

Carbon Tax, Population, and Poverty: An Act-Consequentialist Climate-Change Agenda

Ben Eggleston¹

forthcoming, after revision, in
Moral Theory and Climate Change: Ethical Perspectives on a Warming Planet,
edited by Dale E. Miller and Ben Eggleston
(under contract with Routledge, publication expected 2020)
draft of March 11, 2018
please do not cite or circulate

Abstract

This chapter presents an act-consequentialist approach to the ethics of climate change. It begins with an overview of act consequentialism, including a description of the view's principle of rightness (an act is right if and only if it maximizes the good) and a conception of the good focusing on the well-being of sentient creatures and rejecting temporal discounting. Objections to act consequentialism, and replies, are also considered. Next, the chapter briefly suggests that act consequentialism could reasonably be regarded as the default moral theory of climate change, in the sense that a broadly act-consequentialist framework often seems implicit in both scholarly and casual discussions of the ethics of climate change. The remainder of the chapter explores three topics of policy recommendations: taxing greenhouse-gas emissions (commonly called a "carbon

¹ [acknowledgements]

tax”) to discourage GHG-emitting behavior, slowing population growth to reduce the number of future agents of GHG-emitting behavior, and reducing poverty to lessen personal, familial, and community vulnerability to the harms of climate change.

Introduction

In 2009, the U.S. government implemented the “Cash for Clunkers” program, providing rebates for replacing fuel-inefficient vehicles with relatively fuel-efficient ones. The program was popular among consumers, with nearly 700,000 of them exhausting the allocated funds in a matter of weeks. There were other benefits, too, including reduced greenhouse-gas emissions, reduced emissions of traditional pollutants, and aid to automobile manufacturers who had been hit hard by the recession that had begun two years earlier. On the other hand, the program had costs, such as \$3 billion in public funds, the destruction of drivable cars worth hundreds of millions of dollars (car dealers were required to ruin the engine of every qualifying “clunker” traded in, regardless of the car’s actual functioning), and the increased greenhouse-gas emissions attributable to the increased manufacture and transport of new cars, due to program-induced demand (Lenski *et al.* 2013, 173).

To evaluate this program and other attempts to slow the pace of climate change, one option is to consider what seem to be the most salient facts and form an intuitive judgment. Another option is to adopt the perspective of a particular moral theory and apply it to problems and decisions that relate to climate change. This latter approach

could be useful both in generating a more probing and systematic assessment of the programs and policies under consideration, and in constituting an extensive testing ground for the theory itself. As an instance of this approach, this chapter is devoted to examining the ethics of climate change through the lens of act consequentialism.

Act Consequentialism

A Principle of Rightness

The eighteenth-century religious thinker John Wesley is said to have urged his followers to live their lives this way:

Do all the good you can, by all the means you can, in all the ways you can,
in all the places you can, at all the times you can, to all the people you can,
as long as ever you can. (Shafer-Landau 2018, 120)

This is a paradigmatically act-consequentialist piece of advice, since making the world as good as it can be is the basic idea of act consequentialism.

As its name suggests, the view is often articulated in terms of consequences – specifically, the consequences of acts. To be precise, any moral view that includes the following principle is a form of act consequentialism:

An act is right if and only if its consequences are at least as good as the consequences of any act the agent could have performed.

Act consequentialism is sometimes understood in terms of the *states of affairs* that result from actions: An act is right if and only if it results in the best available state of affairs. Of course, the notion of a state of affairs needs to be understood as temporally extended into the indefinite future, not as a snapshot of the world immediately after the action is performed. Minor pain followed by long-term comfort is a better state of affairs (in this temporally extended sense) than short-term comfort followed by long-term pain.

Assumptions about the Good

Act consequentialism per se is simply a thesis about the absolute and direct dependence of the right on the good: It says that acting rightly is a matter of maximizing the good. In this minimal form, this view is neutral about what the good is, and is compatible with any view about the good that one may wish to hold. A moment ago, I assumed that comfort is better than pain, but one could be an act consequentialist and hold that pain is better than comfort, or that they are equally good, or that they cannot be compared. Act consequentialism per se is just a deontic view, and can be conjoined with any axiological view whatsoever.

Moreover, in order to avoid being practically vacuous, act consequentialism *must* be conjoined with some axiological view or other, in any serious attempt to analyze a real-world ethical issue from an act-consequentialist point of view. Accordingly, in this

section I will put forward some assumptions about the good that will be operative in this chapter. These assumptions will be among the least-controversial options that a theorist could choose in order to flesh out a sufficiently determinate act-consequentialist moral theory, though I do not imagine any of them to be beyond debate.

The well-being of sentient creatures. I take it that the most important constituent of the goodness of states of affairs is the well-being of human beings and other sentient creatures. This, of course, raises the question of what constitutes a sentient creature's well-being, with the added complication that the answers might differ among types of sentient creatures. For the purposes of this chapter, we can take an ecumenical approach, stipulating that the well-being of a sentient creature (whether human or otherwise) is constituted by some combination of the pleasantness of its conscious existence, the satisfaction of its desires, and its possession of certain objectively good things such as meaningful relationships with others (though obviously the applicability of these criteria will vary among different types of creatures). In some contexts, this approach would be unacceptably indefinite. But here, its pliability averts unnecessary controversy, and the unresolved issues are likely to be moot in the context of this chapter. For example, climate-change debates rarely pivot on whether people's lives should be pleasant even if devoid of objective goods, or vice versa.

