

A “Famine, Affluence, and Morality” for Climate Change?

(PENULTIMATE DRAFT – PLEASE CITE FINAL VERSION)

1. Introduction

Peter Singer wrote his landmark paper “Famine, Affluence, and Morality” (hereafter *FAM*) in 1971 in response to the refugee crisis facing Bengal at the time, concluding that affluent people are obligated to contribute a large portion of their wealth to help those in need throughout the world.¹ The core of *FAM* is a now-famous example: it would be wrong not to save a drowning child even if it means ruining one’s clothes, and Singer argues that our obligation to people in need in developing countries is no different than our obligation to the drowning child nearby. This example is easy to grasp, and although Singer’s argument has certainly not been universally persuasive, it is hard to doubt that it, rephrased in books and in major newspapers and magazines by Singer and others, has had significant effects on many people.

Presently, one of the most prominent global issues is the potential that due to anthropogenic global climate change (AGCC), many people (and non-humans) will be severely harmed. In this paper I assess the prospects for an argument analogous to Singer’s to apply to AGCC, and after providing some background in §2 and §3, I argue in §4 that the prospects are good. In particular, there is a striking commonality between our current situation and that discussed in *FAM*: affluent people can sacrifice certain things which are not morally significant and thereby prevent or avoid the occurrence of some suffering of others far away. I also provide a thought experiment analogous to Singer’s drowning child case which should make our responsibilities regarding AGCC very salient. In §5 I discuss some initial objections similar to those discussed by Singer. In §6 I argue that the argument in §4 shows that most plausible forms of deontology are very stringent in an era of AGCC. In §7, however, the paper takes a potentially

surprising turn: I show that the argument in §4 ought to make little difference for consequentialists, for a very specific reason: Singer's original argument in FAM still applies today, and it trumps, by considerable measure, the argument in §4. However, in the concluding sections of the paper, I discuss some reasons to be cautious about this conclusion.

2. Individual responsibility for climate change

Even among philosophers who are deeply concerned about potential harms that will likely occur because of AGCC, there is much skepticism about the role that individual actions have.² The main paradigm in climate ethics is to view it as a collective action problem, where nations as a whole as well as intergovernmental organizations are the ones truly responsible for dealing with AGCC. Given the highly diffuse causal structure of AGCC, where individual actions which emit greenhouse gases (GHGs) are miniscule in comparison to the total amount of GHGs emitted, it seems plausible to believe that individual actions such as Sunday drives are, on their own, completely negligible. For instance, James Garvey writes:

There are nearly 7 billion people on the planet. Together we emit 28.4 gigatons of CO₂ each year... By comparison, I am responsible for about 4 tons of CO₂ each year. Am I to see myself as responsible for $4/28.4 * 1,000,000,000$ or 0.000000000141% of the harm done to our planet this year? Should I try to do better and aim for 4/28,399,999,999? ... [T]he harm I do is so impossibly teeny that it can't figure into a real motivation for green action.³

The goal of the first half of this paper is to show that the answers to Garvey's questions are a qualified "yes." Even though the proportion of harm an individual does is miniscule, given the scale of harm that may occur due to AGCC, we can use utility calculations to show that the

expected harm is not insignificant. This helps provide an illustrative example which, I believe, can lead to some fairly strong motivation for green action.

Philosopher John Nolt has estimated that the “average American” is responsible, through lifetime GHG emissions, for the severe suffering or death of one to two people, typically in the developing world.⁴ This shows that it is still important to view climate ethics through the lens of individual responsibility. However, Nolt leaves it unclear exactly what the expected harm is from daily individual choices, such as a Sunday drive in a gas-guzzling car. I argue elsewhere that if Nolt’s estimate is correct, then emitting 1 kg of CO₂-equivalent causes an expected .43 person-hours of serious harm.⁵ And if that is correct, then the expected marginal increase in overall harm caused by the emissions from a 25-mile Sunday drive (14.1 kg of CO₂-equivalent) is approximately that of six person-hours – or, in other words, ruining someone’s afternoon. As I put it, using non-technical moral vocabulary, that is not a very nice thing to do.

Several points here require clarification. First, although AGCC will occur regardless of any one individual’s emissions, the situation is not one of causal overdetermination. It is unclear how many people will be harmed by AGCC, and certainly many individuals will be harmed regardless of the precise amount of GHGs emitted. But what matters is *minimizing* harm from AGCC, and it is safe to assume that harms due to AGCC will increase as GHGs in the atmosphere increase. So even if the existence of climate-caused harms (in general) is already overdetermined, the precise amount and nature of those harms is not, and thus our actions can make a difference.

One might argue that this relies on an assumption that there will not be completely catastrophic harm due to AGCC no matter what any individual does. If there will already be catastrophic harm because a certain GHG threshold has passed, then what difference does it

make if there are additional GHG emissions beyond that threshold? There are three responses to this concern. First, there is still a range of catastrophic harms, and it makes sense to say that some states of catastrophic harm are worse than others. Second, even if there is destined to be catastrophic harm, the more we delay that harm from occurring, the better the overall outcome will be. Just because an outcome will be horrible doesn't mean that we should not bother to take reasonable steps that might delay it. Third, it is by no means certain that there will be catastrophic harm. So even if in fact my lifetime's GHG emissions will turn out to make no difference in how much harm occurs, I cannot say in advance that my GHG emissions will not increase the likelihood of harm, and that is what matters for decision-making. (I will discuss catastrophic harm further in §9.)

Further, to say that one's own emissions are not responsible for any of the harm from AGCC makes a mystery out of what does cause the harm. Since my GHG emissions are no different in their nature than anyone else's emissions, and since we know that total GHG emissions are causing harm, then if my own emissions are not causing harm, then no one's emissions are. But there is no way to make that conclusion consistent with the fact that in general, emissions are causing harm.

The only way a similar objection may be applicable would be if my choosing not to emit GHGs were to make no difference whatsoever in the total amount of worldwide GHG emissions.⁶ I discuss this issue in more detail elsewhere; I grant if it is true that my emissions make no difference whatsoever, then, from a consequentialist perspective at least, then it would be permissible to emit GHGs. However, although it is possible that total GHG emissions across time will be the same regardless of what any individual does, this is not certain, and so in our decisions we cannot rely on the inevitability of GHG emissions to excuse our own emissions.

