Collective Action, Climate Change, and the Ethical Significance of Futility[*]

Mark Bryant Budolfson
budolfson /AT/ princeton /DOT/ edu
Version 5.2

Consider this familiar idea: *If you know that a course of action would be both costly and futile, then you aren't required to take that course of action.* I argue that this idea is true when properly understood, and that a careful investigation of related issues leads to significant progress on the ethics of collective action and the nature of moral reasons. I investigate these issues by means of the particularly useful example of whether you are required to reduce your greenhouse gas emissions. Along the way, I present arguments against influential views on climate change ethics including those of John Broome and Walter Sinnott-Armstrong, as well as arguments against influential views on the general ethics of collective action including those of Jonathan Glover and Derek Parfit.

[*] Special thanks to Chris Griffin, RJ Leland, Tristram McPherson, Eliot Michaelson, David Plunkett, Ryan Robinson, Gideon Rosen, George Rudebusch, Peter Singer, and Jane Willenbring for especially generous discussion of these issues. Thanks also to Sameer Bajaj, John Basl, Alexander Berger, Heather Berginc, John Bistline, Tom Blackson, Friedhelm von Blanckenburg, Will Braynen, Cheshire Calhoun, Eamonn Callan, Richard Yetter Chappell, Tim Clark, Josh Cohen, Christian Coons, Chiara Cordelli, Michael Cox, John Devlin, Jeff Downard, David Faraci, Ada Fee, Blake Francis, Sarah Hannan, Andy Hultgren, Ryan Jenkins, Hyunseop Kim, Lily Lamboy, Melissa Lane, Hugh Lazenby, Sarah McGrath, Alison McQueen, Kristi Olson, Govind Persad, Ángel Pinillos, Joe Rachiele, Zaira Razu, Rob Reich, Dan Russell, Carolina Sartorio, Debra Satz, Dave Schmidtz, Dan Shahar, Daniel Silvermint, Michael Smith, Brent Sockness, John Thrasher, Ian Vandeventer, Chad Van Schoelandt, Steve Wall, Brandon Warmke, Michael Weber, Lisandra West, Sara Worley, and audiences at Stanford Political Theory Workshop, the University of Arizona Center for the Philosophy of Freedom, Bowling Green, and Northern Arizona University, where I first presented this material to a public audience in 2009.

1. The Familiar Futility Argument

Climate change may be the most serious problem facing humanity today. If we do nothing about it billions of people will likely suffer in one way or another, and a worldwide catastrophe may ensue. In short, climate change is a serious problem.

However, it doesn't follow that you are required to do anything about it. To see why, suppose for a moment that you knew that nothing could be done about it; then you certainly wouldn't be required to do anything about it. For example, if there were a giant asteroid headed toward the earth that was sure to kill us all, that would count as a serious problem, but if you knew that nothing could be done about it, you wouldn't be required to do anything about it. So, if there is a serious problem, it doesn't follow that you are required to do anything about it.[1]

Of course, there are important differences between this asteroid story and our situation with respect to climate change. One difference is that we *can* do something about climate change, whereas in the asteroid case just described there is nothing we can do about the threat. Does this difference alone make for a difference in what individuals are required to do about climate change?

To investigate this, consider a different version of the asteroid story. This time, we discover that a giant asteroid is headed toward earth and will kill us all if we do nothing, but this time we can do something to stop it: we can stop it if and only if we all cooperate and do something together. In particular, suppose that we will avoid destruction from the asteroid if and only if we all congregate in Kansas and fire our ray guns into the sky in unison; if some of us are

[1] Here and in the rest of this paper by 'required' I mean 'ethically/morally required'.

absent, the collected beam from our guns will be inadequate to destroy the asteroid and it will kill us all; but if we are all there, the collected beam will be just powerful enough to vaporize the asteroid.

In this case, it seems like you would be required to go to Kansas and fire your ray gun *if* you had good reason to think that everyone else would do the same thing. However, if it was clear that not nearly enough people would show up, then you would not be required to sacrifice your remaining days on earth by taking a futile trip to Kansas. This shows that even if we can collectively do something about a serious problem, it does not follow that you as an individual are required to do anything about it, because you might know that although we can collectively do something about it, in fact we won't.

There are further important differences between the last asteroid story and our situation with respect to climate change. For example, in the last version of the asteroid story we assumed that the threat from the asteroid was a natural threat, whereas in the real world the threat from climate change is a threat that *we have caused*; in addition, in the last story we imagined that an individual going to Kansas would have no effect at all unless everyone else did the same, whereas in the real world you know that reducing your greenhouse gas emissions *will have some very small effect* on the level of such gasses in the atmosphere. Do these differences make for a difference in what you are required to do?

Suppose again that a giant asteroid will kill us all if we do nothing. As before, we can save our planet if and only if we all cooperate, but it is clear that we won't all cooperate, and so our planet is doomed. However, this time the destruction of the Earth will be *delayed* based on the number of people who show up in Kansas and fire their ray guns in unison at the asteroid. In particular, for every person who shows up and fires their gun from Kansas, the destruction of the

Earth will be delayed by 1/1,000 of a second.[2] In addition, suppose that the only reason we are threatened by the asteroid is that one of our mineral-extracting expeditions to the asteroid knocked it into its threatening orbit, and we do not have time to deflect it using any other means.

Do these changes make any difference to what you as an individual are required to do? They do not. It remains true that if you know that only a small percentage of people will show up in Kansas, thereby sealing our fate, then you aren't required to sacrifice your remaining days on earth by taking a trip to join them. If this isn't immediately clear, focus on the fact that if you went to Kansas, most of your friends, neighbors, extended family, and others would remain at home, enjoying much of their remaining time on Earth with their friends and families, while you'd spend the next week on an unpaid leave from work, on a costly and unpleasant journey filled with traffic congestion and uncomfortable anguish, which would culminate in a bizarre scene in a strange place, surrounded by strange people – and all of that merely to delay the inevitable end of everything by 1/1,000 of a second. Then, after the depressing ride home, many would ask you why you wasted so much of your remaining time and money in such a foolish way. Could you really be required to incur such significant costs to benefit the rest of the world in such a trivial way, given that most of the rest of the world is not willing to do the same? No. In such a situation it would be permissible for you, like the others, to skip the trip to Kansas and spend your time doing something more valuable to you and your family instead.

This is the basis for the familiar futility-based argument that you, as an individual, are not required to reduce your greenhouse gas emissions by a significant amount – just as you wouldn't be required to make the trip to Kansas in the story above, so too you aren't required to reduce your emissions by a significant amount, given that billions of people in China, India, and the rest

[2] Why can we delay the end of the world in this way? Imagine that this is our only opportunity to influence the orbit of the asteroid, and that each shot will result in the asteroid taking a slightly wider orbit, thereby very slightly delaying the moment when it will return, unstoppable, to destroy the Earth.

of the world don't care about climate change and won't care until it is too late. More specifically, you know that the lack of emissions reductions by others will lead to a terrible outcome regardless of what you do, and you know that your own emissions reductions would not have any noticeable effect on the outcome for anyone except to make it much worse for you; therefore, since there are no other ethically significant factors that differ from those in the story above, you are not required to reduce emissions, just as you would not be required to take the costly and futile trip to Kansas in the story above.

Although this kind of argument is familiar, the issues it raises are usually not taken seriously by ethicists – which is unfortunate, because taking the issues seriously yields important insights into both the ethics of collective action and the nature of ethical reasons more generally.[3] In what follows I take the issues seriously, and I mine all the important insights that the familiar futility-based argument and related arguments have to offer. In the next section, I show that there is no obvious incoherence involved in the familiar futility-based argument, and I note a number of potentially surprising facts about morality that are highlighted by uncontroversial examples involving collective action problems, thereby making an initial case for taking these issues seriously. Then, in the following sections I examine a range of objections to the familiar futility-based argument that purport to show that it is misguided. In these sections, I focus, as above, on the particular instance of the argument that claims that individuals are not required to reduce emissions. I focus on this issue about climate change and individual action because the various objections to the futility-based about individual emissions reductions invoke all of the very wide range of factors that commentators have taken to show that the familiar futility-based argument is misguided in other contexts. Furthermore, the empirical facts related to climate change are

[3] Of course, some ethicists have taken the futility argument seriously – for example, I consider general discussions by Jonathan Glover and Derek Parfit below, as well as discussions in connection with climate change by John Broome and Walter Sinnott-Armstrong.

particularly useful for testing various competing explanations of what individuals are required to do in collective action situations. After carefully evaluating the variety of deontological, consequentialist, and virtue-based objections to the futility argument regarding the climate issue, I discuss a few other illuminating test cases, and I use the results to inform a general analysis of the ethical significance of futility and the ethics of collective action. I conclude by commenting on the practical significance of these insights.

2. Taking the Futility Argument Seriously: The Sophisticated Version of the Argument

With that overview in hand, in this section I argue that there is no obvious absurdity or incoherence in the sort of familiar futility-based argument noted above. Toward that end, I focus again on the familiar argument that you are not required to reduce your greenhouse gas emissions by a significant amount.

Again, the intuitive idea behind the futility-based argument is that you are not required to do something if it would be both costly and futile. However, this idea could be interpreted in a number of different ways, some of which are implausible. For example, it could be understood to imply that it is permissible for you to steal other people's property if you know that someone will steal or destroy the property regardless of what you do; it could also be understood to imply that it is permissible for you to ignore the suffering of distant strangers if you are unable to make a difference to the underlying problems that cause such suffering to exist in the world; it could even be understood to imply that it is permissible to join in a direct physical assault if an additional assailant would leave the victim no worse off. To avoid interpreting the argument as having such implausible implications, while preserving the ability of the argument to generate

interesting conclusions, it is most interesting to interpret the core idea behind the argument as follows:

> **Neutral Interpretation of the Ethical Significance of Futility**
> If you know that a course of action would be costly to you, and that there are no significant welfare-based reasons that support that course of action, and that there are also no significant deontological reasons that support that course of action, then you are not required to take that course of action.

When the core idea is clarified in this way, it avoids the implausible implication that it is permissible for you to steal another person's property even if that property would be stolen by someone else anyway, and it avoids the implication that it is permissible to join in a direct physical assault when doing so would make no difference – because there are very significant standard rights-based reasons not to do such things. It also avoids the implication that it is always permissible for you to ignore those who are suffering when your aid is futile with respect to the underlying problem, because there are significant welfare-based reasons to aid whenever you can make a difference to even one person, even if you cannot do anything about the underlying situation that is causing such general suffering. As a result, the familiar futility-based argument remains of interest from both deontological and utilitarian perspectives, as it is consistent with both.[4]

In addition, the familiar argument should not be interpreted as implying that *we* shouldn't do anything about climate change. Instead, it should be interpreted merely as arguing that *individuals* are not required to reduce emissions, which is perfectly consistent with the claim that

[4] Here it may help to note that the Neutral Interpretation of the Ethical Significance of Futility states only a sufficient condition for the permissibility of inaction, which is why it remains consistent with both deontology and consequentialism.

we are *collectively* required to reduce emissions. At first glance, this might seem incoherent, because it might be seem that if we are required to reduce our emissions collectively, then it follows that each of us is required to reduce our emissions individually. However, that does not follow, because the ethically relevant properties of a group are often different from those of its members. For one thing, a group of people Φ-ing often has significantly different effects than an individual Φ-ing; as a result, in a situation where the effects of Φ-ing are ethically important, it might be impermissible for the group to Φ, but permissible for each member of the group to Φ.