Total well-being, and its distribution. Some people who value well-being, as just proposed, believe that the greater the total quantity of well-being, the better. Although I endorse that principle myself, it is one of the more extreme views on this topic. More

popular, I think, is the idea that while total well-being matters, some facts about its distribution matter, too. For example, many people think that if well-being is distributed highly unequally in some state of affairs, then some (not too large) decrease in total well-being would, morally, be a price worth paying for some (not too small) increase in the quantity of well-being enjoyed by the worst-off. These unresolved issues deserve continued attention. As it happens, however, the arguments of this chapter will be sufficiently general as to be largely neutral with respect to them.

Goods beyond well-being. Although the axiological view sketched so far is inclusive in certain ways (e.g., by allowing that the well-being of many nonhuman animals matters, morally), it does not specifically affirm the value of anything other than well-being. One might, however, ascribe value to rational agents' having opportunities to exercise their autonomy, or to the preservation of naturally occurring biodiversity, or to other things beyond and separable from well-being. Here I take no stand on such further possible constituents of the good, other than to assume that they are, at best, much less important than well-being.

No temporal discounting. Many of the costs and benefits of climate change will occur far in the future. In some disciplines, such as traditional cost-benefit analysis, it is customary to give future costs and benefits less weight than present ones. This is done by “discounting” costs and benefits in proportion to their remoteness, typically at a rate of about 5 percent per year (Gardiner 2011, 267). For example, if a particular harm (say, a

major hurricane passing through Houston) will occur either now or twenty years from now, the latter scenario is regarded as being $(0.95)^{20}$, or 36 percent, as bad as the former.

Although such discounting is customary in certain fields, it is “the most controversial issue in climate economics” (Gardiner 2011, 267). Here I reject temporal discounting because of its confused conceptual foundations, often sloppy implementation, and ethically suspect implications. To take the last point first, any nontrivial discount rate implies that we can virtually disregard the interests of all but the nearest future generations. For example, if we were making a policy decision that could affect whether Mumbai would become uninhabitable a century from now (say, due to higher sea levels), a 5-percent discount rate would imply that, in our deliberations, we should regard that event as being less than 1 percent as bad as if it were to occur today: we should prefer the century-delayed loss of up to 100 such cities over the loss of one such city today.

Proponents of temporal discounting give many conceptual rationales for it: temporally remote events are less certain, future people will be richer, resources invested now will be worth more later, the value of money tends to decrease over time due to inflation, and others. But as Tyler Cowen and Derek Parfit point out in their sharp critique of such rationales, these considerations are all distinct from *temporal* discounting (Cowen and Parfit 1992, 145–50). For example, if a remote event is less certain, we discount it for its uncertainty, not its remoteness. Similarly, if some resource can go to a rich future person or a poor present person, we discount the former scenario because of the diminishing marginal utility of money (or because we believe benefits to the rich are worth less,

morally, than same-size benefits to the poor), not because the rich person is in the future. Regardless of the merits of discounting for greater uncertainty, greater wealth, and other factors, these are not rationales for temporal discounting as such. At best, these factors correlate with remoteness in time. But as Cowen and Parfit point out,

Remoteness in time correlates with a whole range of morally important facts. So does remoteness in space. Those to whom we have the greatest obligations, our own family, often live with us in the same building. We often live close to those to whom we have other special obligations, such as clients, pupils, or patients. ... But no one suggests that, because there are such correlations, we should adopt a spatial discount rate. ... The temporal discount rate is, we believe, as little justified. (Cowen and Parfit 1992, 159)

Conspicuously lacking from the temporal-discounting literature is the claim that temporal remoteness is itself a morally important fact – perhaps indicating the indefensibility of that thesis as well as the confusion clouding the proffered rationales.

A dogged cost-benefit analyst might acknowledge these conceptual distinctions but hold that, as a matter of implementation, remoteness in time is a good proxy for morally important facts such as greater uncertainty and greater wealth. Obviously such a claim can be assessed only in light of the details of its proposed application, but it warrants skepticism because of the many ways the alleged correlations can break down. Events are unpredictable not just because of remoteness in time but also because of other factors such as novel circumstances, poorly understood processes, and the dynamics of chaotic

and complex systems. The latter factors can reverse the usual correlation between nearness in time and predictability: for example, it is probably easier to plausibly predict the hydroelectric output of the Colorado River two generations from now than U.S. economic output two decades from now. Similarly, using temporal discounting as a proxy for rising wealth is undermined by the fact that today's extreme economic inequality shows no sign of diminishing; consequently, many of the world's future people will look back with envy on the living conditions of many of the world's present people.

The literature on discounting is vast, and here I must pass over some potentially important views, such as Cowen's argument that the rejection of temporal discounting is not merely defensible but, rather, logically entailed by the axioms implicit in any moral theory "that attempts to evaluate and compare outcomes" (Cowen 1992, 162). Here I have just aimed to say enough to explain and motivate my decision to exclude temporal discounting from the conception of the good assumed here.

Objections

Although the idea of making the world as good as it can be has obvious intuitive appeal, there are two broad objections to act consequentialism that deserve to be mentioned in this overview of the theory.