Benjamin Hale argues that delaying GHG emissions 10 years simply means that the same harms will occur, just delayed 10 years, and recalibration will occur so that the total amount of time with elevated atmospheric levels will be the same,⁷ but it is not true that the total amount of harm would be the same without such a delay. First of all, earlier GHG emissions will cause faster climatological changes than later ones; the problem with AGCC is that it is happening so quickly that humans (and other species) are unable to adapt to the changes. Furthermore, because of carbon cycle feedbacks, the sooner CO₂ is emitted, there may be exponentially more total GHGs emitted.⁸ Thus this objection is not a strong one. Lastly, again, what is at issue is the expected disutility of additional GHG emissions. Unless we are fully confident that our emissions will make no difference whatsoever, we should still take steps to limit our emissions.

Another clarification of my argument is that it does not depend upon the untenable view that, to determine individual responsibility for harms caused by group behavior, we can simply divide the total harm by its parts.⁹ Although I have calculated the amount of harm of an individual action by dividing the total expected harm by the proportion of the total GHGs which are emitted by the action, this figure serves as a proxy for what is important for my purposes, which is the expected marginal increase in overall harm due to the action. Indeed, it is not true that the average amount of harm per unit of GHG emitted *always* is equal to the marginal increase in harm per unit of new emissions, given that the relation between GHGs emitted and harm caused is not fully linear. However, in calculating expected marginal increase in harm, the average harm will provide a reasonable estimate. Again, if we expected catastrophic harm regardless of what any individual does, then the marginal effects of each action will be small and the averaging method would not provide the correct answer. But since there is a great deal of uncertainty of exactly how much harm will be caused, each additional unit of GHG emitted can

have an expected increase in disutility approximately equal to the average. In general, even if there are thresholds in amounts of harm caused, given that we are unaware of where these thresholds are, the relation between GHGs and expected harm will be close to linear.¹⁰

One might object to the idea that one can divide the overall harms into units of ruined afternoons. For ruining one life is not the equivalent of ruining 100,000 people's afternoons, even if one life is composed of 25,000 days and an afternoon is taken to be a quarter of a day. On the one hand, this objection misunderstands my use of the notion of a ruined afternoon; it is simply intended to represent the amount of harm equivalent to 1/100,000th (approximately) of a life. But the objection has purchase since I intend the analogy to be intuitively gripping and thus it cannot ultimately depend upon a somewhat technical notion of a *person-afternoon of harm* (which may be simply defined as the equivalent of ¼ of a person-day of harm, where a *person-day of harm* is defined as the 1/100,000th of the harm of the loss of one life).

There are two responses to this concern. First, it should be pointed out that the harm caused by a ruined afternoon is not necessarily intuitively much different than 1/100,000th of the harm of a ruined life. Second, a defender of the spirit of the view advocated in this paper may simply grant the point, and claim that Sunday drives have an expected disutility of 1/100,000th of a ruined life, which is the same as having a 1 in 100,000 chance of ruining someone's life. In other words, there is a tiny chance that, due to the excess GHGs in the atmosphere from one person's Sunday drive, it will make a difference in the climate in such a way that an additional person's life will be ruined.¹¹ And it is also not a very nice thing to do to perform an action which has a 1 in 100,000 risk of ruining someone else's life. I shall return to this point below.

3. Famine, affluence, and morality

Here is one way to formalize Singer's original FAM argument:

(FAM1) “Suffering and death from lack of food, shelter, and medical care is bad.”¹²

(FAM2) “If it is in [one’s] power to prevent something bad from happening without sacrificing anything of comparable moral importance, [one] ought to do it.”¹³

(FAM3) By giving away one’s excess money to a relief organization, one can prevent something bad from happening.¹⁴

(FAM4) By giving away one’s excess money to a famine relief organization, one will not lose anything of comparable moral significance (to the bad prevented in FAM3).¹⁵

Thus: **(FAM5)** One ought to give away one’s excess money to a famine relief organization.

Singer supports premise (FAM2) by giving his example of the drowning child. Since we believe that one ought to save a drowning child even if it means ruining one’s clothes, presumably that would support the more general premise of (FAM2). If one denies that we have obligations to help those in need far away, one would either need to give a principled reason why it is permissible to not help those far away even though it is impermissible to not help the drowning child – and as Singer persuasively argues, distance should make no moral difference – or one must accept the seemingly abhorrent claim that one has no obligation to save the drowning child. This dilemma makes opposition to Singer’s argument difficult.

There are two aspects of this argument worth noting presently. First, the notion of “excess” money makes (FAM4) trivially true – “excess” could simply be defined as the amount by which one doesn’t cause a loss of comparable significance. But, non-trivially, Singer claims

that giving away one's money to the point of marginal utility will involve lowering one's standards virtually to the point of living like a refugee.

Second, one might argue that the drowning child case, although it lends some intuitive support to something like (FAM2), does not support (FAM2) to its letter, since (FAM2) is a truly radical premise requiring those who have excess goods to donate the entirety of the excess to help others. However, Singer mentions that a weaker form of (FAM2) could still have an enormous effect on our moral outlook. If we omit the word "comparable" from it, the resulting weaker argument would be:

(W-FAM1) "Suffering and death from lack of food, shelter, and medical care is bad."

(W-FAM2) "If it is in [one's] power to prevent something bad from happening without sacrificing anything of moral importance, [one] ought to do it."

(W-FAM3) By giving away a significant amount of money to a relief organization, a typical individual in a wealthy country can prevent something bad from happening.

(W-FAM4) By giving away a significant amount of money to a famine relief organization, a typical individual in a wealthy country will not lose anything of moral importance.

Thus: **(W-FAM5)** Typical individuals in wealthy countries ought to give away a significant amount of money to a famine relief organization.

The drowning child example does seem to provide strong intuitive support for (W-FAM2) even if it does not provide strong enough support for (FAM2). Although Singer himself, as a

utilitarian, accepts the stronger argument, even the weaker argument entails that wealthy individuals ought to make significant changes in their lifestyles.

4. A “Famine, Affluence, and Morality” for climate change (FAMCC)

Now, here’s an argument I shall call *FAMCC*:

(FAMCC1) Someone’s afternoon being ruined is a bad thing.

(FAMCC2) If it is in one’s power to prevent or avoid something bad from happening without sacrificing anything of moral significance, one ought to do (i.e., prevent or avoid) it.

