this assumption, there is a straightforward reason for thinking that most luxury emissions are wrong according to act-consequentialism. Since reducing the aggregate level of emissions could improve on the status quo, and choosing to forgo luxury emissions would reduce aggregate emissions without causing much hardship, it seems things would be better if we choose not to emit luxuriously. And if so, some luxury emissions must be wrong according to act-consequentialism. (Exactly which luxury emissions will be wrong will be a function of how far above the optimal level of emissions we currently are. Here is one way to visualize how this would work: imagine ordering all emissions on our spectrum, from the most beneficial to the least. Now, starting from the most beneficial emissions, add the emissions in order until the optimal level of emissions is reached. Any emissions that would lie beyond this point would be wrong, provided the straightforward analysis holds.)

Avram Hiller makes a very similar point. The crucial issue, he recognizes, is how much harm a luxury emission like a Sunday drive creates. And, he says, “there is a fairly elementary way to determine the answer:

(Step 1) Estimate the amount of GHG emitted by the one drive, d.

(Step 2) Estimate the total amount of GHG emissions responsible for climate change, e.

(Step 3) Estimate the total amount of harm that climate change will cause, h.

³ Where any given emission falls will depend, of course, on the correct theory of well-being. While resolving this will sometimes be important in practice, it won't be necessary for the more theoretical points being made here.

(Step 4) Calculate $(d/e) \times h$.”

Performing that calculation, he suggests, yields the result that a Sunday drive produces harm that “is the moral equivalent of ruining someone's afternoon.” (Hiller 2011, 357-8)

This straightforward analysis trades on the thought that the disvalue of global emissions can be simply disaggregated: that the disvalue of each individual emission is simply its share of the disvalue of total emissions. To this we can add that, if we grant simple disaggregation, and we grant that it is better for aggregate emissions to go down, then it must be better if those individual emissions that produce relatively little good are eliminated. The aggregate value, we might think, is just the sum of its parts, and so if the aggregate value is sub-optimal, then the least beneficial of its parts must be sub-optimal as well.

Implicit or explicit scepticism about simple disaggregation, however, is what leads many to the conclusion that individual decisions to avoid luxury emissions do not make things any better. How could this be? Two lines of thought seem initially plausible. The first is that many of our luxury emissions will occur whether we choose to avoid them or not. A compelling example is airline travel. On the simple disaggregation model, when one flies to Europe for vacation, one is responsible for emitting the amount of carbon that the plane emits divided by the number of passengers on the plane. But, as several have pointed out, attributing this fraction of the plane's emissions to your choice is suspect.⁴ If you had chosen against taking the European vacation, they say, your flight would have gone anyway and the same amount of carbon would have been emitted. So whether or not you choose to be on the plane makes no difference to the amount of carbon released into the atmosphere.

This reasoning can be generalized quite broadly. When we buy things, even including fuel, it is typically true that if we had not bought them, somebody else would. So the emissions that are connected to that product will occur whether we are the purchaser or not. It is not as if the particular carbon atoms that are released into the air when someone fires up their jet-ski would have stayed in the ground otherwise; those atoms were already in the energy pipeline, and if the jet-skier hadn't burned that gas, someone else would have. Given this, it can be hard to see what good forgoing a spin around the lake will do, at least as far as the

⁴ For instance, Joaquim Sandberg asks “How much of a [passenger plane's] pollution does an individual passenger cause?” and answers “...it would seem that my behavior actually has no marginal effect here. Just as much carbon dioxide will be emitted irrespective of whether I [am on the plane].” (Sandberg 2011, 232).

climate is concerned – that carbon is going to find its way into the atmosphere no matter what.

The second source of scepticism about simple disaggregation is more philosophical. It starts from the idea that what makes luxury emissions bad in the aggregate is their affects on welfare. The next step is to claim that welfare is affected for the worse only if there is some *discernible* difference in the subject's life: if their experience is different, or their desires are less satisfied, or so on. But, it seems, the climate is such a vast system, and each individual decision to emit carbon so small, that individual emissions cannot make any discernible difference. While trillions of luxury emissions might change the climate in ways that make a noticeable and negative difference to the lives of many Bangladeshis, for instance, no Bangladeshi will ever feel any different if you choose to lower your thermostat by a degree on a particular day.