For example, imagine that we find ourselves in the middle of a stampede. If the stampede continues, it is clear that an increasing number of innocent people will be killed. Luckily, each of us has a button in our hands: once more than 50% of us press our buttons, a painless 'sleep' device in all of our spines will be activated, causing the stampede to stop peacefully. Other than pressing the button and then continuing to stampede, the only other options we have are to continue stampeding without pressing the button, or else to stop stampeding and certainly be run over and killed. In this case, the thing for each of us to do is to press our button and then continue stampeding, hoping that enough others press their buttons to end the stampede.

So, suppose you press your button, and the stampede continues; this means that less than a sufficient number of the others have pressed their buttons. You know that the stampede will stop once a sufficient number do press their buttons, but now, in the meantime, the only relevant options that you have are to continue stampeding and live, or else stop stampeding and die, given that you have already pressed your button. In this case, you are permitted to continue stampeding, and so is every other individual; at the same time, we are collectively required to stop stampeding. This shows that the objection to the intuitive argument under consideration fails, because individual requirements do not follow from collective requirements.

Perhaps more importantly, in the case just described, individuals are not required to *stop stampeding,* although individuals are required to *press their button*. A proponent of the intuitive argument should insist that analogous remarks apply to our situation with respect to climate change. On such a view, individuals are not required to *reduce emissions,* although individuals are required to *favor effective public policy solutions to climate change*. It is crucial to distinguish the second political act from the more direct act of reducing your emissions, just as in the stampede story it is crucial to distinguish the more political act of pressing your button from the more direct act of stopping stampeding. With this distinction in mind, a proponent of the familiar argument should agree that even individuals are required do *something* about climate change; however, that something is different than it might initially appear – the thing that individuals are required to do about climate change is *political*, and not directly tied to personal emissions.

Finally, the intuitive argument might seem to overgeneralize and imply that it is always permissible to pollute, no matter how wasteful or egregious the pollution. However, the familiar argument should not be interpreted as having that implication if it is to be understood charitably. Instead, the main point of the argument is that when a reasonable individual goes about her daily life she causes emissions that *do not noticeably harm anyone*, even when repeated over many years, and that *could be avoided only by making significant sacrifices of things that are important*; furthermore, these emissions occur in a context in which it is clear that *almost everyone else would be unwilling to make such sacrifices if they were in similar circumstances*. The claim that emissions that meet all of these conditions are permissible does not in any way imply that wasteful or egregious pollution is permissible. For example, wasteful emissions are emissions that could be avoided *without significant sacrifices*, and so the interpretation of the

intuitive argument above does not license such emissions; similarly, egregious pollution is pollution that *tangibly harms others*, and so the interpretation of the intuitive argument above does not license such pollution. For these reasons, the intuitive argument does not overgeneralize and lead to absurd conclusions. Instead, it appears to point the way toward a carefully reasoned basis for condemning wasteful and egregious pollution, while allowing that reasonable emissions are often permissible – something that other accounts often fail to do.[5]

Another thing to note is that the familiar argument is most plausible when its conclusion is understood along the lines of the claim that *as things stand individuals are not required to perform the collectively beneficial 'cooperative' act given the current general non-cooperation of others* – which is perfectly consistent with the claim that in the future individuals would be required to perform the cooperative act and make costly sacrifices if most of the others started cooperating. For example, the conclusion that as things stand individuals are not required to reduce their greenhouse gas emissions is consistent with the view that each would be required to make such reductions if the rest of the world somehow managed to cooperate and reduce emissions in a way that had some hope of making an important difference to the climatic outcome. Such claims are perfectly consistent because there is an important difference between, on the one hand, what individuals are required to do in a collective action problem in which almost everyone else is *defecting*, and, on the other hand, what individuals are required to do in an otherwise similar collective action problem in which most others are *cooperating*. Refusing to

[5] Existing accounts do not plausibly distinguish between permissible and impermissible emissions. For example, as I argue in detail below, John Broome's deontological account seems to overgeneralize and yield the verdict that emissions are never permissible. As another example, Henry Shue has noted that emissions that are necessary for survival are undoubtedly permissible, and must in effect be guaranteed to each person on the planet as part of any ethically permissible intergovernmental emissions permit allocation scheme. Although true, this does not provide a general distinction between permissible and impermissible emissions, because even if everyone is given a right to emissions that are necessary for survival, there is still the further crucial question of what *additional* emissions would be permissible (Shue, "Subsistence Emissions and Luxury Emissions", in Stephen Gardiner et. al. (ed.) *Climate Ethics*.) As this reveals, distinguishing permissible and impermissible emissions is a difficult task that requires careful argument and reflection, which I provide in the following discussion.

cooperate when almost everyone else is defecting is merely refusing to be a *sucker*, which is not generally morally objectionable, whereas refusing to cooperate when almost everyone is cooperating is *free-riding*, which is often morally objectionable. The view that you are not generally required to be a sucker is perfectly consistent with the view that free-riding is impermissible in otherwise similar cases in which almost everyone is cooperating, and a proponent of the familiar argument should therefore stress that the argument is intended only to apply to cases involving general non-cooperation such as our actual situation with respect to climate change.

Along these lines, a further point is that even if one remains suspicious of an argument based on the Neutral Interpretation of the Ethical Significance of Futility, the familiar futility-based argument could still be grounded on a foundational ethical principle that it is even more difficult to deny, such as the following:

> **Non-Cooperation Principle**
> If an individual faces a collective action problem in which she knows that most others will fail to choose the 'cooperative' option, thereby ensuring a collectively bad outcome, she is not required to choose the cooperative option herself if she knows that:
> (i) choosing the 'cooperative' option would involve her making significant sacrifices,
> (ii) there are no significant reasons for her to choose the 'cooperative' option, including either welfare-based reasons or deontological reasons,[6]
> (iii) most of the others would not choose the 'cooperative' option even if they were in her exact situation, and
> (iv) choosing the 'cooperative' option would not allow her to be part of a successful solution to a different collective action problem.

[6] This condition grants (without endorsing) the possibility of reasons of 'fair play' that (arguably) apply in the case of voting and successful provision of public goods. The standard conditions that give rise to such reasons are not present in the climate case, as I argue below.

It is difficult to see how this principle could be false, because it seems to capture what is most compelling and defensible about the intuitive idea that even if *free-riding* is impermissible in response to the general *cooperation* of others, at the same time one is not required to be a *sucker* by cooperating in the face of the general *non-cooperation* of others – and this principle together with, for example, the facts of our situation with respect to climate change seems to imply that individuals are not required to reduce their emissions by a substantial amount as things stand. At the same time, this principle is consistent with the view that you *would* be required to reduce emissions by a significant amount if the rest of the world somehow managed to cooperate and reduce emissions in a way that had some hope of making an important difference to the outcome; however, because the actions of others are not currently or foreseeably securing a cooperative outcome, you are not required to make such sacrifices as things stand. Again, there is an important difference between, on the one hand, what individuals are required to do in a collective action problem in which almost everyone else is not cooperating, and, on the other hand, what individuals are required to do in an otherwise similar collective action problem in which others are cooperating. The view that you are not required to be a sucker is perfectly consistent with the view that free-riding would be impermissible in an otherwise similar case in which almost everyone is cooperating.[7]

[7] As Robert Nozick has shown, it is sometimes permissible to engage in behavior that could be described as free-riding. For example, suppose that your neighbors spend a very large amount of time cleaning the sidewalks around your houses – in particular, imagine they get down on their knees and use toothbrushes, etc. to ensure that the sidewalks are as clean as humanly possible. In such a case, you are not required to help pay for this excessive maintenance even if you enjoy having the sidewalks so clean and thus benefit from their cooperative actions; similarly, your neighbors would not be justified in forcibly taking money from you to pay 'your share' of such a project. Nozick uses cases like this to show the falsity of the principle (due to H.L.A. Hart and endorsed by John Rawls) that if you benefit from the cooperative actions of others, you are required to pay your proportional share of the costs, and/or the state is justified in extracting those dues from you by force. (See *Anarchy, State, and Utopia*, pp. 90-95.) A more plausible principle regarding free-riding cases is suggested by my discussion above: free-riding is impermissible when and only when you owe it to the others to cooperate; in the Nozick-style case just discussed,

In light of all of the preceding, there is no obvious absurdity or incoherence in the sort of familiar futility argument noted above. On the contrary, such an argument can be articulated in a sophisticated and powerful form, and has the humane implication that every individual is required to favor and support a just solution to the problem of climate change, which suggests that the futility argument may even be consistent with a Rawlsian view of individual duties in the face of tragically inadequate institutions.[8] This provides a prima facie case for taking the argument seriously, and for thinking more carefully about the ethical significance of futility. Of course, many important issues and objections remain, which are addressed in the following sections.

3. Deontological Objections to the Futility Argument

In this section I consider a number of important deontological objections to the futility argument, which arise in an especially clear way in connection with the individual ethics of climate change. In the following sections I consider other deontological, consequentialist, and virtue-based objections, including objections that appeal to aggregation of imperceptible effects and the possibility of threshold effects, which often seem crucial to the ethics of collective action. These objections are of general interest because collectively they invoke the wide range of factors that commentators have thought most relevant to the ethics of collective action.

---

you do not owe it to the others to cooperate in their sidewalk-cleaning scheme, and so you are not required to cooperate. (Of course, this leaves unanalyzed the conditions under which you owe it to others to cooperate in funding some public good.)

[8] Compare John Rawls, *A Theory of Justice* (revised edition), pp. 293-4. It also suggests that the futility argument does not imply that morality is *directly collectively self-defeating*, in the sense of Derek Parfit, *Reasons and Persons*, chapter 2.

One kind of objection might be raised by Kantians, rule utilitarians, and other theorists who believe that individual requirements are determined by a principle of *universalizability*. Some of these theorists may be tempted to argue that the futility-based argument must be mistaken, because (for example) if *everyone* fails to reduce emissions by a substantial amount, then a disaster will ensue, which makes it the case that everyone is required to reduce emissions by a substantial amount.

The problem with this reasoning is that it seems to lead to the wrong verdict on the asteroid case and stampede case discussed above and many other collective action problems in which most people are not 'cooperating'. For example, in the asteroid case it is also true that a disaster will ensue if everyone fails to go to Kansas; however, it does not follow that you are required to go to Kansas.

This shows that a dilemma exists in connection with universalizability theories and collective action problems: the theories can either be understood in a way that clearly entails that the futility argument is mistaken, but in that case they deliver (what most would judge to be) mistaken verdicts on many cases involving general non-cooperation; alternatively, they can be understood in a way that does not generate any clear objection to the futility argument in cases that involve widespread non-cooperation. The upshot is that universalizability theories do not give rise to any clear objection to the futility argument in connection with cases that involve widespread non-cooperation. Furthermore, most actual universalizability theories are intended to be understood as not implying counterintuitive verdicts on cases of widespread non-cooperation, and are thus consistent with the futility argument.[9]

[9] I discuss related issues in much greater depth in my paper "Why Morality and All Other Forms of Normativity are Sometimes Dramatically Directly Collectively Self-Defeating".