(FAMCC3) Not going on a Sunday drive would avoid the equivalent of someone’s afternoon being ruined.

(FAMCC4) Not going on a Sunday drive is not a morally significant loss.

Thus: **(FAMCC5)** One ought not go for a Sunday drive.

In case one doubts FAMCC, consider the following illustrative example. It would obviously be wrong for me plan to go for a Sunday drive with someone where that person rearranges her afternoon to join me, but then leave her stranded somewhere bad for the whole afternoon while enjoying the drive. Even if I have left the person with a note of assurance saying that I will pick her up six hours later, it would be a lousy thing to do. However, if the data underlying (FAMCC3) are correct, then that is the approximate effect of a Sunday drive. As with the drowning child example in FAM, anyone who believes that it is wrong to go for a drive while leaving someone stranded ought to provide some reason why the case of AGCC is different from that case in a morally significant way. I shall explore several responses in the next section.

However, if the foregoing is correct, then our entire moral orientation must be changed.¹⁶ In an era of AGCC, even ordinary activities in a carbon-based economy are morally wrong.

There are many choices one can make about one's actions. To take just one example, flying 2500 miles round trip to an academic conference emits a significant amount of GHGs.¹⁷ Although some good is caused by the conference, one should recognize at the same time that the GHG emissions from travel will cause non-negligible harm. At the very least, those who participate in conferences have some obligation to consider whether the benefits accrued from such activities outweigh the harms.

If FAMCC is sound, we need to give up on a great deal of our carbon-based system – individuals should not contribute to it, even for most seemingly small actions. The main differences between FAMCC and the argument in FAM are the following. First, I have included a disjunction in the antecedent of the second premise; this should be uncontroversial. Second, the bad in question in FAMCC is much less significant than the bad involved in FAM. That is because the effects of an individual act of emitting is not as significant as the example in FAM. Still, this does not mean that FAMCC is invalid, although I shall discuss this issue in greater detail in §7. Third, I have mirrored the weaker argument rather than the stronger, since it is all that is needed. (I shall discuss the insertion of the word “comparable” in greater detail in §6.) The most important difference between FAMCC and the argument in FAM is that while Singer does not appeal to anyone's causal responsibility for the suffering of those in need, in the case of AGCC, affluent people are in fact causally responsible for harms, albeit in a diffuse way. Since one objection to FAM is that people do not have a strong duty to prevent harm for which they are not causally responsible, this difference should make my intuitive case even *stronger* than that in FAM. So the application of Singer's argument to climate change will in one way be more similar to Thomas Pogge's view of third-world poverty (according to which we are in fact causing the

suffering of others far away)¹⁸ than to the view in FAM, but it is structurally similar to Singer's in that it uses a vivid example and draws its conclusion from it.

In emitting GHGs, we are not directly causing harm; we are merely diffusely contributing to a harm. Thus one might claim that there is a significant difference between leaving someone stranded on the side of the road for an afternoon and emitting GHGs which enter the atmosphere and in a very complex way lead to climatological changes which cause harm. However, as above, what underlies my argument is a claim of expected harm, using a standard mechanism of a (dis-)utility calculation. It could be that the possible world which is identical to ours (up to now) except in that I go for a Sunday drive would not be any better than the actual world in which I do not, but that would just be a matter of luck. On average, each increment of GHGs emitted can be expected to cause an increment of harm. To deny the use of utility calculation in this case would be unprincipled; there is no reason why it should be inapplicable in this case while it is perfectly applicable in other cases. One should not play a Russian roulette-type game, even with very good odds and even if those who are not killed in the game are given a slight reward, because of the expected disutility. The case of AGCC is more diffuse than Russian roulette, but the difference is merely a matter of scale.

Another form of argument uses the point made above that Sunday drives could also be understood as having a slight risk of causing serious harm:

(FAMCC-R1) Someone's life being ruined is a bad thing.

(FAMCC-R2) If it is in one's power to avoid a non-negligible risk of something bad from happening without sacrificing anything of moral significance, one ought to avoid the risk.

(FAMCC-R3) Not going on a Sunday drive would avoid the equivalent of a 1 in 100,000 risk of someone's life being ruined.

(FAMCC-R4) A 1 in 100,000 risk of someone's life being ruined is not a negligible risk.

(FAMCC-R5) Not going on a Sunday drive is not a morally significant loss.

Thus: **(FAMCC-R6)** One ought not go for a Sunday drive.

In intuitive support of this argument, we might consider an example where, if one goes for a Sunday drive, a random number between 1 and 100,000 will be generated, and if it matches a pre-set number, someone's life will be ruined. This, certainly, should at the very least raise significant doubts about the morality of Sunday drives.¹⁹

5. Initial objections and responses

In "Famine, Affluence, and Morality", Singer anticipates and responds to a number of objections which may be leveled against his view. As it turns out, a number of these objections apply to FAMCC.

First, one might claim that we have a limited understanding of the long-term effects of our actions. Famously, it is difficult to quantify harms, and it is unclear what, exactly, a person-hour of harm is. Furthermore, there is a tremendous amount of uncertainty concerning both the climatological effects of greenhouse gas emissions and the amount of harm that any AGCC will in turn cause.²⁰ However, given the present state of climate science, it is indisputable that there will be significant harms because of AGCC. So if one doubts the estimate above, then one still ought to determine more carefully the effects of their emissions and not to use uncertainty as an excuse to deny individual responsibility. But since our current best estimates do still show that there is *expected* harm from AGCC, the fact that we cannot determine the precise effects is no

objection against FAMCC. Even if one is unsure whether a grenade is live, one should not throw it in the direction of a group of people.

It might also be objected that many other people are causing AGCC, too. People throughout developing countries, especially the US, will continue to emit GHGs at high rates, and China and India are rapidly increasing their GHG emissions. Why should *I* be held morally responsible? Although this notion is very common, it is highly morally dubious: would my leaving someone stranded on the side of the road be any less worse if others left other people stranded in other places? “One has only to ask this question to see the absurdity of the view that numbers lessen obligation.”²¹

One might also claim that AGCC is best seen as a problem for society in general and it should be the job of governments and not individuals to respond to it.²² It is true that AGCC is, to a large extent, a collective action problem, and I do not deny that we do need international agreements to limit GHG emissions. However, there is too little emphasis on AGCC as a matter of individual responsibility. As Singer argues in FAM, governments are more likely to promote a cause when they perceive that citizens are deeply interested in it.²³ Furthermore, we should not underestimate the difference that individuals can make in lessening the harms due to AGCC. It is easy to be overwhelmed by the enormity of the atmosphere into thinking that one’s own emissions make no difference, but if the highly plausible empirical assumptions underlying Nolt’s work are correct, then one’s own emissions do make some morally significant difference.