Unless we can dissolve these two challenges, act-consequentialism is in trouble vis-à-vis Gardiner's test. If, for either reason, no good comes from most individual decisions to forgo luxury emissions, then act-consequentialism will not give us grounds to object to them. But in the aggregate, such emissions are very bad. So act-consequentialism would tell us not to do what is necessary in order to avoid what is a morally terrible outcome, even in consequentialism's own terms. The task for us then, is to show what is wrong with both of these challenges.

§3. Market Thresholds

In a discussion of the relationship between utilitarianism and vegetarianism, Peter Singer showed how act-consequentialists have the theoretical resources to justify moral criticism of behaviors such as airline travel.⁵ The problem for Singer was why utilitarians should not buy meat, given that a single decision not to buy meat seemed to make no difference to the number of animals raised and killed for food, but his solution is directly applicable here. We can start by imagining that you would like to go to Europe on vacation, and that the flight you would take typically carries 100 people. True, if you choose to take the flight, it is unlikely that this will make any difference to how much jet fuel is burned. But if 100 additional people want to fly to Europe on the same airline, the airline will add an additional flight to meet the demand. Of those 100 people, there must be one whose decision made the difference, crossing the airline's threshold for adding another flight.

⁵ Singer 1980, pp. 335-336. The strategy is also invoked in Matheny 2002 and Kagan 2011.

Might you be that person? You have, let us suppose, no evidence that would license you to conclude that you are any more or any less likely than anyone else to be the passenger who causes another flight to be added. So it would be most reasonable for you to assign a subjective probability of 1% to this outcome. But if you are the threshold passenger, then the marginal impact of your decision is that one more flight is added. In this case, your individual choice adds an amount of carbon to the atmosphere sufficient to transport 100 people to Europe. So the expected amount of carbon to be emitted as a result of your decision is 1% of the carbon normally used to transport 100 people: in other words, exactly the average amount of carbon normally emitted per passenger.

How this analysis of market thresholds figures into an act-consequentialist account of the morality of luxury emissions depends on whether we are interested in wrongness in the fact-relative or evidence-relative senses. If we were concerned with the fact-relative sense, our focus would be on the actual value of the outcomes resulting from each person's choice. For 99 of the people in our example, their airline travel produces no emissions. But the trip taken by the one person that pushed the airline over the threshold produced, individually, the amount of emissions that an entire flight to Europe produces. Far from being innocuous, then, this particular person's choice made a significant difference to the amount of carbon in the atmosphere. Provided the other challenges can be met, that person's choice of vacations is thus likely to be wrong in the fact-relative sense.

How our travel decisions will impact airline schedules is a matter of considerable uncertainty, however, and in such circumstances, what we are really interested in is what is wrong in the evidence-relative sense. Here it is not the actual value of our choices, but their expected value that is relevant. In this context, the effect of Singer's reasoning is to bypass completely the effect of skepticism about simple disaggregation. Instead of naively disaggregating by merely assigning an average share to each passenger, we compute the probability of being the threshold passenger, and then put all the negative consequences caused by crossing the threshold onto that outcome. Mathematically, however, the result will turn out to be precisely the same. Each passenger should deliberate as if he or she was causing a proportionate share of the damage; and if the value produced by the choice is less than that share, the passenger should not take the flight.

Since this reasoning can be extended to all cases in which a person's luxury emissions seem to make no difference to the amount of carbon that is released into atmosphere, the first challenge to the initial act-consequentialist analysis of luxury emissions can be met.

§4. Meteorological Thresholds

Despite its success with the example of the European vacation, though, it may well seem that Singer's threshold analysis can't help us with luxury emissions like recreational weekend driving. A small probability of causing harm can give us a reason not to act. But if, as many claim, your driving on the weekend has no probability at all of causing harm, then an argument based on expected value will have no purchase on it. And it seems intuitively plausible that, given the vast scope of the global climate crisis, your individual contribution will make no difference. But is this plausible claim really true?

The reasoning for the behind the idea that my contribution makes no difference can be broken down into two steps. First is the idea that the effect of individual emissions on the weather is miniscule: the planet's meteorological system is so large, and the size of individual emissions so tiny, that whatever impact an individual emission has on the weather must be vanishingly small. The second is that vanishingly small impacts aren't morally relevant because no one could possibly tell the difference between such an impact occurring and it not occurring. Perhaps your emission could add one droplet of water from a flood or remove one droplet of water from a drought. But if no one can tell the difference such tiny impacts make, then there is nothing bad about them occurring. Events are only harmful if they make a difference to someone's experience, and if no one could ever tell if a drought or a flood contained one more or less droplet of water, then this kind of impact can't be experienced and so isn't harmful.