First, the theory is difficult to implement. The most obvious way of attempting to apply the theory to an ordinary decision-making situation is to identify the acts that can be

performed in that situation, ascertain the state of affairs that would result from each act, identify the best of those states of affairs, and choose the act that causes that state of affairs. But for many possible acts, the predictions that can reasonably be made are rough and conjectural, at best. For example, it is typically far from certain what will be the consequences of a high-school senior's selection of a particular college to attend, a parole board's decision to grant early release to a particular prison inmate, or a political party's decision to select a particular candidate as its nominee for president.

Second, the theory has implications that clash with what many people take to be basic truths of morality. For example, it implies that promises ought to be broken, that friends and family members ought to be betrayed, and that human rights ought to be violated, whenever doing so would produce sufficiently large benefits. Similarly, act consequentialism seems to require extreme self-sacrifice on the part of many people who, by giving up much of their personal property (including their internal organs), could benefit other people more than they would be harming themselves. Thus, act consequentialism seems to conflict with many of the prohibitions and many of the permissions of common-sense morality.

In principle, a defender of act consequentialism as a principle of rightness could conjoin it with a theory of the good specially tailored to make the resulting theory less vulnerable to the foregoing objections. For example, to address the prediction problem, one might hold that sufficiently remote effects are neither good nor bad – at least, for the purposes of evaluating acts. And to address the problem of immoral implications, one might hold

that breaking promises, betraying friends and family, and violating human rights are grave bads, while the protection and exercise of considerable autonomy with respect to personal property is a major good.

But such ad-hoc fixes are unpromising, for several reasons. First, they create new problems: the devaluing of remote effects may address the prediction problem at the cost of making the revised view even more substantively implausible than ordinary forms of temporal discounting are. And any elevating of human rights is likely to make the revised view vulnerable to new objections about its relative neglect of non-rights-related goods. (Consider, for example, how a strongly formulated right of self-ownership can preclude the taxation necessary for basic government services.) Second, such fixes frequently work at cross-purposes: consider, for example, the revised view's apparent ambivalence about temporally remote violations of human rights. Also, the more promiscuously a theory designates various goods and bads as *especially important*, the less of a difference that designation makes in the theory's comparison of alternative possible states of affairs. (When everything is special, nothing is.) In the end, it seems unlikely that these objections to act consequentialism can be avoided to any meaningful extent by way of a carefully formulated conception of the good.

There are, however, other replies to these objections that are often regarded as having some merit. In regard to the prediction problem, two points are commonly mentioned. First, act consequentialism is a theory about what makes acts right and wrong, not a decision procedure for selecting an act in any given situation. The difficulty of predicting

possible acts' consequences – which is simply a fact about the natural world and our ways of knowing about it – does not preclude or impinge on the *value* judgment that those consequences are what make those possible acts right or wrong. Second, moral perspectives other than act consequentialism (whether they are other theoretical perspectives or simply common-sense morality) almost invariably regard consequences as being important to the rightness and wrongness of acts. (They just tend to differ from act consequentialism in holding that other aspects of acts matter, too.) But the prediction problem is not seen as vitiate the consequentialist strands of those other perspectives – and rightly so, since people make prediction-based decisions constantly and in all facets of life without feeling that they are tying their fate to the output of an impossible epistemic exercise. On reflection, then, the prediction problem is hardly the special burden for act consequentialism that it might initially seem to be.

The second objection – the problem of immoral implications – can be addressed here only in the most cursory fashion, because of the variety of the topics comprised by it. Ultimately, though, most replies to instances of this objection are deployments of a few recurring argumentative strategies. One is to emphasize that act consequentialism prescribes breaking promises, violating individual rights, and other *prima facie* immoral behavior *only* when such conduct would produce *sufficiently large benefits* – a critical proviso. A second strategy is to point out that breaking promises, violating individual rights, and so on tend to have very bad consequences, not only through their direct effects in individual instances but also indirectly, by eroding background conditions of trust and security. Like bodily health, these background conditions are easy to take for granted and

may not be foremost in one's mind until they start to break down. But in any rational consequentialist assessment of various kinds of acts, it would be shortsighted not to acknowledge the grave (albeit often remote) long-term effects of breaking promises, violating individual rights, and so on. A third and final strategy for responding to objections – quite different from the two strategies just discussed – is to cast doubt on the reliability of the nonconsequentialist intuitions that the objections rely on for their dialectical force. A standard way of doing this is to argue that our moral intuitions were largely shaped by forces of natural selection that operated in an era in which most people lived in small tribes, and in which reproductive fitness was not enhanced by character traits such as being concerned about the impact of one's actions on far-away or future people, or simply being inclined to see oneself as no more important, morally, than anyone else. Of course, there is much more to be said about this and the previous strategies for responding to objections to act consequentialism, but this will conclude our overview of the theory.

The Default Moral Theory of Climate Change

Before applying act consequentialism to the ethical issues raised by climate change, it is worth noting the extent to which something like act consequentialism is already presupposed by scientists, environmental advocacy groups, public-policy professionals, and politicians in their discussions of climate change. I have to say “*something like act consequentialism*” because very few people are thoroughgoing act consequentialists. I do claim, however, that if we were to try to interpret the moral concerns animating the

global conversation about climate change as expressions of some moral theory or other, we would find act consequentialism to be a more natural fit than, say, rule consequentialism, contractualism, or Kantianism. Consider, for example, the following questions:

1. In light of probable ongoing and future climate change, what acts that we can perform will have the best consequences?
2. In light of probable ongoing and future climate change, what are the rules whose general internalization would make the world go best?
3. In light of probable ongoing and future climate change, what rules would be agreed upon by persons deliberating in an appropriately characterized initial situation?
4. In light of probable ongoing and future climate change, what maxims for action are universalizable?