Last, might object that by setting a moral standard too high, people will end up ignoring all their moral responsibilities.²⁴ This is a genuine concern since many people already dismiss what they believe are overstatements and fear mongering about AGCC. But this does not undermine the validity of the argument; it merely shows that we need to be careful in telling

people about it. This too is a problem for aid to developing countries, where people disbelieve that they will do any good by giving money. But that doesn't mean we should stop trying to convince people to do it. Perhaps FAMCC, along with its vivid example, can be persuasive.

6. Upshot for deontology: it is *highly* demanding in an era of AGCC

The conclusion of FAMCC is that simple acts such as going for Sunday drives are wrong. Consequentialism is held by some to be an overly demanding normative theory,²⁵ but FAMCC also applies to any non-consequentialist view which accepts a "Do no easily-avoidable harm" principle, which underlies the second disjunct of the antecedent of (FAMCC2). I do not wish to argue the point here, but it seems that the most plausible forms of deontology ought to accept such a principle.

It is worth noting that whether either type of normative theory is overly demanding is based upon contingent facts. For example, there is a possible world in which virtually every action which causes the agent 10 units of pleasure also causes someone else 8 units of pain. It may still be appropriate to perform those actions under consequentialism (unless there was a way to have a greater margin of pleasure over pain), but likely not under deontology. Or, it could be that in some possible world, we are hooked into a Parfit-style torture device, where every action one may take, including ones which are pleasurable to the agent, harms someone. In these possible worlds, agent-neutral consequentialism is likely to be less demanding than deontology. What FAMCC may suggest is that the actual world is a world quite like that one – nearly all actions in a carbon-based economy will have some amount of expected harm.

In support of this, recall that in my presentation of FAMCC, I used Singer's "weaker" principle, which omits the word "comparable" from premise two. However, Singer and other utilitarians believe that the correct moral principle does include the term "comparable." What

happens when one includes the term “comparable” in (FAMCC2)? Arguably, consequentialists could reject the conclusion of FAMCC on the grounds that the happiness that one causes oneself in driving does outweigh the harm caused. And consequentialists could also appeal to the fact that climate change, though it harms some individuals, will likely also help others. So while including the word “comparable” in premise two of the original FAM argument makes it more stringent, in FAMCC it makes it *less* so. In FAMCC, if an action causes more pleasure than the harm it causes, then it would be permitted on the basis of the premise which includes “comparable,” but would not be permitted without it, since it might be that one could give up on the pleasure without giving up on anything of comparable significance. It is perhaps an odd fact about FAM-type arguments that for small-scale actions, the so-called weaker argument is stronger than the so-called stronger argument. But that does nothing to undermine either of the arguments in FAM or FAMCC.

However, deontological views won’t allow an individual just to do an overall utility calculation in performing an action. On a deontological view which accepts a “do no easily-avoidable harm” principle, ordinary actions which contribute to our GHG-based culture are wrong because they cause easily-avoidable harm.²⁶ Many ordinary actions do so, to some non-negligible extent. Thus these ordinary actions are wrong. This is by no means an argument against non-consequentialism – demandingness is a dubious reason to reject a normative theory. But those persuaded by the overdemandingness argument against consequentialism into accepting a form of deontology may wish to reconsider their commitment to deontology – at least to one with a “do no easily-avoidable harm” principle. In an era of AGCC, the most plausible forms of deontology are extremely stringent.

7. Problems for the consequentialist’s FAMCC

Is deontology really more stringent than consequentialism? In this section, I show that there is a significant problem with the line of argument in §4, and with FAMCC if it is understood from within a maximizing act-consequentialist framework. FAMCC is problematic in that it has competition from another argument: importantly, the original FAM argument still applies today. To see this, let me first note that there is something odd about the original structure of Singer's FAM argument. Singer's argument structure, both in the original article and in recent work, is intended to appeal to a general audience and does not assume utilitarianism. Arising from this is a formal problem which can be seen by considering the following. Right now, giving one's excess money to OXFAM (let us assume) will help prevent harm without causing comparable harm. Thus, as a substitution instance of premises (FAM3) and (FAM4) in the original argument, we can conclude in a more specified (FAM5) that one ought to give an amount of money equal to that of one's excess above the point of marginal utility to OXFAM. But there are numerous other organizations which can also be substitution instances within (FAM3) and (FAM4). Thus one ought to give that amount to those organizations. But one cannot give that amount of money to more than one organization. Thus, assuming an ought-implies-can principle, one cannot be obliged to give that amount of money to more than one organization. Thus Singer's original argument, at least as it is stated in §3, leads to an absurdity.

What Singer really intends is that one ought to give one's excess money to the organization which *best* prevents harm while causing the least amount of harm. This is why Singer is careful in recent work to discuss specific organizations to which one can donate. And Toby Ord has recently stressed the importance of finding efficient ways to help others, since some charitable organizations are thousands of times more effective in helping individuals than other organizations.²⁷ A better form of the FAM argument for a maximizing consequentialist is:²⁸

(FAM-MC1) If one can prevent the most bad by doing x , one ought to do x .

(FAM-MC2) One can prevent the most bad by doing n .

Thus, **(FAM-MC3)** One ought to do n .

Premise (FAM-MC2) is not clearly supported by the drowning child example, since that example only appeals to intuitions about saving an individual's life when the sacrifice is not significant, and does not indicate that one ought to maximize the prevention of harm. But it is more true to FAM's root in a maximizing form of consequentialism. The goal is simply to solve for n .

How is this relevant to FAMCC? If (FAMCC2) is replaced by version which supports a maximizing form of consequentialism, then the argument would still conclude that one ought not go for a Sunday drive. But one could also conclude from a premise of maximizing consequentialism that one ought to be working full-time to prevent the most harm, say, by working full-time to give one's money to OXFAM. For a true maximizing consequentialist, it is worth asking the following question: how does the good one achieves by refraining from Sunday drives compare to the good one achieves by donating one's time and money to organizations that help those presently in the developing world?