This kind of reasoning, however, is invalid. Indeed, both steps are incorrect. In this section we will discuss the first, and in the subsequent section, the second.

Can the effects of individual luxury emissions be anything but miniscule? Sinnott-Armstrong believes not: "No storms or floods, or droughts or heat waves can be traced to my individual act of driving," says he. (Sinnott-Armstrong 2005, 299) Unfortunately, this empirical claim is false. Scientists have known for decades now that the weather is a chaotic system, in the technical sense of exhibiting sensitive dependence on initial conditions. As scientist Richard Kautz writes, "By 1964 ... numerical experiments with full-scale weather models confirmed the idea that small perturbations grow exponentially in time." (Kautz 2011, 161) In other words, if we start with two descriptions of the state of the atmosphere that differ only in a tiny way, and then allow each description to evolve in accordance with the equations that we believe explain how the atmosphere changes, they will grow more and more different with time. Indeed, according to Kautz, in our most up-to-date weather models, the difference between the two descriptions will double every few days. As a result, no matter how well we come to understand the atmosphere, scientists believe that it will remain impossible in practice for us to predict the weather beyond about two weeks in the future. (Kautz 2011, 161)

This means that your decision to take a recreational drive may set off what is sometimes called the “butterfly effect.” The effect is named after Edward Lorenz’s famous question: Could a butterfly flapping its wings in Brazil lead to a tornado in Texas? (Kautz 2011, 13) If we compare the Earth’s future in the scenario in which you drive to the future in the scenario in which you don’t, the initial segments of these two futures will be almost exactly similar. But the chaotic mathematics of the equations that describe the atmosphere will cause these scenarios to grow more and more different from each other with time. There is some very small but non-zero probability that the scenario in which you do not drive features a hurricane season next year that contains three hurricanes, whereas the scenario in which you do drive features four hurricanes. An act-consequentialist interested in wrongness in the evidence-relative sense must take into account the product of this very small probability with the huge amount of damage that an additional hurricane would cause.⁶

An opponent such as Johnson or Sinnott-Armstrong could object at this stage that our argument threatens to lead to paralysis. Suppose that, having abandoned your intention to take a recreational drive, you instead begin to deliberate about whether to take a walk in the park or a stroll by the lake. There is some tiny probability that, due to the butterfly effect, your walk in the park will lead to a future in which there is one extra hurricane. Does this probability generate a serious moral objection against absolutely anything you might choose to do?

Well, exactly the same consideration applies to a stroll by the lake: it too might generate an extra hurricane. Equivalently, we could say that your walk in the park might prevent one hurricane. And 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. So when you consider the vast probability distribution of all the results that might come from your choice throughout the future of the earth’s weather, you should be able to see that your subjective model of the distribution is rationally required to be symmetrical with respect to this choice. All the possible positive indirect effects and all the possible negative indirect effects of your decision to walk in the park instead of the lake will cancel each other out, leaving you with no climate-related reason to choose one or the other. As far as the ethics of climate change are concerned, you are free to walk in whichever place you would enjoy most.

⁶ Hiller 2011 suggests that the possibility of the planet being pushed across ‘tipping points’ where feedback loops are started that lead to relatively sudden and drastic changes should be handled the same way (361).

Similarly, Kautz points out that the answer to Lorenz' question is yes: the tiny effects of a butterfly's wings could interact with the unstable dynamics of the weather to produce a tornado that would not otherwise have occurred at that time. And yet, in another sense, butterflies don't cause tornadoes:

Because the solar heating isn't affected, the weather exhibits the same kinds of events, from tornados to rainstorms, that we're accustomed to. The effect of the butterfly is simply to change the schedule of those events. (Kautz 2011, 162)

We could say the same about the walk in the park as opposed to the stroll by the lake.

Matters are otherwise with the choice to take a recreational 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.