Even setting aside universalizability theories, in the case of climate change it may seem that an important objection to the futility argument arises from the distinction between *actively causing harm* and *merely allowing harm*. According to this objection, there is a crucial disanalogy between the asteroid case and our actual case with respect to climate change, because in the asteroid case by not going to Kansas you would merely *allow* harm to befall others, whereas by continuing to emit greenhouse gasses you would *actively* harm others. John Broome offers a clear and influential statement of such a view:

> …almost everyone recognizes (with some exceptions) the elementary moral principle that you should not do something for your own benefit if it harms another person. ... Climate change will cause harm. … In going about our daily lives, each of us causes greenhouse gasses to be emitted. … In that way, what we each do for our own benefit harms others. Perhaps at the moment we cannot help it, and in the past we did not realize we were doing it. But the elementary moral principle I mentioned tells us we should try to stop doing it and compensate the people we harm.[10]

Broome notes that there are "some exceptions" to the principle about harm that he invokes in this argument. Unfortunately, these exceptions undermine his argument. The problem is that such a harm principle, if it is to be plausible, cannot be applied to cases of *accumulative harms* such as

[10] John Broome, "The Ethics of Climate Change", pg. 69. The context makes it clear that Broome is correctly interpreted as arguing that *individuals* are required to reduce emissions. I discuss Broome's views in much greater detail in "Global Justice, Political Realism, and the Ethics of Collective Action: Review of John Broome, *Climate Matters: Ethics in a Warming World*". Another problem for harm-based arguments like Broome's is that they seem to imply that seizing property via eminent domain is impermissible in a wide range of cases in which it is in fact permissible. Setting that aside, a related point in the current context is that it might seem, for example, that Bangladesh will have to be inundated with water for the greater good of others *in a way that is relevantly similar to other justified uses of eminent domain*, which means that it may be permissible for us *collectively* to act in a way that we know will inundate Bangladesh with water as long as we take sufficiently dramatic steps to resettle and compensate Bangladeshis. If that is right, then even some issues about what we are *collectively* required to do might come out differently than ethicists typically assume. I do not mean to endorse this line of reasoning – I merely intend to point out its significance, and the fact that it does not involve any *obvious* mistake. (Thanks to Gideon Rosen for noting the connection to eminent domain– I don't know that he'd favor any aspect of its consideration here.)

those associated with greenhouse gas emissions in the way that is needed for Broome's argument to go through. In his classic discussion of the sort of Millian harm principle invoked by Broome's argument, Joel Feinberg argues for this explicitly:

> When pollutants and similar threats to the public interest accumulate as the result of the joint and/or successive uneven contributions of numerous parties, each of which is "harmless" in itself except that it moves the condition of the environment closer to the threshold of harm, and most of which are positively beneficial in other ways, then blanket prohibitions of these contributions will often be as harmful as continued blanket permissiveness. The simple harm principle yields no solution to the problem, but it does legitimize the establishment of a regulatory system that can provide meaning to imputations of harmful conduct (e.g. excessive emissions), which can in turn support authoritative orders backed up by criminal sanctions. To be "wrongful," a contribution toward public accumulative harm must be a violation of an authoritative scheme of allocative priorities already in force. *In the absence of such a background there is no nonarbitrary method for imputing accumulative harms to individual parties.* The regulatory system in turn must balance values of efficiency, equity, and feasibility. Its priorities must be determined by the entire range of relevant reasons for policy decisions.[11]

Setting aside the details of this objection from Feinberg, it is perhaps more important to note that there are independently compelling reasons for thinking that individual accumulative harms cannot be treated as analogous to *canonical harms* where a stringent side constraint is violated by one's contribution, or where the effect of one's contribution makes particular individuals noticeably worse off, contrary to what Broome's reasoning assumes. For example, imagine that there is a single power plant in your region that emits water vapor as its only by-product, and

[11] Joel Feinberg, *The Moral Limits of the Criminal Law, Volume One: Harm to Others*, pg 244, emphasis added. Broome does not address the large literature on harm by Feinberg and others.

which as a result causes a snowstorm about once per year in the lesser-developed nation that borders your nation. In such a case, are you required to abstain from consuming power, simply because your consumption of power makes an intangible causal contribution to harm via the snowstorms that are caused by the power plant? In answering this question, it is important to keep in mind that snowstorms cause significant harm to people – for example, they cause transportation accidents in which people are harmed and killed, they cause power disruptions that can have substantial harmful effects, and they cause other undesirable things to happen. However, at the same time, even if such significant harms make it the case that the power plant's customers must pay *compensation* to those who are negatively impacted by the snowstorms, that does not settle the question of whether you are *required to reduce your consumption of power*. With that in mind, could you be required, at significant cost to yourself, to reduce your consumption of power simply because the power plant in your region occasionally causes a snowstorm in such a way?

To make the case more concrete, suppose that if you continue to consume electricity from the power plant you will cause a tiny amount of water vapor to be released which would not otherwise be released, where this additional vapor will have no effect on the once-per-year snowstorm that is caused by the power plant, except to cause a few extra flurries to fly. Given that assumption, it seems clear that you are not required to reduce your power consumption by a costly amount, even though you know that your power consumption plays a direct causal role in significantly harming some people in the adjacent country. This shows that Broome's principle fails to apply in just the sort of cases that are relevant to the ethics of climate change, because it shows that when the causal contribution of your action is negligible, and the cost to you of

avoiding that action is substantial, it does not follow from any "elementary moral principle" that you must avoid that action, even if your action plays a direct causal role in harming others.

Analogous cases support the same conclusion and suggest that the sort of de minimis rights infringements at issue can be permissible even when there are no dramatic ethical reasons for that infringement. For example, imagine that you buy a snow-making machine to make snow for your children in your backyard. You know that when you operate it, a few flurries will fly over onto your neighbor's land and will even occasionally fall on his head without his consent. Do you infringe his rights by operating the snow machine? Perhaps – but the infringement is so minimal that it cannot possibly make it impermissible for you to operate the snow-making machine.

This suggests a reply to another objection – namely, that the futility argument must be mistaken because it ignores the fact that you will *violate other people's rights* unless you reduce your emissions. The snowstorm case suggests that this objection is mistaken, because in that case the people who are harmed as a result of the relevant snowstorms have rights, but you do not infringe anyone's rights by continuing to consume electricity from the power plant – or, insofar as you do, they are de minimis infringements that do not imply that you act impermissibly.

As these cases illustrate, your own interests are ethically significant, just like everyone else's interests, and so in some cases where your own interests are *fully at stake* they outweigh the interests of others that are only *imperceptibly at stake* – and this can be true even when your own interests do not constitute dramatic ethical reasons. As a result, it is sometimes permissible to play a direct causal role in harming others even when there are no dramatic ethical reasons for doing so – especially when the causal role you play in harming others amounts to no more than a de minimis infringement of anyone's rights.

In light of this, it is important to note that the futility argument is defensible in the case of climate change even if we agree that individual emissions play a causal role in harming people. This is an important observation, because previous defenses of the futility argument in connection with climate change ethics, such as that of Walter Sinnott-Armstrong, depend essentially on the implausible claim that individual emissions do not play a causal role in harming people.[12]

This is in keeping with the more general goal of the discussion in this paper, which is to consider the most sophisticated objections to the futility argument, including arguments based on causing harm, the danger of crossing tipping points, and the aggregation of imperceptible effects, all of which are ignored or inadequately addressed by Sinnott-Armstrong and many other commentators who have deployed the futility argument in connection with climate change and other collective action problems. The discussion here takes these objections seriously, and does not rely on implausible claims in responding to them. This allows not only for a more responsible evaluation of the ethics of climate change, but also of the general factors relevant to the futility argument and the ethics of collective action.

Another objection might arise from the idea that you *owe it to others* to reduce your emissions by a substantial amount, even if you do not violate anyone's rights by failing to reduce your emissions. An initial reply to this objection is that in the snowstorm case discussed above, although it is plausible that you owe it to others to *compensate them* to an appropriate extent if they are harmed by a snowstorm caused by your power plant, it is not plausible that you owe it to them *not to use power from the power plant*. Similarly, even if it is plausible that you owe it to others to *compensate them* to an appropriate extent if they are harmed by climate change caused

[12] Walter Sinnott-Armstrong, "It's Not My Fault: Global Warming and Individual Moral Obligations", in Gardiner et. al. eds *Climate Ethics*, esp. pp. 334-337). Compare also Joakim Sandberg, "'My Emissions Make No Difference': Climate Change and the Argument from Inconsequentialism".

by global emissions, it is not plausible that you owe it to them to *reduce emissions*. Although such a combination of claims can seem unmotivated in the abstract, consideration of the snowstorm case and other cases above indicates that it is the correct thing to say upon reflection.

In light of this, it might seem that you are required to pay *compensation* to those who are harmed by your emissions, even if you are not required to avoid playing a role in causing the harm. However, even this is not necessarily true. For example, imagine that in the story above the power plant from which you get your power is owned by your government, and that your government refuses to compensate those who are harmed by the snowstorms that are caused by the power plant. Are you then personally required to compensate those who are harmed? What if the costs to you of providing the compensation are much larger than the amount of compensation that you owe? What if it is not your fault that your government decided to build such a power plant rather than a different power plant that would not harm anyone, and you consistently vote against your government's mistaken policies? Reflection on such questions reveals that if the costs to you are significant enough and the benefit to others is small enough, and especially if someone else or some other institution is directly at fault for the harm, and you support just alternative institutions, it can even be true that you as an individual are not required to compensate others for the harm you knowingly play a role in causing in pursuit of your own goals.[13]

At the same time, there is an important disanalogy here to the climate case, because in the climate case you can arguably compensate those who are harmed by your emissions-generating activities at little cost to yourself by purchasing offsets that reduce your carbon footprint. In other words, even if it is granted that you are not required *to make costly reductions in your emissions*-

[13] Compare the implications of Feinberg's view and a Rawlsian view of the duty to promote just institutions *when it is possible to do so at little cost to oneself*.

generating behavior, nonetheless it could be argued that you are required *to reduce your carbon footprint to zero*, on the grounds that you can do so in a way that is not costly by purchasing offsets. This view is consistent with the *Non-Cooperation Principle* because if there is a way of *reducing your carbon footprint to zero* that is sufficiently inexpensive, then *significant sacrifices* are not required (and thus condition (i) of the *Non-Cooperation Principle* is not satisfied relative to *reducing your footprint*), even if at the same time significant sacrifices would be required *to reduce your emissions*. Although this view relies on a subtle distinction, such a distinction is crucial for thinking carefully about the ethics of climate change.