Recall, from above, that Nolt argues that the average American is causally responsible for the suffering and/or death of one to two people. That sounds quite bad. But what about the original argument in FAM, as applied to today? In Ch. 6 of *The Life You Can Save*, Singer discusses several analyses of how much it costs to save one life, which all are between \$200 and \$2000.²⁹ Taking Singer's highest estimate of the cost of saving a life and Nolt's highest estimate of how much harm the average American's GHG emissions cause, one who fails to give \$4000 dollars is allowing as much harm as one who emits an entire American's life's worth of GHGs. But for the average American, the significance of the loss of \$4000 is next to nothing when

compared to the loss of a lifetime's worth of GHG-emitting activities. In other words, the amount that one can help others by giving money away trumps, by significant measure, the amount of harm one causes by emitting GHGs. Although there may be some dispute regarding the precise amount of harm caused by AGCC, as long as the estimates are accurate to within an order of magnitude, then helping others immediately is going to have more of a beneficial effect than not causing harm due to one's GHG emissions.

At environmental ethics conferences, there is often an awkward moment where a speaker (impolitely?) brings up the issue, as I did earlier, of the carbon emissions from participants' transportation to the conference, which typically causes participants to feel guilty. But if the empirical claims underlying the present argument are correct, the harm due to GHG emissions from an airplane flight is miniscule compared to the amount of good one could do with the money used for the airplane ticket.³⁰ And thus if the argument in the preceding paragraph is correct, it is very strange for a consequentialist to worry about GHG emissions when there is presently rampant avoidable severe present harm from famine, disease, etc.

I am not claiming that the considerations in this section show that we do not have an obligation to limit our GHG emissions; strictly speaking, the idea that one can help others by donating money does not directly compete against the fact that one can avoid harming others by not emitting GHGs (except insofar as the money that one earns in the first place that one is giving away come from activities which emit GHGs). However, most consequentialists do not, in their own lives, act in such a way as to maximize overall good. Singer himself, for instance, does not donate his own money to organizations so that he gives to the point of marginal utility.³¹ Such "practically-non-maximizing" maximizing consequentialists should likely give less priority in reducing their GHG emissions relative to other sacrifices they already should by their own

lights be making but are not. Or those who are scalar consequentialists³² should consider donating more money to aid organizations prior to limiting their own GHG emissions.

8. The Skeptical Environmentalist?

Bjørn Lomborg has argued that we should not expend much effort mitigating harms due to climate change, and instead focus our attention on other global issues.³³ Lomborg has been vilified by the environmental community for making such claims;³⁴ however, if the present argument is correct, we should not be so quick to dismiss Lomborg's views. For example, Lomborg argues that it is much more efficient to put money towards preventing malaria than towards preventing climate change.³⁵ There is a highly plausible case to be made for that conclusion.

On the one hand, I agree that Lomborg is asking precisely the right question: how should we best do the most good? There seems to be a kneejerk reaction about climate change among some people that makes them committed to climate action even when they are not moved by appeals to give to development non-profits. If one rejects a doing/allowing distinction – and most consequentialists do – then Lomborg's general project is quite reasonable.

However, there are several ways to distinguish my view from Lomborg's. First, it should be noted that it is not true that if the marginal effects of my own giving, say, \$10,000 to help others in need now is much more beneficial than my spending \$10,000 to significantly lowering my GHG emissions, then the effects of the US government giving, say, \$10,000 per capita in global aid will be less beneficial than if it spends at a rate of \$10,000 per capita in lowering GHG emissions. This is a complicated empirical issue, but it might be the case that there are decreasing marginal benefits the more money is given to assist those presently in need compared

to the marginal benefits of spending more money to mitigate climate change. There are good reasons for believing that this is the case: climate change is an enormous problem which will affect many generations of people and which may require altering our entire civilization, whereas, if economists such as Jeffrey Sachs are to be believed, the problem of current global poverty could in principle be solved without too much expense.³⁶ Thus, even if my \$10,000 does more good if put to malaria prevention than to climate mitigation, \$10 trillion may do more good when put to the latter than for all issues of current global poverty.

There are several other important issues connecting the ethics and economics of climate change and international development which go well beyond the scope of this paper, but I'd like to briefly mention two here. First, one reason why Lomborg concludes that our money is best spent helping those presently in need is that he uses a 5.5% (and sometimes 6%) discount rate. The discount rate is the annual rate at which future goods are devalued relative to present goods. At such a high discount rate, goods in a mere 20 years will only be valued at approximately 1/3 the rate as present goods. Thus it should then be no surprise that Lomborg stresses that we should try to solve current problems rather than put efforts towards long-term climate mitigation. However, having a 5.5% discount rate is absolutely ethically indefensible.³⁷ Nicholas Stern's Review,³⁸ officially sanctioned by the British government, uses a 1.2% discount rate, and Stern argues that we do need to take strenuous steps to lower carbon emissions, despite the costs of doing so.

Additionally, many efforts to provide present aid to the developing world also help with regard to stopping future climate harms. For example, the better wells that are built now in developing countries, the more likely future individuals in those areas will be to survive AGCC-

induced droughts. Thus the contrast in this paper between the result of FAM and the result of FAMCC need not be to separate the two arguments. One can incorporate FAMCC within FAM.

If these data are correct, then AGCC should indeed be seen primarily as a collective action problem. If (A) the marginal difference that most individuals in developed countries can make in avoiding climate harms is trivial compared to the marginal difference that such individuals can make in helping others by giving monetary aid and (B) the doing/allowing distinction is not ethically significant, then it follows that individuals who are not already maximizing consequentialists in practice should not be overly concerned about climate harms they cause. However, if we assume that action at an international scale can avoid significant harms, then we should indeed be highly concerned with solving the collective action problem.