How should you assess the slight asymmetry in the climate future of the earth in the scenario in which you take a recreational drive, as compared to the future in the scenario in which you take a walk in the park? Suppose you reasonably believe that you have some scientific understanding of what 3°C of global warming would look like. Suppose that in that scenario, various harmful weather events would take place that would not have occurred in a zero-warming scenario: more hurricanes, more tornadoes, more floods, more droughts, and so on. Each of these harmful weather events kicks in at a certain level of greenhouse gas emissions. But you have no idea what that level is for any of the events. In fact, if you were to break up all the greenhouse gas emissions between zero warming and 3°C of global warming into packets equivalent to the amount of greenhouse gases your car would emit during a recreational drive, you would have no reason to think that the packet from this particular recreational drive is either any more, or any less, likely than any other to cause any one of these harmful weather events.

Given this setup, we can apply Singer's threshold analysis in a more complicated way, and get the same kind of striking conclusion we saw before. When you multiply the harm caused by each of the weather events you are considering by the probability that your recreational drive will cause the threshold for that event to be crossed, and then sum all the products, you will end up with an amount of expected harm that is precisely the overall harm of the increased incidence of harmful weather events caused by 3°C of climate change multiplied by one recreational drive's share of the emissions that would cause that much climate change. (For expository reasons, this analysis focuses on the number of significant

adverse weather events that occur, and not on the precise degree of their intensity, a topic to be considered in the next section.)

Of course, 3°C of climate change would also cause some beneficial effects. You could take credit for your recreational drive's tiny share of the probability that each of these beneficial effects occurs. But we are assuming for the sake of argument that the overall harm caused by 3°C of climate change greatly exceeds the benefits. So you still have a consequentialist moral reason not to drive whose strength corresponds to your recreational drive's proportionate share of the net harm caused by 3°C of climate change.

We have simplified the analysis by looking at a binary choice between no warming and the specific 3°C scenario. To be fully rigorous, we would want to take into consideration the probability distribution of the various amounts of warming that could occur as a result of human activity. The resulting calculation would be of ferocious complexity. To actually derive numbers would require all kinds of information that we clearly can't collect. But suppose we are allowed to assume that large amounts of global warming would have many bad consequences, and more intense global warming would be even more damaging than less intense warming. We know that the actual calculation, which we can't do, would assign to the contemplated act of recreational driving its tiny proportionate share of a wide range of outcomes, ranging from bad to cataclysmic, with only a few positive outcomes to offset them. So even if we can't know what the overall result would be, we can have high confidence that the result of the expected harm calculation would be bad enough to give you a moral reason not to take the weekend recreational drive.

Let's sharpen up that last claim. You would enjoy your weekend drive, let's suppose, a bit more than your walk in the park. But part of what we are assuming is that if everyone were to engage unconstrainedly in activities that emit greenhouse gases, the results would be so bad for future generations that the history of the human race would go much worse than in a low-emissions scenario, even taking into consideration the benefits that present people would derive from enjoying their unconstrained polluting activities. You have no reason to believe that you are special in this regard. So when we use threshold analysis to assign to your weekend drive its proportionate share of the damage from climate change, we are licensed to regard that proportionate share as morally outweighing the enjoyment of driving that is your share of the benefits from polluting. Therefore, if you are an act-consequentialist, you should conclude that it is morally wrong for you to take the weekend drive.

§5. Perceptibility Thresholds

So far we have argued that there is a small chance that individual emissions can have significant effects on the weather. It might be that this, when combined with Singer's analysis of the moral relevance of thresholds, is enough to undermine the argument that act-consequentialism can't oppose weekend driving. But we can't be sure. This is because not all the bad consequences of climate change can be attributed to increases in the number of extreme weather events. Some of the bad consequences are much more graduated and hence the changes over time are much more subtle: temperatures and sea levels gradually rise; water tables and crop yields gradually fall; and so on. Moreover, some bad consequences will be due to weather events being more intense than they would otherwise be. In order to be confident in our inference from the badness of aggregate emissions to the badness of individual luxury emissions, we must be sure that *all* the bad effects of climate change can be disaggregated.

As noted above, there are some who think that this is not possible when it comes to very gradual changes or changes in the intensity of harmful weather events. The reason they are skeptical is that they believe that the difference any individual emission makes with respect to such changes is so infinitesimal that no one's life will be any different whether the change associated with a single emission occurs or not.⁷ And if no one's life is made better or worse by an infinitesimally small change, then a consequentialist cannot morally object to an emission that causes it.