Here it is worth returning Broome's view in more detail. Broome's view rests important weight on this kind of subtlety, and evaluating such subtleties within Broome's view provides a useful route into a more complete examination of deontological arguments. In more detail than above, Broome's argument is the following:

> *Broome's Harm Principle*: If one seriously harms others as a predictable result of avoidable, deliberate acts done for one's own benefit, knowing that it is not practically possible to compensate them, then one acts in a way that is wrong.[14]
>
> *Empirical Premise*: Most individuals in developed nations engage in avoidable emissions generating activities for their own benefit that predictably cause serious harm, knowing that it is not practically possible to compensate those who are harmed.
>
> *Conclusion about High Carbon Footprints*: Therefore, the actual emissions generating activities of these individuals are wrong, and they are wrong because they cause harm in the way captured by *Broome's Harm Principle*.[15]

[14] *Climate Matters*, pp. 55-57. Here Broome is focusing attention on deliberate self-interested acts, in contrast to omissions, accidental acts, and acts done for the greater good that can more easily be permissible even when they cause harm. Broome also intends his harm principle to be understood as implicitly containing other obvious conditions, such as that the harm is not consented to, is not just punishment for some wrongdoing, etc. (pg. 55).

> *Conclusion about Offsetting*: Therefore, these individuals must either reduce their emissions-generating activities or else offset their emissions in a way that makes their overall activities carbon neutral. If individuals go carbon neutral, then they can act permissibly even if they do not reduce their emissions-generating activities, because then their emissions do not harm to anyone, and thus are not wrong according to *Broome's Harm Principle*.[16]

In the discussion above, I argued that *Broome's Harm Principle* is implausible if it is interpreted in the way necessary for this kind of argument. In what follows I argue that even if *Broome's Harm Principle* is granted and the argument for the *Conclusion about High Carbon Footprints* is also assumed to be sound, it nonetheless follows that the *Conclusion about Offsetting* is false. This suggests that Broome's overall argument for the *Conclusion about Offsetting* fails by its own lights. After evaluating Broome's argument, I consider other objections to the futility argument based on the intuitive notion of *what we owe to each other*.

To see the problem with the *Conclusion about Offsetting* on the assumption that the argument for the *Conclusion about High Carbon Footprints* is sound, consider the following story: imagine that Ben has always wanted to throw darts into Joe's body for fun, but has not done so because he figured this would be wrong. However, Ben now sees several people who have independently decided to throw darts into Joe's body, and having read Broome's book, he figures that this is his chance to throw darts into Joe's body in a way that is permissible: all Ben needs to do is to pay one of these other people not to throw his darts into Joe's body, and then Ben will throw those darts into Joe's body himself; then, according to Broome's reasoning,

---

[15] *Climate Matters*, pg. 59.
[16] *Climate Matters*, pg. 87.

although Ben will throw several darts into Joe's body himself, he will not harm Joe because he has purchased offsets that make the overall effect of his acts harm neutral.

In this story, it would be absurd for Ben to reason in this way to the conclusion that he does not harm Joe on the grounds that his overall activities are neutral with respect his harm footprint on Joe. Unfortunately for Broome, his argument for the *Conclusion about Offsetting* seems to involve exactly this kind of reasoning – for example "The private morality of climate change starts by recognizing that your own individual emissions of greenhouse gas do serious harm" … "Offsetting does not remove the very molecules that you emit, but… If you successfully offset all your emissions, you do no harm by emissions".[17] This reasoning seems to invoke a conception of harm on which you do no harm if you cause injury via some actions but prevent harms elsewhere in a way that ensures that no one ends up worse off than they would have been if you had not engaged in any of those activities. However, as the dart example shows, this conception of harm seems indefensible. Once that becomes clear, the correct conclusion to draw seems to be that if an individual's emissions cause harm without offsetting, then they continue to cause harm even after offsetting, given the empirical fact that the particular molecules that constitute those emissions continue to exist in the atmosphere and continue to have the same injurious causal effect with or without offsetting. But if that's right, then it follows that Broome's *Conclusion about Offsetting* is false if his argument for the *Conclusion about High Carbon Footprints* is sound.

In response, it could be claimed that there is a crucial difference between the climate case and the dart case: namely, that there is *a very complex causal chain* in the climate case but not in the darts case. However, this reply is not open to Broome, because the entire point of Broome's argument is to insist that the complexity of the causal chain in the climate case does not

[17] *Climate Matters*, pp. 74 and 85.

constitute an ethically relevant difference to other cases in which harm is caused more directly.[18] Furthermore, independent confirmation that Broome cannot avoid the objection by appealing to the complexity of the causal chain is provided by imagining a version of the dart story that is the same as above, except that in this version Ben knows that the darts, once thrown, will pierce Joe's body only after passing through a complex causal chain analogous to the complex causal chain involved in the climate case.[19] Introducing this difference cannot make it permissible for Ben to cause harm to Joe in such a way, even if Ben offsets his harm by paying someone else not to throw additional darts as in the story above. Again, this shows that Broome's argument for the *Conclusion about Offsetting* fails by his own lights, and again the problem is that even if one purchases offsets, that does not change the nature or causal effects of one's emissions in any way: so, if an individual's emissions cause harm without offsetting (as the *Conclusion about High Carbon Footprints* claims), then they continue to cause harm even after offsetting, given the empirical fact that the particular molecules that constitute those emissions continue to exist in the atmosphere and continue to have the same injurious causal effect with or without offsetting.

Why does Broome's ethical view undermine his own *Conclusion about Offsetting* in this way? The source of the problem is that Broome assumes a deontological ethical theory on which there are stringent side-constraints that make it impermissible to cause harm under the conditions described in *Broome's Harm Principle*, and he assumes that the emissions of average individuals in a developed nations violate such constraints.[20] On such a deontological view, if one violates such constraints, then one acts wrongly regardless of whether one thereby brings about greater

[18] To deny that individuals cause harm via their emissions would undermine the entire 'causation of harm' story that Broome wants to tell, and would open the door to arguments such as Walter Sinnott-Armstrong's that depend essentially on the claim that individuals do not cause harm.

[19] Thanks to Blake Francis for suggesting this example that conjoins the initial darts example with a process involving a complex causal chain.

[20] "…our moral duty as private individuals is determined by the duty of justice not to harm… …this duty of justice is stringent. It requires each of us to avoid emitting any greenhouse gas at all" (*Climate Matters*, pg. 14).

good and regardless of whether one thereby or subsequently prevents such violations from being done by oneself or others.[21] So, if it is true that one's emissions violate such constraints without offsetting – as the argument up to and including the *Conclusion about High Carbon Footprints* implies – then it is irrelevant to the permissibility of those activities whether one subsequently offsets the similar violations done by others, contrary to what the *Conclusion about Offsetting* assumes.

Setting aside the details of Broome's actual argument for the conclusion that you are required to reduce your carbon footprint to zero, it is worth wondering whether any other argument for that conclusion might be sound. Toward that end, a different deontological argument that appeals to the intuitive notion of *what we owe to each other* appears able to escape the objections above:

> *Moral Debt Principle*: Contributing to a harm results in a moral debt in proportion to the magnitude of that harm and one's contribution to it at least insofar as one benefits from that harm, even if it is very costly to avoid contributing to that harm because of background facts that are beyond one's control and for which one is not responsible.[22]
>
> *Stringent Duty to Compensate Principle*: If one acquires a moral debt in accord with the *Moral Debt Principle*, then one is required to discharge that debt by compensating those who are harmed to the greatest extent possible insofar as one can do so at reasonable cost.

[21] A classic discussion of 'side constraints' in this sense is Robert Nozick, *Anarchy, State, and Utopia*, pp. 28-30. For a more subtle and comprehensive account, see Judith Jarvis Thomson, *The Realm of Rights*.

[22] Walter Sinnott-Armstrong does not consider anything like this principle in his arguments. To make this principle maximally defensible, one might add, as in Broome's discussion, that one knows that contributing to that harm will cause serious and incompensable harm as a predictable result of that avoidable act done for one's own benefit, etc.

If we assume that offsetting offers a reliable and inexpensive way for individuals to reduce their carbon footprints to zero, then it follows from these two principles and the actual facts about climate change that individuals are required to reduce their carbon footprint to zero.

However, this argument fails for reasons similar to those that explain the failure of Broome's actual argument. The problem is that a *stringent* duty to compensate is most plausible in *canonical harm cases* where a stringent side constraint is violated by one's contribution, and perhaps also when one's contribution makes particular individuals noticeably worse off – but a stringent duty to compensate is *not* plausible in cases where one's contribution makes no difference to any individuals and also does not violate stringent side constraints, as in many accumulative effects cases. Instead, the following principle is more plausible in light of such cases:

> *Defeasible Reasons to Compensate Principle*: If one acquires a moral debt in accord with the *Moral Debt Principle*, then one is required to discharge that debt by compensating those who are harmed to the greatest extent possible insofar as one can do so at reasonable cost, *unless there is something else one can do instead that is much better from an ethical point of view – especially insofar as such an alternative is (a) essentially a Pareto improvement over the option that is best from a compensation perspective in the sense that it makes some individuals much better off without making any individuals noticeably worse off than the compensation option, (b) leads to a much better outcome overall than the compensation option, and (c) doesn't itself involve any violations of side constraints, etc.*

The intuitive idea here is that if one acquires a moral debt, it is sometimes permissible to discharge that debt by doing something that is better from a moral point of view than the thing that would provide the greatest compensation to those owed the debt. This is especially difficult

to deny when, as the *Defeasible Reasons*… principle claims, the thing that is better from a moral point of view also does not leave anyone noticeably worse off than the option that provides the greatest compensation.

But if we reject the *Stringent Duty.*.. principle and endorse this *Defeasible Reasons.*.. principle (as we should), it follows that individuals should *not* reduce their carbon footprints to zero by purchasing offsets, because purchasing those offsets would not do any noticeable good for any individual and would have the opportunity cost of failing to do much more good for many individuals. This is true because charities that offset emissions do much less good with each dollar received than many other charities, such as those that treat malaria, provide clean drinking water to those who lack it, and many others – as Broome and other experts agree.[23] As a result, because we should accept the *Defeasible Reasons.*.. principle rather than the *Stringent Duty.*.. principle, we should reject the conclusion that individuals should reduce their carbon footprints to zero by purchasing offsets even if we accept something like the *Moral Debt Principle*, because a much better way to discharge those climate-related debts is by giving to charities that do as much good for the world as possible, which turn out to be charities that have nothing to do with climate change.[24]

The upshot is that there is no sound deontological argument that you are required to make substantial emissions reductions or reduce your carbon footprint, because you do not have a stringent duty to compensate others for your emissions – and insofar as you *are* in debt to others as a result of your emissions, it is not the sort of debt that involves a stringent duty to compensate, and so you can discharge that debt by giving to charities that have nothing to do

[23] Broome, *Climate Matters*, pp. 65-7 and 91-2. See also Bjorn Lomborg ed., *Global Crises, Global Solutions*, second edition, Posner and Weisbach, *Climate Change Justice*, chapter 4. [others: givewell].
[24] See also Dan Shahar, "Treading Lightly on the Climate in a Problem-Ridden World".

with reducing your carbon footprint, and that accomplish much more good for the world per dollar invested than can be accomplished through emissions offsetting.