9. Additional considerations

My argument has been the following. There is an updated form of Singer's FAM argument that applies to climate change. This argument, FAMCC, makes deontological views which accept a plausible "do no easily-avoidable harm" principle quite stringent. This is because even short drives cause a non-negligible amount of easily-avoidable harm. Consequentialists, on the other hand, ought to be concerned about AGCC, but that concern should not alter in a significant way their already deep commitment to helping others. An explanation of why climate change has become a hotly debated ethical issue may simply be that most people do accept a doing/allowing distinction, and in today's world, it is not hard to be reminded of one's actions which emit GHGs and of their effects. It takes much more of an argument, such as that given by Pogge, to show how we are contributing to suffering far away independent of climate change. Our harms from doings are much more salient than harms we could be preventing. Thus consequentialists who reject a doing-allowing distinction need to be sure that they are not unduly

affected by the normal human importance we attach to doings and thus pay too much attention to climate harms and too little attention to harms that they do not cause.

However, there are several further reasons why AGCC may deserve more attention than the argument in §7 might lead one to believe. One important aspect of AGCC is that it is more than just humans who will be harmed; the entire planet will suffer the consequences of AGCC. Lomborg gives virtually no consideration to non-human harms, and John Nolt's calculation, which forms the basis of the data used in this paper, only involves harms to humans. (Nolt himself is indeed concerned with non-human harms, but avoids discussion of them for ease of calculation.) So perhaps consequentialists' concern about individual contribution to AGCC may be warranted if one factors in the harm caused to non-human animals. However, it is unclear whether an individualist conception of harms to non-human animals will suffice for this purpose. Surely the non-human world will be substantially different if there is significant AGCC. But perhaps there will not be negative effects in terms of total non-human utility if all that is taken into account are individual sentient animals. Perhaps a happy, herbivorous, and fast-multiplying form of rodent will come to dominate a warmer world, squeezing out other species but maximizing total sentient non-human goodness. There is in fact good empirical evidence that in times of rapid global warming, small mammals do increase in number of individuals even if many animal species go extinct.³⁹ So it is unclear whether an appeal to harms to non-human animals can be used as a reason for a classical utilitarian to take more interest in climate change.

Of course there are other pluralistic and eco-centric forms of consequentialism which can appeal to species extinction and ecosystemic collapse as being important additional factors in consequentialist calculations.⁴⁰ If such views about valuable states of affairs are true, then it is possible that for a maximizing consequentialist, FAM will not trump a revised form of FAMCC.

Since I myself accept a form of consequentialism which accepts that ecosystems as a whole have value,⁴¹ I will leave this possibility open for further work. And for non-consequentialist theories which prohibit harms to non-humans, a revised form of FAMCC will become even more stringent than what I suggest in §6.

A final important issue must be addressed. As I note above, it is a possibility that due to AGCC, there will be catastrophic global harm. Even though the risk of this occurring is quite small, the harm due to it may be so great that the expected disutility of GHG emissions becomes quite significant.⁴² This issue involves extremely complex empirical and ethical considerations. First, it is extremely difficult to calculate the risks of such harms.⁴³ Second, it raises the deep issues made famous by Parfit (1984) about whether what matters in performing a utility calculation is average wellbeing or total wellbeing. If the latter is the case, then the harm of human extinction would be hugely significant. Like Broome (2010), I am unsure exactly how to factor in the risk of catastrophic harm into the calculation of expected disutility of actions.

There are several points to be made in response, however. Nolt himself uses an estimate of two billion people harmed by AGCC. This is already very large amount of harm. Even if AGCC leads to human extinction over the next millennium, the amount of people harmed will not exceed Nolt's estimate by a vast amount.

However, one might still claim, if one believes that what matters is total human happiness, that the lost good from the absence of humans on the planet will be overwhelming.⁴⁴ There are two responses to this. First, the extinction of humans still leaves the possibility that another species will eventually evolve with many of the same capacities (and goods) as humans. Second, it is quite possible that humans will go extinct for other reasons than AGCC. There may be a nuclear catastrophe, global pandemic, asteroid impact, or global harms from chemical or

biological warfare, to name a few.⁴⁵ Thus the net expected utility loss from AGCC even in the case of human extinction is not the same as simply the loss of the human species.

This last consideration leads to one further point that supports the overall spirit of this paper. If one is truly concerned with global catastrophic risks and believes that such risks deserve more attention than even our current problems of global poverty, then one ought to dedicate one's time and resources to preventing catastrophic risks. Thus even if consideration of global catastrophic risks shows that we should be more concerned with AGCC than with current global poverty, it doesn't then imply that "practically non-maximizing" maximizing consequentialists ought to prioritize FAMCC; rather, such individuals should prioritize funding general projects which do the best job of avoiding catastrophic harm in all its possible forms.

10. Conclusions

This paper has taken a circuitous route. First, I argue, contrary to James Garvey and others, that a case can be made using an expected utility calculation involving a simple example which may provide motivation for individuals to avoid even small-scale actions which emit GHGs. However, I claim that this argument has a limited scope because, even independent of AGCC, there are many actions which can prevent the same amount of harm with even less sacrifice – the very actions recommended by FAM. Last, I provide reasons why AGCC should still be of great ethical concern for us. In particular, even if individual actions which emit GHGs are less morally important than other actions individuals might take, AGCC as a collective action problem is still one of the most pressing ethical issues of our time. This is especially so if non-human harms are morally important, and if we also countenance the dangers of possible catastrophic harm due to AGCC. However, the risk of catastrophic harm from other sources also looms large as a global ethical concern.

Singer, in concluding FAM, encourages philosophers to discuss and act upon issues of global poverty.⁴⁶ I should say the same about AGCC. Perhaps because of the complex empirical issues in climate economics and in climate science, philosophers have been loath to undertake a rigorous comparison of potential harms and benefits of different schemes. But the arguments of this paper, despite the multiple caveats I have made and the tremendous uncertainty involving AGCC, still do show that AGCC is a subject of great ethical concern.

Portland State University

An earlier version of this paper was presented at the Rocky Mountain Ethics Congress, 2011, and I would like to thank the audience on that occasion and my commentator, David Killoren. An earlier version was also presented at the International Society for Utilitarian Studies conference, 2011, and I would like to thank the audience there, especially Peter Singer, Toby Ord, William MacAskill, and Leonard Kahn. I would also like to thank two anonymous referees from this journal for very helpful comments.

¹ Singer, "Famine, Affluence, and Morality." I will assume for the purpose of this paper that the reader is familiar with Singer's argument in this work.