Frank Arntzenius and David McCarthy, however, have shown in another context that this line of reasoning, too, is invalid.⁸ Their paper deals with a fictional example involving the administration of electrical shocks of varying intensity, but we can translate their argument into the terms of the present debate. Start by focusing on a single Bangladeshi who would suffer terribly from the accumulation of many subtle effects of climate change caused by future emissions. There is a very finely graded spectrum of greenhouse gas concentrations in the atmosphere

⁷ This might be what Jamieson has in mind when he draws analogy between individual emissions and the following case: "I, along with many other people, toss an invisible smidgen of something into a blender. A man takes a drink of the resulting mixture. Am I responsible for the graininess of the texture, the chalkiness of the taste, the way it makes him feel after drinking it, his resulting desire for a Budweiser? You might think that I am a smidgen responsible, since a smidgen is the amount that I tossed into the blender. But I am tempted to say that I am not responsible even for a smidgen of the result because there are so many thresholds, non-linearities, and scalar differences that intervene between my action and the outcomes." Jamieson 2014, p. 164.

⁸ Arntzenius and McCarthy 1997, pp. 132-135. Norcross 1997, §IV and Kagan 2011 develop essentially the same point.

that starts from the concentrations there would be if emissions stopped right now and ends with the concentrations sufficient to cause our Bangladeshi terrible suffering. (For the present purposes, we'll have the concentrations at adjacent points on the spectrum differ by the amount of greenhouse gases emitted on a recreational drive, but nothing hinges on this: we could have them even more finely differentiated.) Now imagine that we could run the following experiment: we allow the Bangladeshi to experience how things would be under the climate conditions produced by each of those many, many possible atmospheric concentrations while holding everything else constant. We give her a very long series of experiences, where each one is of how she would feel under a randomly chosen greenhouse gas concentration on our spectrum. And for each of these experiences, we ask her to report as accurately as she can how things are for her. We then tabulate the results: we list the greenhouse gas concentrations in order, and for each concentration, we group together all the reports she gave when she experienced the conditions produced by this level of atmospheric carbon.

What would this table look like? There are two possibilities to consider: either the Bangladeshi's reports are always the same whenever the greenhouse gas concentrations are the same, or they are not. But we can be sure that at at least one point on our table (and probably many), the reports for adjacent concentrations are different. This is because they are different at the end points: if the Bangladeshi would suffer terribly from the accumulation of many subtle effects of climate change caused by future emissions, then she must feel much worse under the concentrations at the end point than at the beginning point. That difference must show up somewhere along the spectrum. Consequently, at at least some point (and probably many points), the Bangladeshi really can feel the difference that the emissions produced by a recreational drive makes. The claim that no one will feel such differences would be false.

Of course, it is unlikely that our Bangladeshi's powers of introspection are so fine-tuned and well-calibrated that the reports she gives at any given concentration are always the same. Instead we should expect her to sometimes over- or under-report what she feels – either because she is paying imperfect attention or because she makes mistakes in how what she feels now compares to other feelings in the past. But the *frequencies* of the reports between adjacent concentrations will at least sometimes have to be different because they will have to be different at the end points. Suppose, for instance, that the Bangladeshi chooses to give her reports using a numerical scale where 0 is the number she usually assigns to the starting concentration and 10 is the number she usually assigns to the terrible suffering she experiences at the highest concentration. As our table progresses up the spectrum, the mean of the reports will have to change from near 0 to near 10. And so for at least some adjacent concentrations on the spectrum (and probably many, if not

all), the mean – or the frequency with which she uses particular numbers to report her experience – will have to be different.

Now to say that the emissions produced by a recreational drive can affect the frequency with which a person perceives her suffering to be greater is not to say that she will always feel the difference that any particular recreational drive makes. But it is to say that at least some recreational-drive-sized emissions will change the expected value of her felt quality of life for the worse. This difference should make a difference to her, and it does make a difference to our moral assessment of the consequences of the recreational drives. It should make a difference to her because it gives her a self-interested reason to prefer that those recreational drives that reduce the expected value of her felt quality of life not occur. This is because she should conclude that the level of suffering she experiences from adjacent concentrations whose mean report is different is in fact different: the facts about the frequency of her reports give her reason to believe that things actually feel worse at the higher of the two concentrations. It is just that the difference she feels is so slight that she sometimes misses it or makes mistakes in comparing it to how she remembers she felt at other concentrations.