This completes my general response to deontological objections to the futility argument for the conclusion that you are not required to make substantial emissions reductions. However, a number of important objections remain, which I examine in the following sections. More generally, the discussion in this section helps to clarify the subtle contours of deontological reasons in collective action situations by clarifying the reach and limits of the futility argument in such situations. The upshot is that although there appear to be many genuine sources of deontological reasons for action in collective action situations, none of them imply that an individual is required to make costly emissions reductions as things stand, and they do not even imply that individuals are required to make non-costly reductions to their carbon footprint via offsetting.

4. Objections to the Futility Argument based on the Aggregation of Imperceptible Effects

A different kind of objection arises from the idea that large numbers of imperceptible harms can be just as bad as smaller numbers of substantial harms. In the case of climate change, the important objection based on this idea is that you are required to reduce your emissions because the aggregate harm caused by your personal emissions, via many imperceptible harms to others, outweighs the significant good your emissions do for you and others. For example, according to many calculations, the lifetime emissions of an average individual in a developed nation 'adds up' to shortening a single life by several months.[25] In light of this, it might seem that individuals

[25] *Climate Matters*, pg. 74. See also John Nolt, "How Harmful are the Average American's Greenhouse Gas Emissions?"

are required to offset their emissions if they can do so at low cost, regardless of the philosophical quibbles above.

But here again, it is important to note that this conclusion has the implausible implication that individuals are required to purchase offsets even if they know that their money could do much more good elsewhere than if offsets were purchased without making anyone perceptibly worse off and without violating any side constraints. What reason could there be for thinking that individuals are required to waste the opportunity to do much more good in such a way? Clearly, such a conclusion would not follow from a traditional consequentialist perspective. Presumably, the argument must be that offsets are required on deontological grounds even though they will do much less good. But as the discussion above makes clear, this is a claim sorely in need of an argument.

Setting that fundamental problem aside, to make further progress on the independently important issue of evaluating the significance of imperceptible effects, it is important to distinguish two different issues. One issue is whether imperceptible harms *count for something*, while another issue is whether imperceptible harms *count for enough to outweigh other considerations*. In the climate case, this translates into the issues of *whether the imperceptible effects of your emissions provide genuine reasons to reduce your emissions*, and the further issue of *whether the imperceptible effects of your emissions provide strong enough reasons to make such reductions required, outweighing the good that those emissions do for you and others*. Previous critical discussions of arguments from imperceptible effects – such as Walter Sinnott-Armstrong's – tend to give quick arguments for the conclusion that imperceptible effects do not count for anything, generally on the grounds that only welfare effects, side constraints, and other familiar considerations provide genuine reasons for action, and that imperceptible effects make

no difference to any of these. In contrast to such a quick response, my approach is to grant that imperceptible effects count for something, but then provide an argument that even from such a perspective it is not plausible that they count for enough to generally outweigh the reasons you have for your emissions-generating activities. More generally, I argue that it is implausible that they aggregate in anything like the simple additive way that is assumed by Jonathan Glover, Derek Parfit, John Broome, and many other commentators on imperceptible effects and collective action.

Assuming then that imperceptible effects count for something – and with the example of climate change in the background along with the claim that the lifetime emissions of an average individual in a developed nation 'add up' to shortening a single life by several months – to make further progress it is useful to focus on two further issues:

(1) Aggregation of Good Issue

Is (a) just as bad as (b):

(a) shortening everyone's life imperceptibly in a way that 'adds up' to shortening a single life by two months given a simple aggregation principle

(b) shortening a single person's life by two months

(2) Wrongness of Harm Issue

Other things equal, is choosing (a) just as wrong as choosing (b)? – i.e., does (a) involve just as serious a violation of deontological constraints as (b)?

Even if we assume for a moment that Glover, Parfit, and Broome are correct about (1) and thus that (a) is just as bad as (b), it is important to see that the argument that emissions are impermissible without offsetting does not go through unless one also agrees with commentators like Broome about (2) – namely, that (a) involves just as serious a violation of deontological

constraints as (b). Otherwise, if (a) does not involve a violation of deontological constraints, then there is no reason why individuals must refrain from doing (a) or even compensate those who are imperceptibly harmed by (a), since they could do much more good by doing (a) and also giving to non-climate charities that are a much more effective way of doing good than purchasing emissions offsets.

Here it is worth noting that in addition to the arguments in the previous section, in the current context there is a simple way to object to Broome's view about (2) even assuming that the Glover/Parfit/Broome view is correct about (1). To see it, suppose you are forced to choose between (a) and (b). In such a case, it is obvious that you should choose (a) rather than (b). This shows that it is implausible that (a) involves as serious a violation of deontological constraints as (b). Furthermore, there is another impressive but independent objection to Broome's view about (2): namely, (a) involves de minimis *infringements* that do not constitute impermissible rights *violations*, whereas (b) involves a serious violation of stringent rights; therefore, contrary to Broome's view about (2), (b) involves more serious violation of deontological constraints than (a). Either way, the upshot is that principles about the aggregation of harm do not support Broome's view about (2) and thus do not support the conclusion that emissions are impermissible without offsetting.

Setting aside (2), it is important from a philosophical point of view to evaluate the Glover/Parfit/Broome view about (1), according to which ordinary people who think (a) is not as bad as (b) are making a 'mistake in moral mathematics'.[26] Here it is important to note that, despite the rhetoric of 'mistakes in moral mathematics', some thoughtful philosophers agree with ordinary people who make this 'mistake' and disagree with the Glover/Parfit/Broome view. For

[26] See Derek Parfit, "Five Mistakes in Moral Mathematics", in *Reasons and Persons*, following Jonathan Glover, "It Makes No Difference Whether or Not I do It".

example, in a passage quoted in the previous section, Joel Feinberg argues that accumulative harms such as (a) must be addressed differently than canonical harms such as (b), and his analysis entails the negation of both the Glover/Parfit/Broome view about (1) and Broome's view about (2). As Feinberg might say: *elastic as it may be, the notion of harm cannot be stretched as far as is needed by such arguments*.[27]

Even setting aside Feinberg's substantive view about accumulative effects, it is possible to offer many examples that illustrate the implausibility of the sort of simple additive aggregation principles assumed by the Glover/Parfit/Broome view that imply that (a) is just as bad as (b). In what follows, I offer one such example, and then I offer a more theoretical diagnosis of the mistake behind such principles.

So, suppose that 10,000 men are currently set to be tortured with intensity i for the next hour, and that you must choose between either increasing the intensity of their pain by an imperceptible amount (whatever that means – I set aside such worries here), or else causing another man, who is currently set not to be tortured, to also be tortured with intensity i for the next hour along with the 10,000 men. As you make your choice, one question is whether there is a reason that counts against burdening the men with an additional amount of imperceptible pain; however, even if there is such a reason, a further question is whether that reason is outweighed by the reason you have not to burden an additional innocent person with an hour of intense torture that he would not otherwise experience.

In this case, it is clear that you must choose to spare the single innocent man from torture, even though that implies imposing on each of the men an imperceptible amount of additional pain. On reflection, this undermines simple additive aggregation principles for imperceptible

[27] Broome and Parfit do not address the large literature on harm by Feinberg and others. The reference here is to the last line of Feinberg's *Harm to Others: The Moral Limits of the Criminal Law*, pg. 245.

effects, as well as the idea that imperceptible harms caused by emissions could aggregate in a way that provides strong reasons against those emissions. To see why, suppose that the torture will be delivered by a device that has adjustable dials that determine how much pain is inflicted on each subject. Each subject's dial can be set anywhere between 0 and over 10,000, and the settings have been carefully designed so that each increment delivers an imperceptible amount of additional pain, with 0 delivering no pain, and 9,999 delivering the fairly intense torture that the men are antecedently set to endure. So, as you make your decision, each man's dial is set to 9,999. If you choose to spare the single innocent man, you cause 10,000 additional units of pain to be inflicted on the others (distributed equally and imperceptibly among the men, one unit each); however, if you choose not to spare that man and have him tortured along with the other men, you cause only 9,999 additional units of pain to be inflicted (all inflicted on the single man). We can also suppose that the increments have been carefully designed so that each additional increment results in linearly decreasing utility in the subjects, insofar as it is possible to determine a utility curve for pain in such a case. Even in light of all these details, it is obvious that you should choose to distribute the pain imperceptibly among the group of men rather than impose intense torture on one man. Furthermore, this remains true even if we increase the number of men who would be affected to 100,000 or more, given that your choice is still between causing someone to be tortured who will not otherwise be tortured and merely doing something that will have no perceptible effect on anyone. This shows that it is false to claim that imperceptible harms aggregate in anything close to a simple additive way. Instead, the case shows that on any plausible view, inflicting substantial harm on particular people who would not otherwise be harmed is other things nearly equal much worse than inflicting an equal 'sum of harm' that is distributed imperceptibly among many people. In light of this, it is hard to see how

there could be a plausible argument that the imperceptible harms caused by your emissions could outweigh the important goods that those emissions bring to you and others.

Apart from appealing to clear intuitions about cases, it is also possible to offer a straightforward diagnosis of where simple additive aggregation principles go wrong insofar as they imply that (a) is equally as good as (b). As the example just given illustrates, the underlying problem is that such principles assume that distributional facts about welfare don't matter, which is just a characteristic utilitarian idea dressed up as an insightful view about aggregating imperceptible effects. Like the utilitarian view itself, such a view can seem attractive when it is deployed to support antecedently attractive conclusions, but it is ultimately just as implausible here as is the utilitarian view elsewhere, especially in connection with distributional issues.

In support of this diagnosis, it is important to see that rejecting the Glover/Parfit/Broome view about (1) is consistent with thinking that imperceptible effects count for something. In the literature, it is typical for philosophers to slide from reasons for thinking that imperceptible effects count for something to the unwarranted conclusion that imperceptible effects aggregate in the simple additive way assumed by the Glover/Parfit/Broome view about (1). However, if such simple aggregation principles were correct, then it would be a matter of indifference for a person who has to steal in order to survive whether he steals all of what he needs from one person, thereby making that person seriously worse off, or instead steals what he needs from a very large number of people in a way that does not noticeably harm anyone but imposes the same 'aggregate sum of harm' on them. As far as I can tell, this consequence of such simple aggregation principles is just as perverse as the other apparent unintuitive consequences of utilitarianism that commentators like Glover and Parfit originally attempted to explain away by means of simple additive aggregation principles – and by the same token it shows that such

principles cannot really explain the entire range of cases that such consequentialist philosophers intend them to explain. This seems to undermine any reason for favoring such principles even from a traditional consequentialist perspective.[28]

5. Objections to the Futility Argument based on the Possibility of Threshold Effects

One of the most depressing facts about climate change is that greenhouse gasses are currently *accumulating* in the atmosphere, and will continue to accumulate into the foreseeable future even if emissions are reduced by as much as 50% from current levels, because the rate of emissions is much greater than the rate at which greenhouse gasses are removed from the atmosphere by natural processes.[29] In light of this, it would be a mistake to equate the ethics of climate change

[28] For important arguments that imperceptible effects do count for something, see Jonathan Glover, "It Makes No Difference Whether or Not I Do It", and Derek Parfit, *Reasons and Persons*, pp. 75-82. Glover and Parfit slide from reasons for thinking that imperceptible effects count for something to the unwarranted conclusion that imperceptible effects aggregate in the simple additive way assumed by their view about (1). For example, Glover advocates a simple additive aggregation principle that he calls 'The Principle of Divisibility' (see pp. 174-175). For further discussion of issues about imperceptible harms, aggregation, and analogous distributional issues, see Larry Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*, Shelly Kagan, "Do I Make a Difference?", and Julia Nefsky, "Consequentialism and the Problem of Collective Harm: A Reply to Kagan".