² See for instance Sinnott-Armstrong, "It's Not *My* Fault: Global Warming and Individual Moral Obligations," Jamieson, "Ethical Obligations in a Tragedy of the Commons," Sandler, "Ethical Theory and the Problem of Inconsequentialism: Why Environmental Ethicists Should Be Virtue-Oriented Ethicists," Garvey, "Climate Change and Causal Inefficacy: Why Go Green When It Makes No Difference?"

³ Here, Garvey is assuming that his emissions are average globally per capita, which is likely an underestimate due to the differences between developed and developing world per capita emissions. Also, Garvey is rounding up the numerator of his equation; to be precise, he should be asking whether he should aim for $3.99999999/28,399,999,999$. Nevertheless, the precise figure is not essential to Garvey's point.

⁴ Nolt, “How Harmful Are the Average American's Greenhouse Gas Emissions?”

⁵ Hiller, “Morally Significant Consequences of Ordinary Individual Actions.”

⁶ Hale, “Nonrenewable Resources and the Inevitability of Outcomes,” and Johnson, “Ethical Obligations in a Tragedy of the Commons,” both make this type of claim.

⁷ Hale “Nonrenewable Resources and the Inevitability of Outcomes,” 385.

⁸ See Anderson and Bows, “Reframing the Climate Change Challenge in Light of Post-2000 Emission Trends,” esp. 3865.

⁹ See Silverstein, “Utilitarianism and Group Coordination,” for a refutation of this view.

¹⁰ I discuss this issue in more detail in Hiller, “Climate Change and Individual Responsibility.” A further issue is that if GHG were extremely low but greater than zero, those emissions would not cause any harm because the atmosphere would be able to adjust. (For instance, it is commonly held that atmospheric CO₂ levels below 350 ppm would not lead to harmful AGCC – see 350.org, “350 Science.”) Thus one might claim that early GHG emissions cause less harm than later ones. Thus the amount of harm expected to be due to current individual actions may be slightly larger than the average harm per unit of GHGs emitted. I should add that the account here is in accord with Shelly Kagan’s defense of a consequentialist account of individual responsibility in situations of group harm in his “Do I Make a Difference?” However, see Julia Nefsky, “Consequentialism and the Problem of Collective Harm: A Reply to Kagan” for concerns about Kagan’s view.

¹¹ Or, perhaps, there is even a smaller (but still non-zero) chance that, due to the GHG emissions, several people’s lives will be ruined. This could happen, for instance, by slightly increased flooding due to slightly higher ocean levels.

¹² Singer, “Famine, Affluence, and Morality,” 231.

¹³ Ibid.

¹⁴ Singer, “Famine, Affluence, and Morality,” 229-230 and passim.

¹⁵ Singer, “Famine, Affluence, and Morality,” 229, 231, 235, and passim.

¹⁶ Cf. Singer, “Famine, Affluence, and Morality,” 230.

¹⁷ Precise figures regarding GHG emissions of airplanes are difficult to determine, but the best estimate is that there is .28 KG of CO₂-equivalent emitted per passenger per KM, so that approximately 1200 KG of CO₂-equivalent is

emitted by 2500 miles of air travel – over 80 times that of a 25 mile drive. See Ross “GHG Emissions Resulting from Airline Travel,” for information about per-passenger airline emissions.

¹⁸ See for example Pogge, “World Poverty and Human Rights.”

¹⁹ Note: this percentage is approximately thirty six times higher than the direct fatality rate for driving, which in the U.S. is approximately 1.1 per 100,000,000 vehicle-miles (or 1 in 3.6 million per 25 mile drive). See National Highway Transportation Safety Administration, “Fatality Analysis Reporting System Encyclopedia.”

²⁰ This is similar to the issue discussed by Singer (240) about the long-term effects of population increases.

²¹ Cf. Singer, “Famine, Affluence, and Morality,” 233.

²² Cf. Singer, “Famine, Affluence, and Morality,” 239.

²³ Ibid.

²⁴ Cf. Singer, “Famine, Affluence, and Morality,” 236.

²⁵ For analysis of this issue, see for example Kagan, “Does Consequentialism Demand Too Much? Recent Work on the Limits of Obligation,” and Mulgan, The Demands of Consequentialism.

²⁶ One could avoid this conclusion by appealing to a doctrine of double effect, but that doctrine is *highly* suspect.

²⁷ See Ord, “The Moral Imperative Towards Cost Effectiveness.”

²⁸ Here I am setting aside the fact that typically, maximizing consequentialists also believe that we have a duty to promote good in addition to preventing bad. Including this complicates the argument even further but does not change the spirit of what I say above.

²⁹ Singer, The Life You Can Save: Acting Now to End World Poverty, p. 103.

³⁰ From the data in fn. 12, 2500 miles of air travel causes 20 person-days of harm; if the plane ticket costs \$400, which may be approximately ¼ of the amount needed to save a life, and a life is assumed to be 80 years, then giving away the value of the plane ticket would help prevent 20 person-*years*. Of course these are very rough estimates; still, the scale of difference is enormous even if the estimates are off by even 500%.

³¹ See Singer, “Questions for Peter Singer.”

³² See Norcross, “The Scalar Approach to Utilitarianism.”

³³ Lomborg, The Skeptical Environmentalist: Measuring the Real State of the World and Cool It: The Skeptical Environmentalist's Guide to Global Warming.

³⁴ See for instance Grist, “A Skeptical Look at the Skeptical Environmentalist.” <http://grist.org/article/of1/>.

-
- ³⁵ For example in Lomborg, The Skeptical Environmentalist: Measuring the Real State of the World, 272.
- ³⁶ Sachs, The End of Poverty: Economic Possibilities for Our Time.
- ³⁷ See Broome, Counting the Cost of Global Warming, Ch. 3 for a seminal discussion of discounting in the context of AGCC.
- ³⁸ Stern, The Economics of Climate Change: The Stern Review.
- ³⁹ See Blois et al., “Small Mammal Diversity Loss in Response to Late-Pleistocene Climatic Change.”
- ⁴⁰ See Carter, “Towards a Multidimensional Environmentalist Ethic.”
- ⁴¹ Hiller, “System Consequentialism.”
- ⁴² See Ng, “Consumption Tradeoff vs. Catastrophes Avoidance: Implications of Some Recent Results in Happiness Studies on the Economics of Climate Change.”
- ⁴³ See Frame and Allen, “Climate Change and Global Risk.”
- ⁴⁴ See Parfit, Reasons and Persons, 453-54.
- ⁴⁵ See Bostrom and Ćirković, Global Catastrophic Risk, for an extensive survey of the possibilities.
- ⁴⁶ Singer “Famine, Affluence, and Morality,” 242-43.