If one were to survey contemporary discussions of climate change – ranging from experts’ technical treatises to informal conversations around the dinner table – one would find trains of thought far more aligned with the first question than with any of the others. Thus, it seems fair to regard act consequentialism as the *default moral theory of climate*

change, and to think that inquiries of the kind conducted in this chapter – exploring the implications of act consequentialism for climate-change ethics – would be especially pertinent to our current predicament.

To be sure, the moral perspectives underlying most discussions of climate change do deviate from act consequentialism in one major way, by focusing on preventing or lessening the *harms* of climate change rather than attending to the broader mandate of making the world as good as it can be. However, this is only to be expected in light of the fact that climate change seems to present us with many problems – many big problems – and very few opportunities. Thus, the focus on harms is most likely a matter of emphasis, arising out of the grim reality of climate change, rather than a repudiation of the idea of making the world a better place.

It might be pointed out that most people actually reject act consequentialism – for reasons having to do with the objections mentioned above – and it might be asked how a theory that most people reject could be called the default moral theory of anything (whether climate change or otherwise). It turns out, however, that those objections do not have much relevance to climate-change ethics, since it is exceedingly rare for people to have the opportunity to substantially improve the climate by, say, breaking a promise or violating someone's rights. Moreover, most non-philosophers (and many philosophers) do not apply a single moral principle to every issue they think about. Instead, they are content to use different frameworks in different contexts. And in the context of large-scale, long-term problems such as climate change, many people seem to operate with

frameworks that align fairly closely with act consequentialism. Robert Goodin persuasively argues that utilitarianism “can be a good guide to public affairs without its necessarily being the best practical guide to personal conduct” (Goodin 1995, 4). Given the similarities between utilitarianism and the well-being-focused act consequentialism presented here, Goodin’s distinction might explain why many people who decline to endorse act consequentialism *tout court* still seem to engage in essentially act-consequentialist forms of reasoning when considering the ethics of climate change.

Policy Recommendations

The remainder of this chapter will be concerned with policy recommendations that follow from the normative standpoint presented above. I will discuss just three topics, though each is a realm in which any decisions made (even by omission) are enormously consequential. The first topic, taxing GHG emissions, is standard fare in discussions of climate change, but its merits cannot be repeated often enough (until it is enacted). The other two topics, concerning population and poverty, are currently somewhat less prominent, but deserve to be at the center of climate-change ethics.

Taxing GHG Emissions

Probably the single most effective step that could be taken to responsibly reduce GHG emissions would be the implementation of a tax on such emissions – commonly called a

carbon tax, though it would apply to all major greenhouse gases in proportion to the harmfulness of their effects. Ideally, this would be done on a global scale, with the participation of all of the countries of the world. But it could also be done regionally, such as by the European Union, or just by a single country such as the United States.

Let me sketch the basic idea of a carbon tax. In theory, it would be a tax paid whenever anyone engaged in an activity that caused GHG emissions. For example, if a power plant burns coal to generate electricity, then it would pay a tax proportional to its GHG emissions. Similarly if an airline burns jet fuel to fly from point A to point B, or if a livestock firm maintains a warehouse full of cows emitting methane. The main purpose of such a tax would not be the usual purpose of taxation – producing revenue – but making GHG-emitting behavior more costly and, therefore, reducing its occurrence. (It would be a Pigovian tax, in the parlance of economics.) Such behavior changes would be expected because if it becomes more expensive for firms to use high-emissions methods of producing their output (whether electricity, transportation, beef, or whatever), then there will be substitution effects, both in production and in consumption. At the production stage, the tax will incentivize firms to switch to lower-emissions ways of producing the same product – whether by retrofitting a power plant, buying more fuel-efficient planes, or (somewhat more speculatively) making beef in a petri dish rather than in an animal. At the consumption stage, to the extent that firms end up with higher production costs and have to raise their prices, the higher prices will cause some consumers to change their purchasing decisions. Without assuming anything about the behavior of any given

individual, we would expect to see, over the whole population affected, small decreases in the consumption of electricity, air travel, and beef.

I mentioned that in theory, a carbon tax would be paid whenever anyone engaged in an activity that caused GHG emissions. Taken to its logical extreme, however, this would have absurd implications. For example, it would imply that just by being alive for another day (and emitting carbon dioxide every time I breathe), I should pay a tax, with perhaps a surcharge if my respiration needs are increased by being out of shape or exercising a lot. (So, they get me either way.) In practice, however, it would be possible to exempt such activities altogether, and focus on industrial production as in the examples above. Such simplifications would not have major distortionary effects, since not many people would, for example, exert themselves inefficiently generating their own electricity on stationary bikes because they want to avoid paying for industrially produced electricity.