[29] From the authoritative Intergovernmental Panel on Climate Change (IPCC) report *Climate Change 2007: The Physical Science Basis*: "The concentration of greenhouse gas in the atmosphere depends on the competition between the rates of emission of the gas into the atmosphere and the rates of processes that remove it from the atmosphere. For example, carbon dioxide (CO_2) is exchanged between the atmosphere, the ocean, and the land through processes such as atmosphere-ocean has transfer and chemical (e.g. weathering) and biological (e.g. photosynthesis) processes. While more than half of the CO_2 emitted is currently removed from the atmosphere within a century, some fraction (about 20%) of emitted CO_2 remains in the atmosphere for many millennia. Because of slow removal processes, atmospheric CO_2 will continue to increase in the long term even if its emission is substantially reduced from present levels. … More specifically, the rate of emission of CO_2 currently greatly exceeds its rate of removal, and the slow and incomplete removal implies that small to moderate reductions in its emissions would not result in stabilization of CO_2 concentrations, but rather would only reduce the rate of its growth in coming decades. A 10% reduction in CO_2 emissions would be expected to reduce the growth rate by 10%, while a 30% reduction in emissions would similarly reduce the growth rate of atmospheric CO_2 concentrations by 30%. A 50% reduction would stabilize atmospheric CO_2, but only for less than a decade. After that, atmospheric CO_2 would be expected to rise again as the land and ocean sinks decline owing to well-known chemical and biological adjustments" (pp. 824-825). For further ramifications for greenhouse gas stabilization levels, see *The Economics of Climate Change: The Stern Review*, Chapter 8; for a clear summary, see Hal Harvey and Sonia Aggarwal, "The Costs of Delay", ClimateWorks Foundation.

with the ethics of 'voting cases' in which your actions have a very small chance of making a dramatic difference, as in the following argument:

> There is some danger that we will cross an emissions tipping point that will lead to a catastrophe; so, unless you reduce your emissions by a significant amount, there is some chance that you will make such a catastrophe happen. Therefore, you are required to reduce your emissions by a significant amount.

The problem with this argument is that it ignores the fact that greenhouse gas levels are currently accumulating, and will continue accumulating into the foreseeable future. To see why this is a problem for the argument, imagine that this week, as you engage in some emissions-generating activity, your emissions cause a catastrophic tipping point to be crossed. Nonetheless, even if you had avoided those emissions and thus hadn't tipped the scales yourself, it is certain that someone else's emissions would have tipped the scales at essentially the same time, because an entire planet of other people would still have been emitting at the same time even if you had not been emitting. This shows that, given the empirical facts, there is no chance that you could delay a catastrophic tipping point from being crossed today or in the foreseeable future by reducing your emissions, and thus there is no good reason for reducing emissions that arises from the possibility of tipping points being crossed now or in the foreseeable future.[30]

This shows that it is a mistake to think that the ethical issues regarding climate change are essentially the same as the ethical issues regarding voting in an election. That is because in an

[30] More precisely: there is no chance that you could delay a tipping point from being crossed today, *at least not by any perceptible amount of time*. Here it is worth noting that catastrophic tipping points associated with climate change are *global* in the sense that they are a function of emissions that tend to be dispersed quickly around the entire earth, which means that it is a mistake to believe that some particular, localized facts about your emissions could conspire to cause a catastrophe that would otherwise have been avoided. It might also be objected that we should focus on your emissions over your entire lifetime, rather than on your emissions on a particular day. This wouldn't affect my point, since a reasonable individual's emissions over an entire lifetime would still be much less than the emissions happening elsewhere on the planet in any single second.

election a crucial consideration is that your vote has some chance of tipping the scales and making a dramatic difference to the outcome, whereas a depressing fact about climate change is that as things stand now and as they will stand into the foreseeable future, your personal emissions have no real chance of making a difference to the outcome, except by making things substantially worse for yourself if you choose to make reductions.

There are also further important differences between voting and climate change. For example, it is not costly to vote in an election, whereas it is costly to reduce your emissions by a substantial amount; in addition, it is plausible that you have important deontological reasons to vote, whereas you do not have important deontological reasons to reduce your emissions by a substantial amount (as suggested above). This shows that the considerations relevant to the ethics of climate change are importantly different than the considerations relevant to the ethics of voting, and that the challenge of explaining how individuals could be required to vote in an election is an easy problem compared to the challenge of explaining how individuals could be required to reduce their greenhouse gas emissions by a substantial amount.[31]

One important qualification to all of this is that there is some chance that in the future humans will change their behavior dramatically (i.e. there is some chance that in the future humans will reduce emissions by more than 50% from current levels), thereby creating a situation in which greenhouse gasses are being removed from the atmosphere faster than they are added. If such a scenario ever obtains, then the argument above would not fail for the main reason just given. However, to argue that individuals are currently required to reduce emissions by a substantial amount because of the possibility of such an unlikely change in the future is like

[31] To say it is an easy problem compared to another problem is not to suggest that anyone has ever offered a convincing argument that citizens are required to vote in elections, especially from a utilitarian perspective. (Thanks to Michael Cox and Richard Yetter Chappell for encouraging me to discuss the connection between the ethics of climate change and the ethics of voting.)

arguing that in an election scenario you would be required to go to the polling place and attempt to cast a ballot even if the polls were already closed and the costs to you were high, because there would still be *some* chance that precedent would be reversed and somehow your vote would count, and thus there would be *some* chance that your vote would end up making a dramatic difference. This shows that insofar as your situation with respect to climate change can be accurately described as analogous to a voting case, it is analogous to a voting case in which you are clearly not required to vote.[32]

The most decisive way of showing that tipping point reasoning is misguided is to note that it overgeneralizes in absurd ways. The basis for this observation is the empirical fact that when emissions cause harm, the most direct cause of that harm in the causal chain is always something like sea level rise, temperature and weather changes, or changes in the risk for other health and welfare issues such as melanoma. In light of the effect that a single individual has on these more direct causes via his or her emissions, and how that effect compares to his or her impact on such things via non-emissions activities, it is hard to see how an individual's emissions could be seriously objectionable as things stand on the basis of these effects. For example, the probability of causing a climatic catastrophe by your personal emissions is roughly the same as the probability of causing a climatic catastrophe by putting a boat in the ocean or having a cup of tea every day, given that the effects of your emissions on the potential causes of climatic catastrophe such as sea level rise and air temperature increase are about the same as the effects on those things of taking a swim in the ocean and boiling water for tea every day, respectively. Similarly, the increased risk of cancer imposed on others by your personal emissions is similar to the increased risk of cancer imposed on others by your use of a cellphone, hairspray, and many other everyday products. However, it would be absurd to think that putting a

[32] Thanks to Joe Rachiele for encouraging me to discuss the issues in this paragraph.

boat into the ocean and boiling water for tea are morally wrong because of the chance that you might thereby cause a climatic catastrophe, just as it is absurd to think that that using hairspray and other everyday products is wrong because of the *de minimis* increased risk of cancer that you thereby impose on others. In other words, if it is impermissibly risky not to reduce your emissions by a substantial amount because of the chance that you will cause a tipping point to be crossed, then it is hard to see why it is not impermissibly risky to put a boat in the ocean or have a cup of tea every day, since a similar risk to others is involved, and the potential benefit to yourself is of no greater significance. But it is absurd to think that putting a boat in the ocean and having a cup of tea are impermissible because they risk climatic catastrophe. This seems to show that there is something irreparably with the tipping point reasoning under consideration, as well as the other simple harm- and risk-based arguments canvassed above.[33]

6. The Average Effects Fallacy

Another familiar objection arising from the arguments like the following:

> If we do nothing about climate change, then at least millions of people will die; so, unless you reduce your emissions by a substantial amount, then the *harm footprint* associated with your emissions will be very great. Therefore, you are required to reduce your emissions by a substantial amount.

On its most natural interpretation, this argument commits the *Average Effects Fallacy*, which is the fallacy of equating the ethically relevant effects of a particular act with the average effects of

[33] Thanks to Jane Willenbring for helping me appreciate the importance of this point.

all of the actual acts of that type.[34] To illustrate the fallacy, imagine that you are leaving a stadium after an event. As you are about to cross the street and walk toward the parking garage, you hear an ambulance. You initially stop and are inclined to wait if other people do, but you see hundreds of other people cross the street in front of the ambulance, and based on your knowledge of the situation you know that your crossing the street will not make any difference, given that there is no way you can convince the crowd to wait for the ambulance. In light of this, you cross the street to get to your car, because waiting for the entire crowd to disperse would be a significant inconvenience for you, and would not do any good for anyone else. Because so many people cross the street in front of the ambulance, the ambulance is delayed, and the patient inside dies, who would otherwise have lived. In this case, we can suppose that the average effect of an act of crossing the street is negative on balance, because the good that accrues to all of the people as a result of getting to their cars faster is greatly outweighed by the harm that is caused by their collective action. Nonetheless, it is permissible for you to cross the street, because the outcome only would have been worse if you had not crossed the street, because then no good would have accrued to you and every other relevant feature of the outcome would have been the same. The explanation is that the ethically relevant effects of a particular act are the marginal effects of that particular act, and not the average effects of acts of that type – and perhaps also that it is not your fault that others are acting this regrettable way in the background. (Similar remarks apply in

[34] In general, the *x footprint* of an act of a particular type is simply the *average effect* of all actual acts of that type for some particular kind of effect x. One version of the Average Effects Fallacy is simply to equate the actual effects of an act with the average effects of actual acts of that type. Many forms of carbon footprint reasoning commit this version of the fallacy. For example, in *The Food Revolution*, John Robbins claims that a pound of beef raised in California has a larger water footprint than taking a shower each day for six months; from this, he concludes that you can save more water by not eating a pound of beef than by not showering for six months (Tenth Anniversary edition, pp. 236-237). Robbins's conclusion is false and his argument is invalid because foregoing a pound of beef will not actually affect the amount of water that is consumed, whereas foregoing a shower for six months would affect the amount of water that is consumed. Robbins's mistake is to equate the actual effects of particular acts with the 'footprint' associated with those acts, where footprints are simply a measure of the *average effects* of acts of that type. In other words, his mistake is to commit the Average Effects Fallacy. I discuss other problems with carbon footprint reasoning in a subsequent footnote in connection with an analogous argument due to Michael Pollan.

more dramatic fashion to the asteroid and stampede cases described in the first sections of this paper.)