The moral assessment of the consequences of recreational drives needs to take into account the changes in the expected value of people's lives they can make. Consider first those recreational drives that make a change in the expected value of the Bangladeshi's life. When we are asking about wrongness in the evidence-relative sense, what we are interested in is the expected value of the outcomes of our action. Since the value of the Bangladeshi's felt experience is morally relevant, changes in the expected value of her experience are relevant to the permissibility of these drives. What should we say about those recreational drives, if any, that don't cause a change in the expected value of her experience? In a fact-relative sense, their relationship to her is not a moral consideration, as they do not affect her. But in an evidence-relative sense, things are different. Since we will never be in a position to know whether our recreational drive is or is not one of those that makes a difference, the question at issue is what their expected value is. And since the alternatives are that our recreational drive either makes no difference or makes the expected value of her quality of life worse, they too have a negative expected value vis-à-vis the Bangladeshi.

We should acknowledge that the impact our recreational drive can have on the Bangladeshi, though morally significant, is still tiny. It is likely to be dwarfed by the positive impact that a recreational drive has on the driver. Recall, however, that we have been considering only one person who can be impacted by a recreational drive. In reality, there are (or will be) billions, and the analysis above applies to each one of them. Our luxury emissions lower the expected value of the

experiences of people around the globe, and will continue to do so for centuries. All those tiny effects thus add up to a significant moral cost in consequentialist terms.

Here, though, a new worry might arise. Can all those billions of tiny negative impacts really outweigh the comparatively substantial negative impact of not going on a recreational drive? Do tiny impacts on many different people really combine to exceed the moral significance of a far more noticeable impact on one person? There are several things to say in response. The first is that, strictly speaking, our position does not depend on the claim that luxury emissions' contribution to gradual climate change outweighs the benefits for the emitter. As we argued, luxury emissions can also cause large-scale weather events. Our claim is that together the expected disvalue of these large-scale and small-scale impacts is greater than the expected value of at least some luxury emissions to the emitters.

It is also important to note that this is not the same worry that motivated several of the contrary analyses we began with. These authors claimed that act-consequentialism provides no ground to object to luxury emissions because they produced *no* morally relevant harms, not that the harms they produced were very many but very small. The current worry, then, is really a worry about whether act-consequentialism gives the correct explanation of the wrongness of luxury emissions, not about whether, on its own terms, it can provide any explanation at all.

Finally, the thought that there is some kind of mistake in thinking that all morally relevant value is commensurable is hardly new, nor unique to climate change. It has been the subject of a great deal of analysis – both positive and negative. This is hardly the place to engage in detail with this complicated question. But we should note that the particular comparison at issue here is importantly different from the kind of comparisons that are normally thought to raise problems. Normally, philosophers have been concerned about whether small benefits to many people could possibly outweigh a fundamentally important benefit to a single person: for instance, some are skeptical that preventing literally any number people from suffering from a minor headache could be more valuable than preventing a single person from being tortured to death.⁹ Denying oneself the pleasures of luxury emissions is a far cry from being tortured, however. It is not just that these harms are of vastly different magnitudes. They differ in kind as well as in degree. Torture and death strike at the very core of an agent; they undermine all that could plausibly make someone an object of moral concern. Having to wear socks in one's house to be comfortable in the winter does not. To the extent that

⁹ Norcross 1997 includes an especially rich discussion.

skepticism about commensurability is compelling, it is so when things of fundamental importance are at stake. Since luxury emissions do not even approach this level of importance, and since nearly everyone will grant that preventing many small harms can be more worthwhile than preventing a single greater harm of the same kind, doubts about the kind of commensurability necessary for an act-consequentialist to object to individual luxury emissions seems relatively unproblematic.

So, if the analysis we have presented is correct, all the harms of anthropogenic climate change can indeed be disaggregated, and this provides act-consequentialists with all they need to show that we should eliminate those luxury emissions that would be sufficient to bring us down to the level of emissions that is impersonally best. Before concluding, however, we should acknowledge that some may think that this does not go far enough. For it may be that there are what intuitively seem like terribly self-indulgent – and hence morally objectionable – emissions whose elimination would not be involved in getting down to the optimal level in the least costly way. There could be someone, for instance, whose sense of self-worth is tied to using their private jet. Given how central burning fossil fuels is to them, it could well be that their emissions contribute enough to their lives so as to offset the expected harm they cause, in which case an act-consequentialist may have to acknowledge that the emissions themselves are not objectionable. But to some, this sounds obnoxious: how could jet-setting emissions possibly be justified when we know that they contribute to harms that will befall innocent Bangladeshis who are struggling to survive?