References

- 350.org, “350 Science,” <http://350.org/en/about/science>, accessed November 1, 2013.
- Anderson, Kevin, and Alice Bows. “Reframing the Climate Change Challenge in Light of Post-2000 Emission Trends,” Philosophical Transactions of the Royal Society 366 (2008): 3863-82.
- Blois, Jessica, Jenny L. McGuire, and Elizabeth A. Hadly. “Small Mammal Diversity Loss in Response to Late-Pleistocene Climatic Change,” Nature 465 (2010): 771-74.
- Bostrom, Nick, and Milan Ćirković, eds. Global Catastrophic Risk. Oxford: Oxford University Press, 2008.

- Broome, John. Counting the Cost of Global Warming. Cambridge, UK: White Horse Press, 1992.
- "The Most Important Thing about Climate Change," in Public Policy: Why Ethics Matters, edited by Jonathan Boston, Andrew Bradstock, and David Eng, 101-16. Canberra: ANU Press (2010).
- Carter, Alan. "Towards a Multidimensional Environmentalist Ethic," Environmental Values 20(3) (2011): 347-74.
- Frame, David, and Myles R. Allen. "Climate Change and Global Risk," in Global Catastrophic Risks, edited by Nick Bostrom and Milan Ćirković, 265-286. Oxford: Oxford University Press, 2008.
- Garvey, James. "Climate Change and Causal Inefficacy: Why Go Green When It Makes No Difference?" Royal Institute of Philosophy Supplement 69 (2011): 157-74.
- Grist, "A Skeptical Look at the Skeptical Environmentalist." <http://grist.org/article/of1/>, accessed November 1, 2013.
- Hale, Benjamin. "Nonrenewable Resources and the Inevitability of Outcomes," Monist 94 (3) (2011): 369-90.
- Hiller, Avram. "Climate Change and Individual Responsibility," Monist 94(3) (2011): 349-68.
- "Morally Significant Consequences of Ordinary Individual Actions," Ethics, Policy, and Environment 14(1) (2011): 19-21.
- "System Consequentialism," in Consequentialism and Environmental Ethics, edited by Avram Hiller, Ramona Ilea, and Leonard Kahn: 54-69. New York: Routledge, forthcoming.
- Jamieson, Dale. "When Utilitarians Should Be Virtue Theorists," Utilitas, 19(2) (2007): 160-83.

- Johnson, Baylor. "Ethical Obligations in a Tragedy of the Commons," Environmental Values, 12(3) (2003): 271-87.
- Kagan, Shelly. "Does Consequentialism Demand Too Much? Recent Work on the Limits of Obligation," in Consequentialism, edited by Philip Pettit, 391-406. Dartmouth: Aldershot, 1993.
- , "Do I Make a Difference?" Philosophy and Public Affairs 39(2) (2011): 105-141.
- Lomborg, Bjørn. Cool It: The Skeptical Environmentalist's Guide to Global Warming, New York: Knopf, 2007.
- , The Skeptical Environmentalist: Measuring the Real State of the World, Cambridge, UK: Cambridge University Press, 2001.
- Mulgan, Tim. The Demands of Consequentialism, Oxford: Oxford University Press, 2001.
- National Highway Transportation Safety Administration, "Fatality Analysis Reporting System Encyclopedia." <http://www-fars.nhtsa.dot.gov/Main/index.aspx>, accessed November 1, 2013.
- Nefsky, Julia. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan," Philosophy and Public Affairs 39(4) (2012): 364-395.
- Ng, Yew-Kwang. "Consumption Tradeoff vs. Catastrophes Avoidance: Implications of Some Recent Results in Happiness Studies on the Economics of Climate Change," Climatic Change 105 (1-2) (2011): 109-27.
- Nolt, John. "How Harmful Are the Average American's Greenhouse Gas Emissions?" Ethics, Policy and Environment 14 (1) (2011): 3-10.
- Norcross, Alastair. "The Scalar Approach to Utilitarianism," in The Blackwell Guide to Mill's Utilitarianism, edited by Henry West, 217-32. Oxford: Oxford University Press, 2006.

Ord, Toby. "The Moral Imperative Towards Cost Effectiveness,"

<http://www.givingwhatwecan.org/about-us/our-research/the-moral-imperative-towards-cost-effectiveness>, 2012.

Parfit, Derek. Reasons and Persons. Oxford: Oxford University Press, 1984.

Pogge, Thomas. "World Poverty and Human Rights," Ethics and International Affairs 19 (1) (2005): 1–7.

Ross, Davide. "GHG Emissions Resulting from Airline Travel,"

http://www.carbonplanet.com/downloads/Flight_Calculator_Information_v9.2.pdf, 2009

Sachs, Jeffrey. The End of Poverty: Economic Possibilities for Our Time, New York: Penguin Press, 2005.

Sandler, Ronald. "Ethical Theory and the Problem of Inconsequentialism: Why Environmental Ethicists Should Be Virtue-Oriented Ethicists," Journal of Agricultural and Environmental Ethics 23 (2010): 67-183.

Silverstein, Harry S. "Utilitarianism and Group Coordination," Noûs 13 (1979): 335-60.

Singer, Peter. "Famine, Affluence, and Morality," Philosophy and Public Affairs 1 (3) (1972): 229-43.

----- "Questions for Peter Singer," New York Times Magazine, December 24, 2006;

<http://www.utilitarian.net/singer/interviews-debates/20061224.htm>, 2006.

----- The Life You Can Save: Acting Now to End World Poverty, New York: Random House, 2010.

Sinnott-Armstrong, Walter. "It's Not *My* Fault: Global Warming and Individual Moral

Obligations," in Perspectives on Climate Change: Science, Economics, Politics and

Ethics, Advances in the Economics of Environmental Resources, vol. 5, edited by Walter Sinnott-Armstrong and Richard B. Howarth, 285-307. Amsterdam: Elsevier, 2005.

Stern, Nicholas. The Economics of Climate Change: The Stern Review, Cambridge, UK: Cambridge University Press (2007).