This helps to explain why the argument displayed above is unsound. On its most natural interpretation, the problem is that the conclusion does not follow from the premise about your harm footprint, given that your 'harm footprint' must be interpreted as meaning, roughly, 'the average effect of failing to reduce emissions' if that premise is to be plausible. The argument is initially seductive because we are often disposed to commit the Average Effects Fallacy in our reasoning by equating the ethically relevant effects of an act with the average effects of actual acts of that type, rather than with the marginal effects of the act itself.[35]

The street-crossing case also illustrates the mistake made by causal consequentialist theories that proportion credit for consequences based on agents' causal responsibility for those consequences. The street-crossing case is a counterexample to such views because, given suitable assumptions about that case, you are as causally responsible for the outcome as every

[35] To reinforce the idea that we should focus on marginal effects rather than average effects, it is worth noting that even from a collective perspective we should focus on marginal effects rather than average effects. For example, consider a policymaker who is deciding whether to advise people to grow vegetables in their backyard, and suppose it is clear that individuals would lower their carbon footprint by growing vegetables in their backyard, because the average emissions associated with store-bought vegetables are higher than the average emissions associated with vegetables grown in a backyard. At first glance, it can seem obvious that the policymaker should advise people to grow vegetables in their backyard in light of these facts if the only relevant goal is to reduce emissions. However, that does not follow, and to think it does follow is to commit a version of the Average Effects Fallacy. To see the problem, consider the realistic assumptions that only a small percentage of the population will actually follow the advice, and that the followers will be fairly evenly distributed with respect to the supermarkets where they currently buy vegetables; if so, the actual effect of such advice would be to increase emissions, because each follower will generate some new emissions in the course of creating and tending their garden, and any collective positive effects will be negligible compared to the collective negative effects, because the followers are dispersed and few in number, thereby having no important effect on the production of vegetables and hence emissions outside their backyards, because their dispersed consumption acts will not make a difference further up the supply chain. For such reasons, advising a population to reduce their carbon footprint in some way will often predictably lead to greater overall emissions, even if everyone who follows the policy succeeds in lowering their carbon footprint, and even if none of the non-followers increase emissions or their carbon footprints. This illustrates why it is crucial to focus on marginal effects rather than average effects at every level of individual and collective abstraction – and why decision makers at all levels should recognize that carbon footprints and lifecycle analysis provide unreliable evidence about what policies would have the best effects on emissions levels. (Michael Pollan commits the Average Effects Fallacy in his article "Why Bother?" when he concludes that growing your own vegetables would have good environmental effects because you would thereby reduce your carbon footprint.)

other person (imagine that everyone else engages in similar reasoning before crossing the street), and therefore your proportional share of the harm that is done is unacceptably negative according to a causal consequentialist view. Nonetheless, there is nothing wrong with your decision to cross the street, which shows that such views are mistaken. (A similar point can be made using the asteroid and stampede cases discussed in the first sections of this paper.)[36]

In response to all of this – and to avoid committing the Average Effects Fallacy – the argument displayed above could be understood in terms of the *marginal effects* of your emissions-generating acts rather than the *average effects* of acts of that general type. However, if the argument is understood in terms of marginal effects, then it mistakenly assumes that the harm that would actually be avoided by you reducing your emissions is more significant than the harm that would be imposed on yourself and others by such reductions – which is implausible, as the discussion in the previous sections revealed.

7. Virtue Ethics and the Futility Argument

Another kind of objection begins with the idea that (roughly) an act is required if and only if a virtuous person would perform that act, and then argues that you are required to reduce emissions, because reducing emission is what a virtuous person would do.

However, upon reflection it is implausible that a virtuous person would have decisive reasons to reduce her emissions by a significant amount, because a virtuous person would not necessarily sacrifice important aspects of her own well-being and the well-being of others in a way that does not benefit anyone else or serve any other weighty purpose, and a virtuous person

[36] For defense of causal consequentialist views, see Alvin Goldman "Why Citizens Should Vote: A Causal Responsibility Approach", and Dan Moller "The Morality of Overdetermination: Wrongdoing When it Makes No Difference", in *Abortion, Killing, and Overdetermination – Three Essays*, PhD dissertation, Princeton University.

would not act on the basis of mistaken reasoning about what is an effective way of combating climate change. For example, if you reduce your emissions by canceling family vacations and family trips, that imposes significant costs on your family, and does not really help anyone else, and is not an effective way of contributing to a solution to climate change. And if you walk or ride your bike to work every day, you impose a non-trivial risk of catastrophe on your family and friends, because there is a much greater chance that you will be killed in an accident if you walk or bike than if you drive – and, again, imposing this risk on your family does not really benefit anyone else, and is not an effective way of contributing to a solution to climate change. Here it is important to note that although unilateral individual action to reduce emissions would have some indirect effects, it is not plausible that the expected effect on the political process is anywhere near significant enough to make emissions reductions required for that reason. Furthermore, as noted above, a virtuous person would not necessarily choose to purchase offsets for her emissions, because she would recognize that there is no stringent duty to do so, and would recognize that she could do much more good by investing those resources elsewhere. In light of all this, a virtuous person would not impose such costs and risks on herself and others in order to reduce emissions, and she would not invest her resources in an unmotivated and ineffective way rather than investing them in very effective ways instead. Instead, a virtuous person would recognize a duty to support just institutions and thus favor and support promising solutions to climate change at the level of public policy, because such solutions are the only ones that have any real prospect of success, and she would give an appropriate amount to the sort of charities that promise to do the most good for those in need, which are all non-climate charities.

A related objection is that you are required to reduce emissions by a significant amount because you should want to be the kind of person who expresses something like a *symbolic*

protest of the unfolding climate catastrophe, even if such reductions are costly to you and others and cannot be expected to have any other important effects. There are two main problems with this objection. The first problem – and a problem for the previous objection from virtue ethics – is that even if we assume you have *sufficient* reason deriving from symbolic value or some other source of virtue to perform the futile but symbolic act of reducing emissions in a way that is very costly to you and others, that does not make it plausible that such an act is *required*, because it is still much more plausible to think that such an act would be *supererogatory*. This is true because although it is plausible that there are genuine reasons for action that arise from considerations of symbolic value, a general fact about such distinctively virtue-theoretic reasons is that they are not generally of sufficient weight to make such acts *required*, as opposed to merely *rational* and *supererogatory*. Of course, in some cases where it is particularly *essential* to an act that substantial and serious violations of rights lie behind it, it is arguable that such reasons could make symbolic protest required.[37] However, it is not in the nature of everyday emissions-generating activities that they *essentially* involve substantial and serious violations of rights, because as noted above no substantial violations of rights lie behind them, and in any event such activities would be completely unobjectionable and in fact highly desirable if only there was a basic adequate regulatory background for such actions, such as a global carbon tax.[38]

Another objection is that the futility argument must be mistaken because it implies that hypocrisy is permissible. However, the futility argument does not obviously have that

[37] I argue for something very much like this view, based on *degrees of essentiality* of harm and the violation of rights, in my paper "The Ethics of the Marketplace and a Surprisingly Deep Question for Normative Theory: What are Consumers Required to Do When Products are Produced in Morally Objectionable Ways?". For a related view with many important connections to human rights issues, see Chiara Lepora and Bob Goodin, *On Complicity and Compromise*.

[38] Thanks to Cheshire Calhoun, Eamonn Callan, Jeff Downard, and Michael Smith for encouraging me to consider this kind of objection. For related discussion, see Thomas Hill, "Symbolic Protest and Calculated Silence". In connection with Hill's paper, it is worth noting that his examples tend to rely on cases in which individuals are intuitively required to "disassociate from irreparably corrupt groups" (92). However, there is no clear analogy between such cases and our situation with respect to climate change. See also Julia Driver, *Uneasy Virtue*.

implication. To see why, consider a serious tragedy of the commons in which your family must choose between either using the commons to ensure their own basic flourishing, or else foregoing use of the commons in a way that is certain to be both very costly and futile, because the commons will be destroyed by the overuse of others regardless of your family's actions. What is the correct and virtuous thing to do in such a situation? On the one hand, it is true that your family and every other family should favor an effective solution to the tragic dilemma, perhaps at the level of public policy. On the other hand, if you know that such a solution is not forthcoming and that the commons will be destroyed regardless of what your family does, then it would be wrong for you to sacrifice your family in a way that is certain to be futile by foregoing use of the commons. In such a situation, you are not a hypocrite if you say "As things stand, I am not required to sacrifice the well-being of my family, even though I favor and support a course of action in which everyone makes sacrifices for everyone's benefit, perhaps by means of mutual coercion, mutually agreed upon". Such a statement expresses correct judgment, not hypocrisy – and is a straightforward analogue of the view about the individual ethics of climate change implied by the futility argument, and arguably also by a Rawlsian view of individual duties in the face of tragically inadequate institutions.[39]

[39] The phrase "mutual coercion, mutually agreed upon" is from Garrett Hardin, "The Tragedy of the Commons". (Thanks to Gideon Rosen for encouraging me to discuss hypocrisy.) In an oft-quoted passage in *A Theory of Justice*, Rawls writes "From the standpoint of the theory of justice, the most important natural duty [of individuals] is that to support and to further just institutions. This duty has two parts: first, we are to comply with and do our share in just institutions when they exist and apply to us; and second, we are to assist in the establishment of just arrangements when they do not exist, at least when this can be done with little cost to ourselves" (revised edition, pp. 293-4). I assume that the qualifier "From the standpoint of the theory of justice…" makes it controversial how Rawls's view would apply here. Compare also the following from Peter Railton: "By altering social and political arrangements we can lessen the disruptiveness of moral demands on our lives, and in the long run achieve better results than freelance good-doing. A consequentialist theory is therefore likely to recommend that accepting negative responsibility is more a matter of supporting certain social and political arrangements (or rearrangements) than of setting out individually to save the world. Moreover, it is clear that such social and political changes cannot be made unless the lives of individuals are psychologically supportable in the meanwhile, and this provides substantial reason for rejecting the notion that we should abandon all that matters to us as individuals and devote ourselves solely to net social welfare", "Alienation, Consequentialism, and the Demands of Morality" in *Facts, Values, and Norms*, pg. 172.

8. Climate Change Ethics, the Ethics of Collective Action, and Public Discourse

At this point, I have replied to a number of objections to the futility argument for the conclusion that as things stand you are not required to reduce your emissions by a substantial amount. None of these replies depend on our prior ignorance of the facts about climate change, our uncertainty about cause and effect relationships, our uncertainty about the long-run effects of climate change, complications regarding the identity of persons in future generations, or agent-relative personal prerogatives to give more weight to our own projects and the projects of those to whom we stand in morally significant relations.[40] Instead, the arguments here suggest that even if we ignore those issues there is still decisive reason to think that the main objections to the futility argument in connection with climate change considered above are mistaken. If these additional considerations are taken into account, then there appears to be even more reason to think that the conclusions defended above must be correct.