At issue here is the question of whether it can be wrong (in the evidence–relative sense) to impose a risk on others even when doing so maximizes expected value. Act-consequentialists say no, while others will disagree. This is not a debate it makes sense to join here, however, and not only for the usual ‘constraints of space’ reason. For the point of our arguments has not been to establish the superiority of act-consequentialist analysis of climate ethics vis-à-vis its competitors. Instead, it has been to show why one influential source of its supposed inferiority is an illusion, and why even those who care only about the effects of their actions from an impersonal perspective ought therefore to care about their personal emissions. And both these points would still stand even if act-consequentialism doesn’t provide an exhaustive account of what makes the difference between those choices that are permissible and those that aren’t.

§6. Conclusion

The arguments we have developed in this article have often been complex, abstract, and counterintuitive. Yet our main conclusions are neither complex, nor abstract, nor counterintuitive. In pursuing our own perceived self-interest, the

rich are doing terrible harm to other people and to the natural systems on which humanity's future depends. We should act, and act now, to reduce the damage we are doing. When and if we can, we should act collectively, through governments, political parties and NGOs. In the meantime, each of us has an individual obligation to do what we can to stop harming others, including by refraining from, or perhaps by purchasing offsets against, our own individual luxury carbon emissions. Once people understand the science of climate change, they immediately and intuitively recognize that they have reason to comply with these straightforward and demanding moral obligations. Yet many have been led away from recognizing what the morality of climate change genuinely requires of them by a set of influential, seemingly compelling, but ultimately fallacious arguments. In climate ethics, it seems, a little knowledge can be a dangerous thing.

Works Cited

Arntzenius, Frank and McCarthy, David. 1997. 'Self-torture and group beneficence'. *Erkenntnis* **47**(1): 129-144.

Gardiner, Steven. 2011. *A Perfect Moral Storm: The Ethical Tragedy of Climate Change*. New York: Oxford University Press

Jamieson, Dale. 2007. 'When utilitarians should be virtue theorists'. *Utilitas* **19**(2): 1-24.

Jamieson, Dale. 2014. *Reason in a Dark Time: Why the Struggle Against Climate Change Failed -- and What It Means for Our Future*. New York: Oxford University Press.

Johnson, Baylor. 2003. 'Ethical obligations in a tragedy of the commons'. *Environmental Values* **12**(3): 271-287.

Hiller, Avram. 2011. 'Climate change and individual responsibility'. *The Monist* **94**(3): 349-368.

Kagan, Shelly. 2011. 'Do I make a difference?' *Philosophy and Public Affairs* **39**(2): 105-141.

Kautz, Richard. 2011. *Chaos: The Science of Predictable Random Motion*. New York: Oxford University Press.

Matheny, Gaverick. 2002. 'Expected Utility, Contributory Causation, and Vegetarianism'. *Journal of Applied Philosophy* **19**(3): 293-297.

Norcross, Alastair. 1997. 'Comparing harms: headaches and human lives'. *Philosophy & Public Affairs* **26**(2): 135-167.

Parfit, Derek. 2011. *On What Matters*, Volume 1. New York: Oxford University Press.

Sandberg, Joachim. 2011. 'My emissions make no difference: climate change and the argument from inconsequentialism' *Environmental Ethics* **33**(2) 229-248.

Sandler, Ronald. 2010. 'Ethical theory and the problem of inconsequentialism: why environmental ethicists should be virtue-oriented ethicists'. *Journal of Agricultural and Environmental Ethics* **23**(1-2): 167-183.

Singer, Peter. 1980. 'Utilitarianism and vegetarianism'. *Philosophy & Public Affairs* **9**(4): 325-337.

Sinnott-Armstrong, Walter. 2005. 'It's not *my* fault: global warming and individual moral obligations'. In Walter Sinnott-Armstrong and Richard Howarth (eds), *Perspectives on Climate Change: Science, Economics, Politics, Ethics* (Oxford: Elsevier): 285-307.

Shue, Henry. 1993. 'Luxury Emissions and Subsistence Emissions'. *Law & Policy* **15**(1): 39-60.