Implementing a carbon tax would solve three major problems that arise from conventional environmental regulations. Currently, most developed countries have vast patchworks of environmental regulations, with distinct sets of rules for the emissions of power plants, the fuel efficiency of cars, the efficiency of air conditioners, the efficiency of light bulbs, and countless other industrial practices and consumer products – not to mention a separate patchwork of grants for purposes ranging from basic research on carbon sequestration to more insulation for old homes. One problem with this patchwork approach is that it relies on a case-by-case determination, by government agencies, of what amount of environmental benefit is worth what amount of increased costs of

production and consumption (or what amount of government funds, in the case of the grants). This approach involves relying heavily on the technical expertise and moral judgment of government agencies (many of which are headed by political appointees rather than career professionals). In contrast, the carbon-tax approach would let efficiencies in every sector of the economy emerge in response to financial incentives, and many regulations could be taken off the books. For example, it would be unnecessary for government regulators to declare inefficient light bulbs illegal, since sales of them would drop in response to consumers' increasing sensitivity to their own electricity consumption. (If a small quantity continued to be sold, to meet either special needs or quirky preferences, that might even be better, from an act-consequentialist point of view, than banning them altogether.) And power plants would be incentivized to invest in carbon-sequestration research – unless their experts thought it was a pipe dream, in which case maybe the government shouldn't be giving grants for it anyway.

A second problem with the patchwork approach is that it is almost impossible to implement with a consistent standard of how much environmental benefit is worth how much financial cost. Given the vagaries of human psychology (and the particular corruptions that typically hobble regulatory agencies, such as being captured by lobbyists), it is almost certainly the case that some environmental regulations are much more stringent than others, and that actors in some sector of the economy are being asked to make sacrifices far out of proportion to the sacrifices being required of actors in other sectors of the economy. For example, given that industries tend to have better lobbyists than consumers do, it is likely that many industrial activities are held to lower standards

of environmental concern than are implicit in many consumer-products regulations.

Replacing the patchwork approach with a carbon tax would cut off these unfair inconsistencies at their source, and replace them with the same standard for everyone.

A third problem with the patchwork approach is that, despite the enormous amount of information that it requires for its input (as I mentioned two paragraphs ago), it produces a frustrating dearth of information among its output. Specifically, it provides very little information for people to use when trying to make choices in a GHG-minimizing way. Suppose, for example, that I plan to go to Chicago, and I am deciding whether to drive, fly, or take the train. Suppose also that I care about ascertaining which option would minimize GHG emissions, either because I am otherwise indifferent among the options or because I value the environment enough to incur some extra cost for it. The existing system of regulations does not put me in a position to know anything more than that each option is not so bad for the environment that it's *illegal*. I might be able to do some research about the GHG effects of the different modes of transport and apply that research to the specific details of my trip, but that would quickly get unwieldy. In contrast, a rationally determined carbon tax would more closely align the prices of the different options with their GHG emissions. There would not be perfect alignment, since the prices would reflect the fundamental production costs as well as the carbon tax. But even that complication is desirable from an act-consequentialist point of view, because people should be discouraged from consuming a particular good or service not only in proportion to its GHG emissions, but also in proportion to its use of physical goods, the time and expertise of human beings, and so on.

One final aspect of a carbon tax that I want to discuss is what to do with the money that is collected. I mentioned above that the main goal of a carbon tax is to change behavior, not to raise money. But unless the tax is set so high as to make all GHG-emitting activities utterly cost-prohibitive – which is not a realistic possibility – there will be tax revenue, and decisions will have to be made about what to do with it. Obvious candidates include reducing the debt, reducing taxes, and funding government services. Here I want to mention another option – namely, just dividing the total revenue by the total number of adults in the country and sending each of them a rebate check for that amount at some convenient interval, such as annually. Every January 31 (or whatever), the government would announce the rebate for the previous year, and every adult would get a check for that amount.

This proposal has some similarities to one of the unique features of living in Alaska: the Permanent Fund Dividend, under which proceeds of some of the state’s mineral revenues (especially oil revenues) are invested in a state-owned fund which then generates dividends paid directly to the residents of the state. The annual amount varies, but is typically on the order of \$1,000 (Griffin 2012, 79). A carbon-tax proposal in the same vein is unorthodox, but it has several advantages over other options. First, it would blunt the criticism – sure to be mounted against any carbon tax – that despite the rationale of trying to discourage harmful behavior, it is really (just like all other taxes) just another way for bureaucrats to grab more money from hard-working citizens. Second (and relatedly), it would make the tax more politically viable to implement, and politically

more difficult to eradicate. Third (again, relatedly), it would pave the way for occasional increases in the carbon tax, which might be desirable either as part of an initial phase-in of the tax or as part of an updating of the tax level in response to scientific and economic findings indicating that the tax should be raised. Fourth – turning now from political considerations to directly moral ones – the rebate program would help low-income people cope with the financial impact of paying more for gasoline, electricity, and other necessities that would become more costly under the tax. Fifth, it would lessen economic inequality, by taxing wealthy people more (because they have higher GHG emissions) but not giving them greater rebates. Sixth, and finally, it would promote a national sense of solidarity, by representing the revenue of the carbon tax as a joint resource from which all should benefit equally.