At the same time, it is important to note that nothing like a conclusive proof of any conclusion has been offered here. Serious consideration of the futility argument is a very recent event, which has not yet run its course. We cannot yet predict whether, as in Mathematics, some previously unconsidered argument might prove to be decisive.[41] For example, many have argued that at the level of public policy, climate change ethics is largely a matter of human rights. Although important barriers have been identified above to extending such an approach in a way

[40] Compare Samuel Scheffler, *The Rejection of Consequentialism*.
[41] Compare Derek Parfit's concluding remarks in *Reasons and Persons*, pg. 454.

that implies that individuals must unilaterally reduce emissions, it would be premature to conclude that such a project has no prospect for success.[42]

With that proviso in hand, the most sophisticated version of the futility argument in connection with climate change might invoke the following principle:

Extended Non-Cooperation Principle

If an individual faces a collective action problem in which she knows that most others will fail to choose the 'cooperative' option, thereby ensuring a collectively bad outcome, she is not required to choose the cooperative option herself if she knows that:

(i) choosing the 'cooperative' option would involve her making significant sacrifices,

(ii a) there are no significant welfare-based reasons for her to choose the 'cooperative' option, including reasons that arise from aggregating imperceptible effects and the possibility of triggering threshold effects,

(ii b) there are no significant deontological reasons for her to choose the 'cooperative' option, where these deontological reasons might include reasons based on rights and what we owe to others, including duties to promote just institutions, and perhaps other deontological reasons based on agent-relative values, principles such as the doctrine of double effect, the principle that acts are wrong when it is very essential to them that serious harm and rights violations lie behind them, and so on,[43]

(iii) most of the others would not choose the 'cooperative' option even if they were in her exact situation,

(iv) choosing the 'cooperative' option would not lead her to be part of a successful solution to a different collective action problem, and

[42] Thanks to Debra Satz for convincing me of the importance of this point. See for example Simon Caney, "Cosmopolitan Justice, Rights, and Global Climate Change", and David Miller, "Global Justice and Climate Change: How Should Responsibilities Be Distributed?", Tanner Lecture on Human Values, 24-5 March 2008, Tsinghua University.

[43] Again, this condition grants (without endorsing) the possibility of reasons of 'fair play' that (arguably) apply in the case of voting and successful provision of public goods. The standard conditions that give rise to such reasons are not present in the climate case.

(v) the only feasible way to secure a socially desirable outcome depends on public policy or some other solution that given the facts of the case, would not plausibly arise through unilateral individual 'cooperative' action.

Again, it is difficult to see how this principle could be false, because it captures what is most compelling about the intuitive idea that even if *free-riding* is impermissible in response to the general *cooperation* of others, at the same time one is not required to be a *sucker* by cooperating in the face of the general *non-cooperation* of others. More generally, the conditions within this principle form the basis for a needed distinction between, on the one hand, cases in which an act is genuinely required despite making no difference to the outcome, and, on the other hand, cases in which an act is not required because it is *completely futile* in the much broader sense involving satisfaction of (i) through (v). Furthermore, the discussion above clarifies and makes progress on many of the subtle challenges involved in applying these considerations, including those related to imperceptible effects and the possibility of threshold effects.[44] In this way, appreciating the arguments for this principle and clarifying its proper application to cases results in significant progress on the ethics of collective action.

Returning to the case of climate change, this principle implies that individuals are not required to reduce their emissions by a substantial amount as things stand because without an effective public policy response such reductions are costly and futile in a sense that makes them not required. At the same time, individuals are presumably required to favor and support just

[44] One important issue not discussed here but much discussed elsewhere is the extent to which traditional consequentialist views generate plausible verdicts on collective action situations simply in virtue of an appeal to imperceptible effects and/or the possibility of threshold effects. For a general argument that they do, see Shelly Kagan, "Do I Make a Difference?". For objections to Kagan, see Julia Nefsky, "Consequentialism and Collective Harm: A Reply to Kagan", and my paper "You Don't Make a Difference: The Inefficacy Objection and the Problem with the Singer/Norcross/Kagan Response". Kagan's arguments largely follow earlier arguments made by Jonathan Glover, "It Makes No Difference Whether or Not I Do It", Derek Parfit, *Reasons and Persons*, and Peter Singer, "Utilitarianism and Vegetarianism".

solutions to climate change and the institutions that go with them. It is crucial to distinguish this *political* act from the more *direct act* of reducing one's emissions unilaterally. With that distinction in mind, the proponent of the futility argument can agree that individuals are generally required do *something* about collective action problems; however, that something is often different than it initially appears – for example, the thing that individuals are required to do about climate change appears political, and not directly tied to direct action.

Is this distinction between direct action and political action important? Does it make a difference in practice? It does. To see why, consider a controversial remark by Barack Obama that might initially seem confused. During a 2008 debate, Brian Williams asked Obama what he was *personally* doing about climate change in his day-to-day life. Obama later described Williams's question as "stupid", telling his staff: "What I'm thinking in my head is, 'Well, the truth is, Brian, we can't solve global warming because I fucking changed light bulbs in my house. It's because of something collective'."[45]

The explicit content of Obama's statement is true: climate change is caused by "something collective", and so it cannot be solved by one man changing light bulbs in his house. However, Obama's point is undoubtedly something further. At first glance, it might seem that Obama is arguing that he isn't required to change light bulbs in his house because a person isn't required to make a contribution toward solving a problem if that contribution won't solve the problem all by itself. However, that is not the best way to interpret Obama's words. (For example, the stampede case near the beginning of this paper shows that such a view is mistaken: in that case, you pressing your button cannot solve the problem all by itself; nonetheless, you are required to press your button.)

[45] http://www.huffingtonpost.com/2008/11/05/obama-we-cant-solve-globa_n_141407.html

A better interpretation is that Obama's view is essentially the same as the view about the individual ethics of climate change articulated here: there is something that we as individuals are required to do about climate change, but it isn't to change light bulbs in our house, or to do other costly emissions-reducing actions at a personal level. Instead, each of us is required *to favor and support* promising public policy solutions to the problem of climate change because such solutions are the only ones that have any real prospect for success. Furthermore, insofar as we are inclined to go above and beyond what we are merely required to do, as individuals we should focus our energy on promoting public policy solutions and encouraging others to do the same, rather than focusing our efforts on changing light bulbs in our houses and other individual-level actions that are not an important part of a genuine solution to the problem. Williams's question mistakenly assumes that climate change is a problem to be solved by the direct action of individuals motivated by the recognition of individual obligations to reduce emissions, when in fact individuals have no such obligations, and would do better to focus on promoting collective-level solutions to climate change instead.

If this is Obama's view, then why doesn't he explicitly say it? Part of the explanation might be that the mistake made by Williams is subtle and could only be demonstrated by lengthy arguments such as those in this paper. But a more more tragic part of the explanation is that most left-leaning Americans have already concluded that Williams's way of thinking about the ethics of climate change is correct, and they reject without examination anything inconsistent with that view. As a result, Democratic politicians are afraid to question that view publicly, and instead ensure that their policies are consistent with and reward those who hold such views. This is tragic because if it continues, it will ensure that the problem of climate change is never successfully addressed. That is because if left-leaning Americans continue to believe that climate

change is a problem to be solved at an individual level by "doing the right thing", and continue to believe that they are "doing their part" by changing light bulbs and so on, then the collective upshot will be that left-leaning Americans fail to understand and fail to support a real solution to the problem (including a global public policy that creates a single global price for emissions of greenhouse gasses), and instead continue to direct their limited political capital toward solutions that conform to a mistaken view about the nature of the problem and the nature of effective solutions to the problem. And if left-leaning Americans never demand a solution to the problem at the level of public policy, then there will never be sufficient support for such a solution in America; and if there is never sufficient support for such a solution in America, then much of the rest of the world will recognize that costly action on their part would be futile, thereby ensuring that the problem of climate change is never successfully addressed. If anything like this is true – and empirical studies suggest that it is[46] – then the arguments in this paper have some practical

[46] For example, see Dena Gromet, Howard Kunreuther, and Richard Larrick, "Political ideology affects energy-efficiency attitudes and choices", 2013 *Proceedings of the National Academy of Sciences*, and Verena Tiefenbeck, Thorsten Staake, Kurt Roth, and Olga Sachs, "For better or worse? Empirical evidence of moral licensing in a behavioral energy conservation campaign", 2013 *Energy Policy*. This reasoning is not undermined by opinion polls, some of which claim to show that left-leaning Americans support some form of collective action on climate change. For one thing, supporting some form of collective action is not the same thing as supporting *dramatic* collective action: in particular, the relevant polls provide some reason to think that left-leaning Americans support a token incentive scheme (which would have no important effect), but do not provide any reason to think that left-leaning Americans support the more dramatic action that would be necessary to change the ultimate outcome in any meaningful way. Furthermore, nothing dramatic will happen unless left-leaning Americans *demand* such action, and there is no chance of that happening as long as they continue to conceptualize the problem as one to be solved largely by individuals "doing the right thing" and "doing their part" by changing light bulbs in their houses, etc. In any event, the results of even these 'token response' climate surveys vary widely because of massive framing effects that depend on whether the questions are framed in a way that emphasizes *costs*, or instead emphasizes *environmental benefits*; as a result, individual polls provide no real evidence for conclusions about individuals' 'real opinions' on token climate policy. (For some related discussion, see Cass Sunstein, *Laws of Fear*, pg. 48.) An additional consideration is that left-leaning Americans tend to demand action to counteract any substantial increases in gas prices, which arguably provides more reliable evidence about their true preferences regarding policy changes that essentially involve such consequences.

Although climate change will not be successfully addressed until left-leaning Americans demand dramatic action, that does not mean that individuals are required to demand such action as things stand now. Instead, individuals are required to favor demanding such action, and perhaps should also stand ready to demand such action conditional on a sufficient number of other people demanding such action – but if few others are demanding such action, an individual is not required to take to the streets and demand it all by him or herself. Participating in the early stages of an important movement is generally a supererogatory act – and participating in the later stages is often a supererogatory act as well, depending on the features of the case.

importance, because they can help us avoid inaccurate and counterproductive views about what individuals are required to about climate change and other collective action problems. At the level of individual ethics, this involves shifting efforts away from the sort of moralized advocacy for carbon footprint reductions that dominate contemporary discourse, and toward more effective advocacy for the very different sort of large-scale public policy that experts agree is necessary for a truly effective solution to the problem.[47]

[END]

[47] As William Nordhaus observes, "Suppose that [a] person proposes regulating the fuel efficiency of cars, or requiring high-efficiency lightbulbs, or subsidizing ethanol, or providing research support for solar power – but nowhere does the proposal raise the price of carbon. You should conclude that the proposal is not really serious and does not recognize the central economic message about how to slow climate change. To a first approximation, raising the price of carbon is a necessary and sufficient step for tackling global warming. The rest is at best rhetoric and may actually be harmful in inducing economic inefficiencies. … Although other measures might usefully buttress this policy, placing a near-universal and harmonized price or tax on carbon is a necessary and perhaps even a sufficient condition for reducing the future threat of global warming" (*A Question of Balance*, pg. 22 and pg. 29). See also *The Economics of Climate Change: The Stern Review*, Section IV; for example, "…a common price signal is needed across countries and sectors to ensure that emissions reductions are delivered in the most cost-effective way" (pg. 352). Broome agrees that a collective-level approach is essential for a real solution to the problem: "The needed reductions will not be achieved by the private initiatives of each of us; it will be achieved by governments using their powers of coercion over us, including their power to regulate and their power to tax. By these means, they can induce all of us together to reduce our emissions. Reductions on the required scale cannot be achieved in any other way" (*Climate Matters*, pg. 100).