It might be thought that if the revenue is just returned to the citizens every year, then the financial impact will be neutralized and no one will be motivated to change their behavior away from high-GHG activities. It is true that the *collective* financial impact will be neutralized, but each *individual* will still have virtually the same incentive to choose low-emissions activities over high-emissions ones. If my city told me it was going to institute this kind of rebate program to disburse the proceeds of parking tickets every year, it would still be in my interest to avoid getting parking tickets. And likewise with a carbon tax. (A few sentences ago, I had to say “*virtually* the same incentive” because when people feel wealthier – which would be an effect of the rebate program – they tend to be a little less responsive to price differences. But this “wealth effect” on consumer choice would be negligible.)

Population Policy

Whereas taxing GHG emissions is a pretty standard climate-change policy proposal (though the rebate proposal is not so standard), the next topic I want to discuss – making choices about the size of the human population – is less commonly discussed in this context. But this topic deserves careful and sustained inquiry, because of both the enormous ramifications of any decisions made in this realm and the difficult theoretical challenges that must be met in order to formulate defensible principles for making such decisions.

Most recommendations for reducing GHG emissions pertain to sacrifices and efforts that people can make throughout their lives, in greater or lesser degrees: using less heat and air conditioning, driving less, recycling more. Underlying the rationales for these lifestyle adjustments is the fact that the life of the typical human being – who chooses among the ordinarily available options for housing, transportation, time allocation, and so on – tends to make the climate worse. Now, the non-climatic moral reasons bearing on the advisability of ending the life of any already-existing person are typically so weighty that it seems unfathomable that climate-change considerations could significantly influence the rightness or wrongness of the intentional death or continued life of such a person: no one should commit murder or suicide just so there's one fewer GHG emitter walking around. But even if we regard the termination of lives as beyond the bounds of climate-change ethics, we should consider the implications of climate-change considerations for

choices about the creation of lives – specifically, prospective parents’ choices of whether to have children.

Because lifestyle adjustments (such as turning down the thermostat) only marginally reduce a human being’s lifetime GHG emissions, the choice of not having a child – or having one fewer child than one otherwise would – has effects at a whole other order of magnitude. As Dale Jamieson writes,

If an American wants to minimize his environmental impact, the most effective thing he can do is to refrain from having children. He can drive around in an SUV, hang out at McDonald’s, take long hot showers and still have much less environmental impact than if he fathers one, good, green, nature-loving American child. (Jamieson 2008, 189)

A recent study quantifies the impacts of about a dozen different GHG-reducing actions more precisely. In this study, perhaps the most striking statement is that “a US family who chooses to have one fewer child would provide the same level of emissions reductions as 684 teenagers who choose to adopt comprehensive recycling for the rest of their lives” (Wynes and Nicholas 2017a, p. 3). But to be honest, although the number 684 is impressive, I’m underwhelmed by the choice of recycling as a reference point, and focusing on a U.S. couple inflates the figure since the U.S. has the highest per-capita GHG emissions in the world. For me, the most meaningful comparison derivable from the study’s computations is that if a couple in the *developed world* (not just the U.S.) were to decide, at some point in time, to have one fewer child, then the GHG-reducing effect of that decision would be about *nine times* as large as the *combined* GHG-reducing

effect of their doing the following *for the rest of their lives*: adopting comprehensive recycling, not eating meat, buying green energy, getting rid of their two cars, and taking two fewer transatlantic flights per year (Wynes and Nicholas 2017b, 6).

Thus, having fewer (or no) children is an extraordinarily effective way for people to reduce GHG emissions. But, of course, we cannot infer from this the moral permissibility of any general constraint on the number of children that people can have, such as China's one-child policy. Unlike the temperature at which one's thermostat is set, or the price one has to pay for gasoline, the number of children one has is among the most life-defining decisions that most people will ever make. Whether one wants many children, or just one, or none at all, any external interference with one's chosen plan would be felt by almost anyone to be an intolerable imposition on what should be a private and autonomous choice. Accordingly, state interference in this realm is generally regarded as a gross violation of human rights – even when it actually does have beneficial GHG effects. On this point, Dale Jamieson strikes an appropriately rueful tone in his observation that “the undeniable fact is that the three main factors that have reduced GHG emissions [since the early 1990s] are, in increasing importance: global recession, the collapse of communism, and China's one child policy” (Jamieson 2014, 59).

Limits on family size are a blunt instrument, and there are markedly less disturbing ways of slowing population growth. For example, countries experiencing rampant population growth can make contraception (and related medical expertise) more widely available to their citizens. Also, the international community can encourage or pressure countries in

which there is widespread sex discrimination to give women the same schooling, employment opportunities, and legal rights that are provided to men. In addition to being required by basic norms of social justice, such reforms often result in reduced family sizes as women become more financially independent as young adults, marry later, and have greater professional aspirations – all of which typically result in correspondingly reduced parenting aspirations. Unfortunately, the impact of population policies such as these will be limited by the fact that high-population-growth countries also tend to have relatively low standards of living and correspondingly low per-capita GHG emissions – meaning that they are the places in the world where births are least bad for the climate. But births anywhere are bad for the climate, so any reduction helps.

I characterized the policies just discussed as “markedly less disturbing” than coercive measures such as China’s one-child policy. However, a fully comprehensive ethical assessment of even non-coercive population-reduction policies would have to tackle what may be the most fundamental question of population ethics: Do more people – assuming they have lives worth living – make the world a better place? Within this chapter’s act-consequentialist framework, there is an obvious route to an affirmative answer to that question: If there are more people with lives worth living, then there is more well-being, and the world is a better place. On this view, when new people come into existence, there may be disvalue because of GHG-related considerations, but there may also be value because of the individual’s personal well-being – and the latter might exceed the former. And when that is the case, any given population-reduction policy might make the world better for existing people, but worse overall. At this point, most readers of this chapter

will recognize the foregoing thoughts as just the opening lines of what could be a book on population ethics. Thus, perhaps the clearest conclusion that we can draw here, on the topic of population policy in climate-change ethics, is that resolving its conundrums is likely to be as difficult as it is important.

Poverty Reduction

The last topic I want to discuss – poverty reduction – is rarely regarded as a priority of climate-change ethics. Of course, poverty reduction has been a major topic of the broader field of applied ethics at least since the 1972 publication of Peter Singer’s article “Famine, Affluence, and Morality.” But it is not typically a focal point of discussion for people concerned about the effects of climate change. Indeed, one can easily get the impression, from contemporary philosophical and popular discussions, that poverty reduction and climate change are two separate areas of applied ethics. Here, however, I argue that reducing poverty should be seen as one of the central imperatives of climate-change ethics.

The standard narrative in climate-change ethics is that we are spoiling the planet for our children and grandchildren. Stated a little more explicitly, the idea is that ever since the industrial revolution or so, our atmosphere has been on a dangerous warming trajectory. For a while – maybe the entire nineteenth century or so – people could be excused for not knowing about the greenhouse effects of burning coal and other fossil fuels. But scientific knowledge of the greenhouse effect (and public dissemination of that knowledge)

advanced tremendously during the twentieth century, and it has been at least several decades since we last had excusable ignorance of the consequences of our actions. Yet we continue to increase our GHG emissions every year, wrongly leaving our descendants an atmosphere much worse than the one we've enjoyed.

Intergenerational justice is certainly an important dimension of climate-change ethics. But the story is not as simple as it might seem. One complication – though a very fortunate one overall – is that in many ways the world we will pass along to our descendants will be better than the one we inherited. There are many resources other than the natural environment that can contribute to people's well-being, including infrastructure such as wells, sewers, roads, and ports; well-functioning institutions such as legal systems, financial systems, and school systems; useful professions such as medical science, engineering, information technology, and journalism; cultural artifacts such as television, film, and literature; and morals and norms that restrain unethical behavior.

Given these many dimensions along which current people can help or hurt future people, it seems likely that the requirements of intergenerational justice have more to do with the totality of them than with any single one of them such as the temperature of the atmosphere. And along many of these dimensions, future people are likely to be better off than current people. Consider, for example, the improvements of recent decades in fields such as medicine, transportation, and telecommunications – we live longer, fly cheaper, and communicate more easily than even our recent forebears. And there is little reason to think this progress is about to come to a halt. Consider, too, that we have an international

order that has not seen a world war for more than seventy years even though the interval between the first world war and the second was only twenty-one years. Finally, consider the progress that has been made, over the last several decades and centuries, in curbing and eradicating unethical practices such as discrimination on the basis of race, sex, sexual orientation, and (more recently) gender identity. Although moral progress is often unsteady – sometimes being “two steps forward, one step back” – it seems likely that future people will inherit an ethos of greater respect and inclusiveness than was experienced by previous generations. All things considered, then, the concept of intergenerational justice does not deliver the unequivocal verdict that the standard narrative often claims.

A second factor complicating the intergenerational-justice narrative is that climate change will not affect all future people the same, or even similarly. Even if it is correct to say that the climate is getting worse for humanity overall, this generalization masks two further facts that must not be overlooked. First, some people will benefit from climate change. Some residential real estate will become more valuable, and likewise for some commercial real estate, ranging from ports to tourism destinations. Some farmland will become capable of growing more valuable crops. And some people will have skills and expertise that will be in greater demand. Second, among people who are harmed by climate change, some people will be harmed a lot less than others. A corporate executive who owns a Miami Beach condominium that loses 20 percent of its resale value will probably have his lifestyle affected far less than a subsistence farmer in Indonesia who owns a small house that starts flooding from time to time. The condominium owner

might have a greater dollar-value loss, but he can probably adapt to the changed circumstances with less effect on his well-being than the owner of the small house that soon might not even be habitable.

These facts not only show the inadequacy of the intergenerational-justice narrative, but also point the way toward a response to climate change that goes beyond just trying to lessen it altogether (though I do support that, as indicated by my discussion of a carbon tax). If we had perfect foresight and frictionless channels for reallocating resources, then it would make sense to simply identify and aid the people who will be harmed the most by climate change. Unfortunately, given the knowledge and institutions that we have, such a fine-grained approach is impractical. But there is a feasible second-best approach, suggested by the example of the Miami Beach executive and the Indonesian farmer. We can expect that, to considerable extent, the poorer a person is, the more his or her well-being will be damaged by climate change. This means that providing aid to poor people is, in effect, a way of lessening the harms of climate change. And it is a much more targeted way of doing that than just trying to lessen climate change altogether.

In response to this argument, it is natural to think that if we were to just prevent climate change, then we wouldn't have to worry about finding exact or approximate ways of targeting the people adversely affected by it. However, climate change is basically unavoidable. Even if global GHG emissions were to drop sharply today, global average temperature would still continue to rise, due to the GHG accumulation already present in the atmosphere. Moreover, given current trends in major GHG-emitting countries such as

the United States and China, the prospects for even modest reductions in global GHG emissions are dim. Thus, absent a technological miracle that could simply reverse past and future GHG emissions (such as massively scalable carbon-sequestration technology), we must acknowledge that significant climate change is going to occur, and that decisions we make now can put people in better or worse positions to cope with it.

Let me turn to two additional reasons favoring an increased focus on poverty reduction in this context. First, although we should be increasing our current GHG-reduction efforts, we must also keep in mind that each additional degree of cooling (or non-warming) will be more costly to achieve than the previous one: there is “low-hanging fruit” that we can grab relatively easily, so that we might achieve 1 degree of cooling pretty inexpensively, but the second 1-degree increment will be more costly to achieve, with the third 1-degree increment being even more costly, and so on. (For a homely analogy, consider that when trying to lose weight, losing a few pounds might be relative easy, losing another few pounds is typically more difficult, losing another few pounds is even more difficult, and so on.) In other words, money spent on GHG-reduction efforts has decreasing marginal effectiveness, in terms of quantity of resulting GHG reductions and temperature reductions. Consequently, there will come a point at which we can more effectively lessen the harms of climate change not by plowing more money into preventing it, but by letting it happen – to some extent – and using our resources to help the most vulnerable people adapt to it, whether by buying a home on higher ground or paying higher prices for energy or food.

A second reason favoring increased poverty-reduction efforts is that greater wealth would not only make poor people more capable of adapting to climate change, but would also help them deflect the impact of other calamities they may face. In addition to the financial burdens that poor people already face, they may face, at any time, any of a Pandora's list of crises that could affect everyone in a certain region or the entire planet, such as fallout from nuclear weapons and dirty bombs (either in conventional warfare or terrorism), chemical and biological weapons such as weaponized smallpox, accidental epidemics such as the 1918–20 influenza that killed 50 million people, abrupt disruptions to the food supply, failures of utility distribution systems (e.g., water, electricity, or natural gas) due to accidental or malicious disruptions to their control systems, failures of other communications systems (personal phones, the global banking system, etc.), and non-climate-based natural disasters such as earthquakes, tsunamis, volcanic eruptions, and meteor strikes. For people confronting the consequences of any of these events, it probably will not matter very much whether the atmosphere is a couple of degrees warmer or cooler than it might have otherwise been. But it probably will matter very much whether they can afford to relocate, or pay for medical care, or repair their homes in some way, or just not default on their mortgages because of unemployment stemming from the economic disruption that many of the above events would cause. Money cannot buy everything, but it can buy things that people need in a wide range of life-altering circumstances – with climate change being, unfortunately, just one of them.

I mentioned at the beginning of this discussion that poverty reduction and climate change are often seen as two separate issues. On this widely held view, any resources committed

to one cause are thereby denied to the other, just as any dollar donated to the local symphony is thereby denied to the local animal shelter. But as I have argued here, poverty reduction and climate change overlap to a considerable extent. Since poverty makes people more vulnerable to being harmed by climate change, the threat of climate change makes poverty reduction an even greater priority than it has long been recognized to be. Thus, poverty reduction, along with more commonly discussed policies such as a carbon tax, should be seen as urgently required by the growing threat of climate change.

Bibliography

- Cowen, Tyler. 1992. "Consequentialism Implies a Zero Rate of Intergenerational Discount." In *Justice between Age Groups and Generations*, edited by Peter Laslett and James S. Fishkin, 162–68. New Haven: Yale University Press.
- and Derek Parfit. 1992. "Against the Social Discount Rate." In *Justice between Age Groups and Generations*, edited by Peter Laslett and James S. Fishkin, 144–61. New Haven: Yale University Press.
- Gardiner, Stephen M. 2011. *A Perfect Moral Storm*. Oxford: Oxford University Press.
- Goodin, Robert E. 1995. *Utilitarianism as a Public Philosophy*. Cambridge: Cambridge University Press.
- Griffin, Jr., Christopher L. 2012. "The Alaska Permanent Fund Dividend and Membership in the State's Political Community." *Alaska Law Review* 29, no. 1 (June): 79–92.
- Jamieson, Dale. 2008. *Ethics and the Environment*. Cambridge: Cambridge University Press.
- . 2014. *Reason in a Dark Time: Why the Struggle against Climate Change Failed – and What It Means for Our Future*. Oxford: Oxford University Press.
- Lenski, Shoshannah M., Gregory A. Keoleian, and Michael R. Moore. 2013. *Ecological Economics* 96 (December): 173–80.
- Shafer-Landau, Russ. 2018. *The Fundamentals of Ethics*, 4th edition. Oxford: Oxford University Press.
- Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1, no. 3 (Spring): 229–43.

Wynes, Seth, and Kimberly A. Nicholas. 2017a. “The climate mitigation gap: education and government recommendations miss the most effective individual actions.”

Environmental Research Letters 12, no. 7 (July): 1–9.

<https://doi.org/10.1088/1748-9326/aa7541>.

Wynes, Seth, and Kimberly A. Nicholas. 2017b. Supplementary materials for “The climate mitigation gap: education and government recommendations miss the most effective individual actions.” <http://iopscience.iop.org/article/10.1088/1748-9326/aa7541/data>.